

The Personal Audio Loop: Designing a Ubiquitous Audio-Based Memory Aid

Gillian R. Hayes, Shwetak N. Patel, Khai N. Truong, Giovanni Iachello,
Julie A. Kientz, Rob Farmer, and Gregory D. Abowd

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

Abstract. Sound is an important medium in our lives, but its ephemeral nature can be problematic when people cannot recall something they heard in the past. Motivated by everyday conversational breakdowns, we present the design of a continuous, near-term audio buffering application: the Personal Audio Loop (PAL). PAL was designed as a truly ubiquitous service to recover audio content from a person's recent past. Initial brainstorming and prototyping for PAL revealed major aspects of the design space that require further investigation, including potential usefulness in everyday life, the level of ubiquity required, the usability features for any instantiation of the service, and the social and legal considerations for potential deployment. We present a design of PAL, informed by a controlled laboratory study, diary study, and examination of pertinent legislation. We conclude with an analysis of the results and some initial observations of the deployment of a prototype developed for a Motorola i730 handset.

1 Introduction

Everyday conversations fill our lives, and we are all very familiar with the kinds of breakdowns suggested by these simple scenarios:

- You are in a conversation with a friend, and one of you is interrupted. When the conversation resumes, neither of you can remember what you were talking about.
- You are at a social event, and you are introduced to someone new. Minutes later, you have forgotten the person's name.

We have a particular interest in automated capture of live experiences for later access, and we are naturally drawn to these scenarios, because they demonstrate the use of audio capture with near-term access. Over the past three years, we have experimented with different technical approaches, and have found that a mixed technological and human-centered approach is necessary to produce a near-term (*i.e.*, less than one day) audio service that would be likely to survive a real deployment. Such a design must answer questions of human significance pertaining to the following issues:

- *Usefulness*: Though motivated by observations from everyday life, how often and in what situations do people actually need a near-term audio memory aid?
- *Ubiquity*: What parameters of such a service would make it available everywhere and every time someone needed it?

- *Usability*: How should the service deliver functionality to maximize its benefit and minimize its distraction?
- *Social and legal considerations*: What social and legal concerns might prevent the successful deployment of an audio recording application for everyday life?

An automatic audio-based memory aid is arguably outside of the realm of a typical person's experience. Therefore, potential users should be able to interact with a working prototype to have a sense of the capabilities, necessitating the answering of engineering questions as well, including important architectural considerations.

From a technical perspective, there are several options for designing an audio-based memory aid to provide the capability motivated by the above examples. Although all designs reflect the same basic notion of replaying a buffer of recently recorded audio, early prototypes varied in terms of distribution of recording and playback capabilities. A fully distributed system assumes an instrumented environment, with microphones, speakers and interface controls placed to maximize opportunities for recording and playback wherever and whenever needed. A fully localized solution provides recording and playback in an all-in-one package carried wherever needed. A hybrid solution might allocate the recording in the environment and accomplish playback through a handheld device that receives streamed audio from a central repository.

In this paper, we present a design study of the Personal Audio Loop (PAL), a solution for a deployed near-term audio reminding service that addresses both the technical concerns of an interesting capture and access application while also answering questions from the four categories described above. The process involved a series of formative studies that led to the design of a self-contained service integrated into a commercial mobile phone handset. Although the decision to build a local solution for PAL came fairly early, it results naturally from an exploration of the usefulness, ubiquity and socio-legal concerns for this problem, and it is justified by our findings.

In the next section, we provide a brief background of technology and of relevant social and legal work in this area. In Section 3, we describe the initial implementation of PAL on a commercial mobile phone handset and outline the various empirical and diary studies that formed the basis for our formative studies. In Section 4, we give preliminary results from an initial deployment study and in Section 5 we summarize the critical design features of PAL. Finally, in Section 6 we summarize the contributions of this work and outline future work.

