Can you see what I hear? The Design and Evaluation of a Peripheral Sound Display for the Deaf

F. Wai-ling Ho-Ching, Jennifer Mankoff, James A. Landay
Group for User Interface Research, Computer Science Division
University of California at Berkeley
{wai-ling, jmankoff, landay}@cs.berkeley.edu

ABSTRACT

We developed two visual displays for providing awareness of environmental audio to deaf individuals. Based on fieldwork with deaf and hearing participants, we focused on supporting awareness of non-speech audio sounds such as ringing phones and knocking in a work environment. Unlike past work, our designs support both monitoring and notification of sounds, support discovery of new sounds, and do not require *a priori* knowledge of sounds to be detected. Our Spectrograph design shows pitch and amplitude, while our Positional Ripples design shows amplitude and location of sounds. A controlled experiment involving deaf participants found neither display to be significantly distracting. However, users preferred the Positional Ripples display and found that display easier to monitor (notification sounds were detected with 90% success in a laboratory setting). The Spectrograph display also supported successful detection in most cases, and was well received when deployed in the field.

Keywords

Assistive technology, deaf, sound visualization, peripheral and ambient and notification displays, non-speech audio

INTRODUCTION

People use sound in many subtle ways to gain awareness of the state of the world around them. For example, sounds of officemates working provide awareness of whether or not one is alone in an office. At home, one might hear the television at night and guess a household member is up late. Similarly, many everyday devices such as doorbells, ovens, and telephones use sound to communicate events.

In 1997 there were 3.4 million Americans with difficulty hearing, of which 227,000 were deaf [16]. The deaf use awareness techniques such as vibration and flashing lights to gain awareness of environmental sounds. However, a large gap remains between the experiences of deaf and hearing individuals. Previous work on assistive technology for the deaf has concentrated mostly on verbal communication, including supporting the automatic

translation of American Sign Language (ASL) [6, 13], but little research has addressed awareness of non-speech sound (defined here to include background speech for which no content analysis or speech recognition is done). This paper describes the design and evaluation of a peripheral visual display of sound to help close that gap.

As hearing people designing for the deaf, the problem of designer bias becomes acute since we are designing for a population very different than ourselves. Although some work has been done outlining guidelines for running usability studies for the disabled [3,14], there is little guidance available for the challenges encountered during the early design stages. Our approach was to employ a number of exploratory design techniques with both hearing and deaf participants during the early design of our display.

We explored the ways in which hearing people use sound in their everyday lives. We analyzed existing techniques employed by the deaf for sound awareness and interviewed deaf participants to discover where their current techniques fell short and how we could improve their experience with new tools. Conceptual drawings of sound by hearing participants and exploration with a deaf participant using paper prototypes formed the basis for our current design.

We implemented two prototypes, one of which, a spectrograph, represents pitch and amplitude of sound over time. The other depicts position and amplitude over time. We evaluated two prototypes in a dual task experiment with eight deaf participants and found that participants were able to peripherally identify notification sounds with both systems while performing a visual primary task. Participants had significantly higher identification rates with the visualization that represented position. Neither visualization resulted in a significant amount of distraction in terms of performance of the primary task.

Although we only used eight participants, our results indicate that we are able to provide the deaf with awareness of sound visually while they perform a primary task with fixed visual focus. This work has been received with much enthusiasm by members of the deaf community and may ultimately result in better support for sound awareness for the deaf in situations of fixed visual focus.

Overview

We begin with a discussion of related work, then present the results of interviews and explorations with hearing and deaf participants indicating what sounds and visualizations

Technique	Application	Characteristics	Cons
Vibration sensing	Awareness of sounds that create vibration, e.g. sensing footsteps, feeling that a computer is on	Does not require focus of attention Supports some ambient and some notification sounds	Depends on infrastructure (e.g., having hardwood floors)
Flashing lights	Awareness of telephones, doorbells (Notification sounds)	Supports notification sounds	Fixed visual attention Must hook up each device
Hearing Dogs	Awareness of all sounds	Supports notification sounds only	Requires ongoing maintenance Requires *a priori* training per sound
Visual Inspection	Multiple applications, e.g., Steam for a kettle, Looking out the window for the arrival of a guest	Sometimes is the only alternative Some ambient, some notification	Different for each sound Polling rather than interrupt based
Hearing Aids, Cochlear Implants	Enhancing existing hearing, but not at the fidelity of "normal" hearing	Enhances awareness of all sounds	Requires training for interpretation of sounds Results vary by case

Table 1: Taxonomy of Existing Sound Awareness Techniques for the Deaf.

are of interest, and where current techniques fail. This is followed by descriptions of the two peripheral displays that we built. Our evaluation comparing those displays, presented next, suggests a preference for positional information over pitch. We end with a discussion of the qualitative experiences of our users, followed by our conclusions and plans for future work.