2 Background and Related Work

Near-term capture and access applications that provide audio reminder services have been previously explored in the office as well as for 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) [9]. Although the system was designed for use in a setting where social protocol allows recording, the authors recognized the privacy issues of subsequent use of archived recordings, and suggested that social expectations change with use. In MERL's real-time audio buffering technique, captured audio persists for the duration of that phone conversation [4]. During the course of the conversation, a user may tap

the phone against the ear to move backwards in the audio and to replay any portion of the discussion. This system does not store conversations and could arguably pass legislative tests and be socially acceptable. Video has been employed for reminder services as well: the *Déjà vu Display* (previously known as the *Cook's Collage*) explores the use of collage displays to show recent activities in the kitchen [13]. In the case of a memory lapse or interruption, the user can rely on annotated snapshots of key steps to remind her of the last few things she did. Although the *Déjà vu Display* was designed for private space (*i.e.* the home), much attention was given to specific privacy-friendly affordances, such as the camera angle, the richness of captured data and avoiding sound recording. Such affordances are determinant in the equilibrium of privacy and will be extensively discussed below.

Legal cases over the past two decades have exposed the contrasting requirements and balances of privacy and utility for recording applications. We draw from the experience in the fields of surveillance in public spaces and of the privacy of private communications.¹ Among other sources we considered European Directive 95/46/EC, [6] together with opinions and rulings by various EU Data Protection Authorities (DPAs) [5, 2] and several US Supreme Court² rulings—the most relevant being the *Katz v. United States* [10] case, which extended the right of privacy to what the individual seeks to protect from the public and the *Kyllo v. United States* [11] case, which indicated that the subjects of surveillance are granted a sufficient expectation of privacy if the surveillance technology employed is not in common use.³

Despite the ongoing debate stressing the differences between the United States and Europe regarding privacy, legislation regulating the recording of communications by electronic means is remarkably similar. The main items are the US Electronic Communications Privacy Act (ECPA) of 1986 [14] and European Directives 2002/58/EC [7] and 95/46/EC. ECPA regulates wiretap and surveillance and applies to any electronic recording device and conversations (“oral communication”) between two persons “exhibiting an expectation that such communication is not subject to interception,” even if the conversations were not transmitted through a telecommunications network. European Directive 2002/58/EC covers only personal conversations transmitted over public telecommunication networks. However, Directive 95/46/EC applies to any personally identifiable information, which includes recorded voice conversations, according to multiple opinions by European national data protection authorities. Although the Directive was originally meant to regulate the management of personal data collected by organizations in large textual databases, recent opinions expressed by DPAs have addressed cases of more limited balancing of individuals’ rights. As is detailed below, Directive 95/46/EC requires a proportionality assessment between potential harm and benefits; however, the personal character of the application might exempt users from many provisions, including informed consent.

¹ Most industrialized nations have pertinent legislation; we limit our inquiry to the US Federal legislation and European Union directives. Note that these laws are not directly comparable: US legislation gives states less discretion than EU law gives to member states.

² The United States do not have DPAs specifically appointed to examine privacy issues.

³ Further information on the details of these and other US Supreme Court decisions can be found at <http://www.findlaw.com/casecode/supreme.html>.

3 Formative Studies of PAL

Based on early interviews and our intuition, we determined that the platform for PAL would need to be mobile, powerful both in processing and development environment, include buttons, and an external or attachable microphone. The mobility, ubiquity and performance of mobile phones make them an appealing platform for this application, but only certain phones support the required capabilities. Our choice, the Motorola iDEN i730 (Fig. 1) is a clamshell phone featuring a J2ME programming environment conforming to the MIDP 2.0 and Mobile Media APIs. The i730 microphone is capable of recording voices in a small room with the phone open or closed in a shirt pocket or attached to a belt, with higher quality than most PDAs. The two formative studies reported were designed to answer questions of the feasibility of using a mobile phone as the interface to an audio-based memory aid and to characterize the frequency and situations of use in everyday life.

3.1 Laboratory Study: Developing a Usable Phone Interface

In its normal operating mode, our implementation of PAL continuously records audio from the user's environment. Audio older than the buffer length (in our initial prototype, 15 minutes) is automatically deleted. Recording automatically halts when the user answers or makes a call. Five buttons are available on the outside to accommodate interactions while the phone is closed (Fig. 1). PAL provides simple audio navigation features (*e.g.* rewind), informed both by previous research on skimming [1] as well as by commercial video recording services like Tivo™. PAL includes a simple timeline visualization on the exterior LCD of the handset indicating application status (recording, playback and direction of navigation) as well as the playback position in the audio buffer relative to the current time (the right edge of the timeline).

We designed a laboratory study using an early prototype to test the usability of the interface from a quantitative performance perspective and a qualitative impression.




Fig. 1. The Motorola i730 handset used for PAL. Three buttons control navigation and record/playback mode. A timeline indicates mode and relative place in the buffer.

Method. The laboratory study included 18 participants (students and faculty from our institution specializing in HCI research, 5 female, 13 male, ages ranging from 18 to 50). Participants with an HCI background were explicitly chosen with the intent of examining heuristics such as the mapping of buttons to functionality and the quality of the visualization. Participants' experience with mobile phones ranged from seven years of consistent use to no experience at all (7 participants). We demonstrated the prototype, encouraging participants to examine the device and ask questions until they expressed comfort with its functions.

PAL's intended use involves the replay of audio for which the user was present initially. The controlled study, designed to mimic this scenario, included a scripted dialog of five minutes. In this script, the participants asked researchers predetermined questions, and researchers replied with the same answer for every participant. The script purposely involved a large amount of detail to increase the likelihood that participants could not recall the answers to all questions from memory. After completing the dialog, the researchers who had been participating in the dialog removed the script and asked the participants a series of questions about the information they had just been provided. Although it was noted whether the participants remembered the information without use of PAL, every participant was asked to find and play every answer. Participants were encouraged to "think aloud" as they used the prototype, and the researchers timed how long it took an individual to find the answer, theorizing that this first time use while discussing their actions would be a worst case timing for most users. Participants answered seven questions, the first two being practice questions not used for computing timing results. An exit survey and semi-structured interview provided a qualitative evaluation of the interface and of their need for this kind of service.

Results. After a short demonstration, all participants were able to navigate the audio well enough to answer our questions. They commented that the device was easy to use with one hand ($\mu = 6.95$, $\sigma = 0.2$, 7 being the highest), and small enough to carry at all times ($\mu = 5.42$, $\sigma = 2.0$ out of 7). They could clearly understand the audio even in its highly compressed form ($\mu = 6.5$, $\sigma = 0.9$, with 7 being "strongly agree").

With an audio buffer of 15 minutes, participants required an average of 34.8 seconds ($\sigma = 22.58$) to find responses for questions that were known to be in the recorded audio while talking aloud about their actions. Participants reported the visualization was somewhat helpful in accomplishing the tasks, but not overwhelmingly so ($\mu = 5.21$, $\sigma = 1.4$, with 7 being "very helpful"). Thirteen out of our eighteen participants used PAL without the visualization, preferring an eyes-free interaction.

Although inquiring about privacy was not a goal of this study, ten of our participants raised spontaneous concerns regarding the social acceptability of a continuously recording system. The most common sentiment expressed indicated that participants were less concerned about recording their own voice than their conversation partners'.

3.2 Diary Study: Determining the Usefulness of PAL

The laboratory study showed the feasibility and usability of PAL on a mobile phone, but it did not inform us about the overall usefulness in everyday life. We undertook a

diary study to explore the extent to which a near-term audio reminder service was needed, looking for frequency and characteristics of potential use. Diary studies balance the ecological validity of gathering such data *in situ* against interruption of everyday activity flow caused by recording personal observations, particularly in mobile settings [3]. We asked for specific information relating to social context including privacy concerns in the diary entries and during the follow-up interviews.

Method. Twelve experienced mobile phone users (5 female, 7 male, ranging in age from 22 to 60 years) participated in the study. Participants' occupations spanned a spectrum of domains, including a psychologist, finance manager, realtor, car dealer, consultant, professor, and full-time homemaker. We demonstrated a fully working version of PAL to participants. We then asked them to carry small pocket-sized diary and record an entry in it for each incident during the following week when they would have needed or liked to use the PAL service. Each page of the diary contained a simple form to complete for the potential instance of use, streamlined after an initial trial period. Each form in the diary included space for describing the content of the audio to retrieve, when and where the incident occurred and whether any persons unrelated to the conversation were nearby. Participants also estimated how far in the past the salient audio content was and rated how important it was to retrieve that information. Fig. 2 shows an example of an incident survey.

At the end of each week, we collected the diaries from participants and conducted semi-structured interviews to examine in detail up to six diary entries per participant per week, including privacy-related questions such as the kind of information being sought, the distance of unrelated third parties from the participant and their assessment of the social appropriateness of using the device in the specific context. We then gave each participant who chose to continue for another week a new empty diary to again record incidents. At the end of the study, we conducted semi-structured interviews with all participants. The weekly and summary interviews allowed us to clarify misunderstandings in the entries as well as to probe particular issues, such as privacy concerns, that were not easily gathered in the chosen diary form factor.