RELATED WORK

Research related to this project falls into two categories: tools for *visualizing audio*, and tools for supporting monitoring and notification of *peripheral information*. Additionally, a range of *sound awareness tools* is available to the deaf, and we present a taxonomy of these tools in Table 1. Most research on assistive technology for the deaf has focused on support for communication between deaf and hearing people, including automatic recognition of sign language using computer vision techniques [13], and translation of spoken language into text signed by an avatar [6]. Sign language is not used in our work.

Visualizing Audio: Tools for visualizing speech have reached the mainstream in the domain of speech therapy [7]. These tools traditionally employ waveforms and spectrographs as visualizations, and are targeted at expert users. Additionally, they are not meant to display non-speech sounds, or to support monitoring and notification.

Monitoring and Notification: Displays that support monitoring of peripheral information without grabbing the users attention are often referred to as peripheral or ambient displays [4,10,17,18]. For example, the *"dangling string,"* an art installation, was a string that spun at a speed proportional to the current network load [17]. These displays lie on the boundary between background and foreground awareness. They must notify the user of interesting information without impeding the performance of a primary task, a goal that requires careful design [15].

To be successful, our system must support *monitoring* by displaying information continuously without distracting the user, unless a sound is of particular interest. In this case, it *should* capture the user's attention, and we refer to this as *notification*. Successful peripheral displays exist in both of these categories. For example, TimeAura supports

continuous monitoring of the passage of time by a lecturer [10], while the Scope system notifies users about communication events of interest [4].

Sound Awareness Tools: There are a variety of techniques currently used by the deaf for monitoring and notification of sounds, summarized in Table 1. The techniques shown were chosen based on interviews with deaf participants, an assistive technology consultant and ASL interpreters.

We can classify the techniques based on whether they support notification or monitoring. A technique such as visual inspection that supports only monitoring requires constant attention if important sounds are not to be missed, while techniques such as flashing lights or hearing dogs that support only notification leave out many potentially interesting events. Vibration sensing supports both notification and monitoring, but cannot provide complete information. Hearing aids and cochlear implants provide the most complete solution but may not work for all people.

These techniques also vary in the initial investment and the amount of ongoing maintenance required for them to work. For example, vibration sensing requires an uncarpeted floor, but has no other cost or maintenance requirements. A hearing dog requires up-front training and ongoing care. A flashing system light requires every phone or doorbell in a house to be connected separately and the light must be visible in different rooms. Each change requires additional time and effort.

In conclusion, there is still a gap between the sound experience of a hearing person and the experience of a deaf person. For example, although there are several methods used to provide awareness of certain notification sounds, there is little effective support for monitoring. Also, many of these techniques require that the sounds monitored be known *a priori*. This is not practical in the case of unexpected or new sounds.

FIELDWORK

We carried out a series of field studies to learn what sounds were most important and what visualizations were most intuitive to our users. Because many office sounds are currently not accessible to the deaf with existing tools, we worked with both hearing and deaf participants to answer
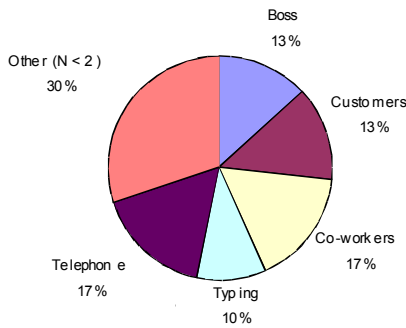
Figure 1: Useful sounds in the workplace as reported by hearing participants.
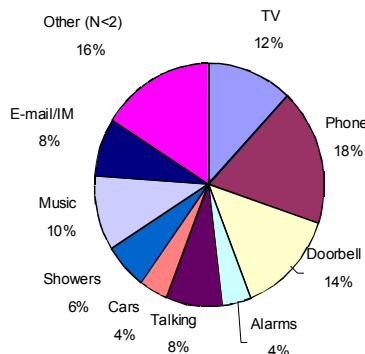

Figure 2: Useful sounds in the home as reported by hearing participants.

the following questions: What tools do the deaf use to perceive sound? What sounds do hearing people find important? Where do deaf people believe current techniques fall short? What metaphors do hearing people use when thinking about sounds? We present the results of our inquiry with deaf and hearing users below.

Use of Sound in the Home and at Work

We surveyed ten hearing adults about the sounds that were important to them at home and at the workplace. Many of the responses were speech related (*e.g.,* Boss, Customers, Talking). Figures 1 and 2 summarize their responses. Both speech and non-speech sounds were reported in each case. Speech sounds such as public announcements and background conversations are an important issue for the deaf. However, we chose to ignore speech content in this work. The technology necessary to do speech recognition in unconstrained settings is not far enough along to support a device for monitoring speech sounds.