Results. Twelve people participated in the first week, eleven of them continued for the second, and eight in the third, for a total of 31 participant weeks and 109 incident reports. Participants reported an average of 3.5 ($\sigma = 2.7$) incidents per week, of which 32% referred to audio from "less than 10 minutes ago", 26% from "10 minutes up to an hour", while only 6% were from over a day prior.

Of the incidents reported, 25% occurred in public, 44% in semi-public spaces (defined as schools, workplaces, *etc.*) and the remaining 31% in private space (predominantly car and

Day of Week	S	M	T	W	T	F	S	
Time of Day	12 am - 6 am	6 am - 9 am	9 am - 11 am	11 am - 1 pm	1 pm - 6 pm	6 pm - 9 pm	9 pm - 12 pm	
Location	Kroger						Unrelated People Present?	<input checked="" type="checkbox"/>
1) What was the audio you needed? 2) Why?								
While shopping I ran into an old friend. We talked for a while and he gave me his phone number. I could not remember it later.								
How long ago was the audio of interest?	< 10 min	10 min - 1 hour	1 hour - 1 day	> 1 day				
How problematic was it to forget this?	1 Neutral	2	3 An annoyance	4	5 Problematic			

Fig. 2. Sample diary entry.

home). In 44% of the incidents, participants indicated that people unrelated to the audio they wished to retrieve (*e.g.*, other customers in a restaurant) had been present during the time they would have liked to record. We collected follow-up information for 83 incidents during the weekly interviews. Participants asserted that they would not have felt rude towards their communication partner using PAL in 52 of these. During the second and third weeks participants were questioned about their reactions had their partners objected to their use of the application. Participants stated that such an objection would be “not likely” in 24 of 26 incidents queried and indicated that they would not have complied with the objection, had there been one, in 19 of the 26 incidents queried. Only in 4 occasions participants asserted that unrelated bystanders could have been concerned had they known that they were using PAL. When asked how far away they would like PAL to record, 67% chose within a small room (10 feet), 22% preferred smaller areas (own voice or arm-length distance), and only one individual requested a large radius, reporting that he “is just nosy”.

During interviews, participants reported on how long they would be willing to search for content rated at various levels of problematic. If they were “neutral” (scoring a one on a five point scale) about the content, they reported being willing to spend an average of 336 seconds ($\sigma = 172$) to search, whereas if the audio content was of vital importance (scoring a five), they reported being willing to spend a minimum of 15 minutes with three users responding “however long it takes” to retrieve it.

4 Preliminary Results from Deployment

We deployed a working version of the application to four of the diary study participants for seven weeks. Although we do not report on their use, four members of our research team have also been using PAL for over two months. During the first four days of the deployment, we asked participants to carry a diary to note their uses of the device. These participants used the device on average 2.5 times per week ($\sigma = 1.9$, pro-rated given the short term of the study). Although this average is lower than what was indicated by the diary study, participants also reported on average 1.5 incidents that they thought about using the device and chose not to ($\sigma = 0.6$). In one case, the user’s conversation partner recovered the information before the user was able to try with PAL. In all other cases, the reason not to use PAL was reported as forgetting it was available. Informal interviews with the users since this initial probe indicate that ordinary use subsequently remained fairly consistent with the rate observed in the first four days, and that the frequency of use for exploring the application or showing it to others has decreased substantially. Overall, satisfaction as reported through qualitative interviews has been high. All four users requested to continue using the devices after the first four days and reported that they believed they would use them more over time. Each user changed the buffer length (ranging from ten minutes to sixty), the initial jump backward (ranging from 15 seconds to 60), or both. Users expressed that configuring the application was important and one user even indicated that he changes the buffer length depending on the situation he is about to encounter.

By deploying the devices to even a small number of users, we expected to be able to observe uses both expected and emergent and gain greater understanding about the

dependency users might have developed on the service. In the initial four day probe, the most frequent reported situation for use was to remember forgotten details (60%). Other unexpected situations have also been reported in the following weeks. Specifically, users have been employing PAL as an instructional aid: recording conversations with customers and then replaying them for an employee in training. One user has also been using it as a medical journal to record data about symptoms requested by her doctor by speaking them aloud when she can not write them down at that moment. Users have begun to expect the service to be available, reporting that they choose not to write information down when it is already spoken aloud.

Social contract issues recurred more often than the results of the diary study had revealed: users expressed that conversation partners aware of the device sometimes responded negatively initially, but relaxed after the application and its buffering and discarding functions were explained. Interestingly, all four users reported informing new conversation partners about PAL less frequently as time went by. After several weeks, users have almost stopped alerting conversation partners altogether. As frequently as users reported negative social repercussions from PAL, they also reported positive cooperative uses of the device. For example, one user's wife consistently uses PAL on his device by walking near to him and speaking when she needs to remember something. We are exploring in depth the changing behaviors of individuals around the owner as the study continues.

5 Critical Features for Use

Informed by the exploration of privacy regulations and by findings from the laboratory and diary studies, we uncovered the critical features of PAL outlined previously.

Making PAL useful. Given the rates of 2.5 and 3 incidents per week as reported by the deployment and the diary study, the need for PAL is justified. Analysis of the stated purpose for recovering the audio provided additional information, synthesized in Table 1. From the legal perspective, the frequency and unpredictability of use of this application could support a positive argument for the proportionality test (as used in [5]) with regards to the issue of continuous automatic recording.

Table 1. Purpose for recovering audio (total 109 entries in diary study)

Purpose category	Occurrences
Forgotten previous details (<i>e.g.</i> , making a list, retrieving details)	36 (33%)
Replaying for conversation partner (replaying for person who either spoke the audio originally or was present to hear it)	20 (18%)
Interrupted (external activity took focus away from important audio)	18 (17%)
Explicit tape recorder behavior (participant was aware prior to the incident that she wanted to record it)	13 (12%)
Distracted (another concurrent activity took attention)	13 (12%)
Relaying information from one partner to another (replaying for person not present when original audio was recorded)	9 (8%)

Information minimization requires collecting the minimum amount of personal information needed by the application. Given that 58% of the diary incidents referred to content within one hour, a buffer up to 60 minutes should suffice, with a 15 minute default. EU and US law diverge in this regard, as ECPA does not make any distinction based on stored information retention time. A more conservative way of looking at this issue would be that of understanding the duration of the “social contract”, implicit among parties engaged in a conversation, to determine how long a recording can be maintained after the end of such conversation. This measure relates to the relation between distance and place (in the sense of [8]): how long does it take to move between places with incompatible social contracts? Because PAL could be abused when crossing place boundaries, the recording should be limited to minimize such risks. While valid from a phenomenological standpoint, we decided to postpone this assessment, given the unsolved issue of gathering reliable contextual data.

Making PAL ubiquitous. As discussed in Section 3, we targeted a mobile/wearable solution for PAL. Our intuition was that the mobile phone would likely be with an individual most of the time (at least during working hours, perhaps also at home). Of the participants in the laboratory study who owned a mobile phone, all but one was carrying it upon arrival for the study. Furthermore, in 79% of the diary entries queried, the participant’s mobile phone was on her or within reach.

The results of both studies demonstrated the need for and appropriateness of this service to be wearable, as opposed to environmental. The argument can be made that an audio buffering service in the environment might be preferable for a variety of reasons, including power concerns, better audio quality, and the convenience of users not needing to wear a device. Every participant reported, however, that there are times when it would not be possible for the service to be environmental. Every participant who recorded any entries recorded at least one at a public place or outdoors, where environmental solutions would be difficult. Participants also expressed control concerns about an environmental version of PAL versus a wearable solution. One participant noted, “[I would] rather have the control of it being on my person.”

While advocating a wearable solution, however, participants were not interested in a completely separate device but instead as a “value added features” to the mobile phone already owned and carried. Although this may seem obvious in retrospect, it implies the fairly strict requirements that PAL must run unattended on the mobile handset, without recharging for at least a day, and it must not interfere with the call functions of the phone. These requirements are met by our currently deployed prototype, resulting in an arguably ubiquitous service.

Making PAL usable. Our final prototype provides asymmetric backward/forward skip features over the recording, with default values of 10 and 5 seconds, respectively. While most participants of the laboratory study liked these defaults, the values can be adjusted, and anecdotal experience shows that individuals do optimize them. We did not observe effective use of fast forward or rewind skipping features during the laboratory study. Considering the limited capabilities of the handset, we opted to support earmarks instead. The user can set earmarks and can use the backward/forward skip buttons to traverse these earmarks or to simply navigate without using them.

One issue identified in the laboratory study related to the mapping of the pair of navigation buttons: there is no “natural association” between the buttons and backward and forward navigation. This issue is exacerbated by the variety of ways the handset can be mounted on a belt or carried in the pocket or purse. We opted for a “never-wrong” mapping. When transitioning from record to playback, the only possible direction of navigation is backward in time. Therefore, whichever navigation button the user first presses is mapped to backward navigation; the other button is used for forward navigation. Once recording resumes, the previous mapping is cancelled.

Making PAL socially and legally acceptable. We do not endorse the common opinion that people necessarily must adapt to technological evolution by changing their social expectations. However, a case could be made that PAL does not impinge on constitutional rights and that, in the long term, practice could show the harmlessness of this application, granted specific guarantees, namely, small recording radius, short buffer length and some form of notification to the conversation partners. We would like to stress that it is not in the scope of this paper to provide conclusive legal opinions – a task best left to courts and DPAs. Our purpose is to provide a balanced, if necessarily concise, overview of PAL’s social and legal impact.

A number of different stakeholders can be identified with regards to PAL; we consider three: the user, conversation partners and unrelated third parties. Considering the third category, diary results indicate that 69% of the entries related to recordings in public or semi-public spaces, and 44% stated that other, unrelated, people were present. These figures support our concern with third-party privacy, which contrasts with the fact that the vast majority of our participants were neither preoccupied with a third party’s privacy nor with that of the conversation partner. These observations are particularly interesting because they diverge from legislation in force. ECPA does prohibit capturing a third party’s conversation when the owner of the device is not part of that conversation and the conversation takes place with reasonable expectation that it is not being intercepted (*e.g.*, non-public space). On the other hand, it must be noted that the perceptual properties of sound might not grant *constitutional* basis (in the US) for an expectation of privacy in public space, as suggested among others by numerous cases adopting the “plain view” rule. This could allow adapting surveillance legislation to permit limited memory aid devices such as PAL.

Interface affordances and information retention policies greatly impact social acceptability. Altering the coverage of the microphone is an essential factor of a proportionality determination, as suggested by analogous DPA opinions involving personal uses of video surveillance (namely, outdoor camera units at home entrances) [2]. Likewise, DPAs have used retention time and deletion policies to evaluate the social impact of surveillance applications. Completely eliminating the risk of recording third parties’ conversations is extremely difficult, given the characteristics of sound transmission, but the retention properties of this application do support the claim that PAL does not serve archival purposes, nor does it vastly facilitate surveillance, since the device is carried around by its user; if concealed or left unattended, the application arguably presents lower risks than traditional audio recorders.

In the relationship with conversation partners, informed consent is one fundamental tool of social action, embodied in privacy law. Its implementation presents though formidable technical and usability challenges. In our case, anecdotal evidence col-

lected during the deployment suggests that our participants have, over time, renounced to preventively explain or ask permission to use the service. At times participants turn off the device due to social pressure. Both observations support our previous findings from the diary study. This could hint at a gradual adaptation to the technology, and the adoption of appropriate social behavior, similarly to what is currently happening with camera phones.

Directive 95/46/EC exempts the personal use of information (*e.g.*, diary) from the informed consent requirements, and the figures reported above regarding control and usage seem to confirm that users view PAL as a preeminently personal application. When asked about objections by conversation partners, one participant answered “I wouldn’t care. It’s a tool for me.” Moreover, all-party informed consent would place an unreasonable burden on the user (a condition which may exempt from the consent requirement). Still, it is not guaranteed that this application would qualify as personal, nor that the Directive’s provisions, thought for textual diaries and address books would transfer in DPAs’ judgment to environmental recording. If not so, DPAs have expressed the need for explicit notification and consent. ECPA provides in the general case the “one-party consent” rule, in which informed consent by conversation partners is not necessary if the user of the recording device takes part in the conversation, without prejudice on the legality of the subsequent use of that information. ECPA acts only as a baseline, however, and many states have introduced various additional safeguards, such as two-party consent and notification cues such as “recorder beeps” (a useful, non-authoritative comparison of US state laws can be found in [12]).

Although we did not receive strong feedback from our participants requesting that PAL provide a notification cue while recording, in view of the above considerations, we decided to incorporate such function in the deployed handsets. When recording, the outer LED integrated in the round ornament on the phone shell (see Fig. 1) lights up red. During playback the light turns green. Although recording is usually associated with a red indicator, we are aware that people might not understand its meaning and that users could obviously conceal the LED as well as the recording device: the user remains ultimately responsible for abiding to the social contract and mores.

Concluding, the legality of PAL in parts of the US with stronger safeguards appears to be more problematic than in Europe because of the greater flexibility granted by EU law to DPA judgment. The lack of precedents and novelty of this recording without archiving do not allow us, however, to reach any definitive conclusion. In any case, characterizing PAL as a memory aid and *not* as a recording device appears to be the juncture through which any argument in favor of social and legal acceptability must flow.

6 Conclusions and Future Work

Based on controlled and field studies of use of a mobile audio-based memory aid, we conclude that not only is the service desirable for users, but also that its implementation on a mobile phone is possible and usable. Users can find the information needed in less time than they reported being willing to spend. They need this service at least once a week, and they are willing to wear a mobile phone at all times to have access to it. Our analysis shows that this application falls within a legal “grey area”, and that we cannot definitively assert or deny its legality. The interface and retention charac-

teristics of the application, along with observation of initial deployment suggest that the application might be socially acceptable. We have deployed PAL on the Motorola i730 platform and plan to report on a long-term study of the emergent uses PAL inspires and on the social contract and mores it influences.

Acknowledgements. Thanks to Motorola Corporation, and specifically Joe Dvorak of the iDEN group, for supplying the i730 handsets used for these studies. This work is supported in part by NSF ITR grant 0121661, the NSF GRFP, the Sam Nunn Security Program and Intel Corporation. All human subjects research conducted for this paper is covered under IRB approved protocols H03221 and H03217 at the Georgia Institute of Technology.

References

1. Arons, B.: SpeechSkimmer: Interactively Skimming Recorded. In: Proceedings of the Sixth Annual Symposium on User Interface Software and Technology (UIST '93). ACM Press, New York (1993) 187–196
2. British Institute of International and Comparative Law: The implementation of Directive 95/46/EC to the Processing of Sound and Image Data. (2003) <http://www.europa.eu.int>
3. Czerwinski M., Horvitz E., Wilhite S.: A diary study of task switching and interruptions. In: Proceedings of ACM Conference on Human Factors in Computing Systems (CHI 2004). ACM Press, New York (2004)
4. Deitz, P., Yerazunis, W.: Real-Time Audio Buffering for Telephone Applications. In: Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology (UIST 2001). ACM Press, New York (2001) 193–194
5. European Commission Art.29 Working Party: Opinion 4/2004 on the Processing of personal data by means of video surveillance, 11750/02/ENWP89 (2004) <http://www.europa.eu.int>
6. European Union: Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data (Data protection directive)
7. European Union: Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications)
8. Harrison, S., Dourish, P.: Re-place-ing space: The roles of space and place in collaborative systems. In: Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work (CSCW'96). ACM Press, New York (1996) 67–76
9. Hindus, D., Schmandt, C.: Ubiquitous Audio: Capturing Spontaneous Collaboration. In: Proceedings of the 1992 ACM Conference on Computer Supported Cooperative Work (CSCW'92). ACM Press, New York (1992) 210–217
10. *Katz v. United States*, 389 U.S. 347; 88 S. Ct. 507 (1967)
11. *Kyllo v. United States* 121 S. Ct. 2038; 150 L. Ed. 2d 94; 2001 U.S. LEXIS 4487 (2001).
12. Radio-Television News Directors Association: Hidden Cameras, Hidden Microphones: At the Crossroads of Journalism, Ethics and the Law (1998) <http://www.rtnfd.org/>
13. Tran Q.T., Mynatt, E.D.: What Was I Cooking? Towards Déjà Vu Displays of Everyday Memory. GVU Technical Report 03-33. Georgia Institute of Technology (2003)
14. United States Electronic Communications Privacy Act of 1986, 18 USC 2510