In our analysis of the survey, we concentrated on non-speech sounds mentioned by at least two people. At home, there are nine commonly mentioned sounds, of which eight are non-speech sounds. In the office scenario, only two non-speech sounds were mentioned by multiple people. Both monitoring and notification were mentioned in each case.

Monitored sounds provided the participants with a general sense of what was happening in a room. Examples of these sounds include background chatter and typing. In contrast, notification sounds required attention or action. Examples are the telephone ring or the smoke alarm.
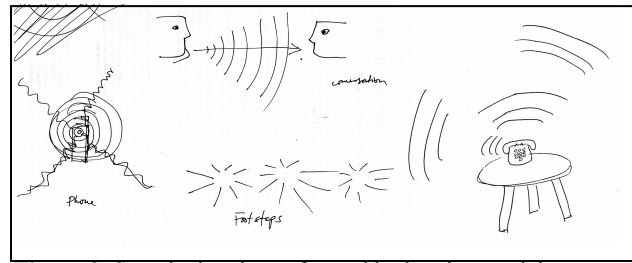

Figure 3: Sample drawings of sound by hearing participants. Note how the object that causes a sound is often included in the drawing. Also, sound is often represented as waves traveling through space.

The Needs of the Deaf

We interviewed ten deaf participants and an assistive technology consultant to learn about issues they felt were not adequately supported by current techniques.

Awareness of the presence of others

Sound provides awareness of the activity and presence of others. Having a sense of when they were alone was important for the participants we interviewed. Participants mentioned that, for this reason, they wanted to be aware of sounds such as the radio or music.

Interaction with sound based appliances

Many appliances such as ovens, kettles, microwave ovens, and smoke alarms are explicitly designed to use sound notification. Others, such as faucets that drip when left on, and printers that stop making noise when printing is complete have implicit sound cues that communicate their state. Awareness of the state of such appliances is limited for the deaf.

Leaving the Home Environment

A deaf person has more control over their home environment than of more public environments such as the workplace. For example, they can choose a home with wooden floors or invest in a light system to hook up to the doorbell. However, these systems are not commonly in place when the deaf person works amidst hearing co-workers. In a few cases the participants worked in an environment where the majority of the office workers were deaf. Their office space was designed so that everyone could see anyone at the doors and phones and appliances were fitted with lights. However, such participants were the minority among those we interviewed.

Intuitive Visualization of Sound

Previous work has focused on information displays of sound for experts. We wanted to explore how non-experts intuitively visualized sound.

An exploratory exercise was performed where six hearing undergraduate and two hearing graduate students were asked to draw a visual representation of three sounds: footsteps, a telephone ringing and a conversation. A similar technique has been used previously to study a user's conception of privacy mechanisms [8].
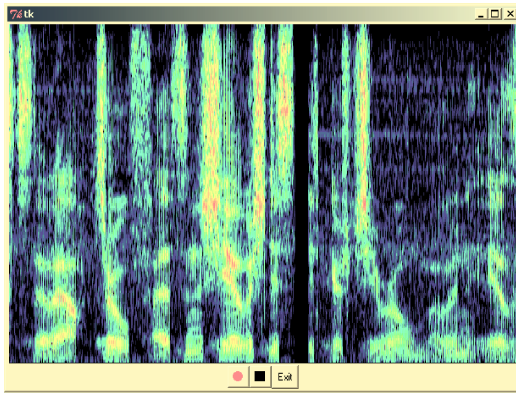
Figure 4: Speech visualized by the Spectrograph prototype. *X*-axis = time (lines move from right to left); line height on *Y*-axis = pitch; Color = volume (bright = loud).

Figure 5: Speech shown in the positional Ripples prototype. It is identifiable by its position in the room. The speech is coming from the desk on the left. For legibility, we have modified the colors in this and other Ripples figures: The original background was black, with the map drawn in white lines.

Figure 3 shows examples of the drawings that were created. 65% of illustrations included the object that caused the sound. This suggests that identification of the source of the sound is important. Also, 59% of drawings depicted sound as ripples or waves that travel through space.

TWO PROTOTYPES

Based on these interviews, we decided to create an everyday device that would provide awareness of monitored sounds and notification sounds in the office in conjunction with office work at the desk. Our goal was to support both monitoring and notification of sounds, as well as discovery of sounds of which the participant was not previously aware.

We explored tactile and visual modalities with the help of a deaf participant, but based on his feedback and the results of our fieldwork we decided to focus on visual designs because of the increased amount of information that could be displayed. We developed two visual display designs in depth (see Figures 4 and 5). The first was based on the spectrograph, a mapping that is commonly used as a focal display by experts for visualizing speech in speech therapy. The second was based on the metaphor of sound as ripples derived from the drawings of our hearing participants (Figure 3). Below, we introduce each prototype, and discuss how they were designed to support both monitoring and notification of sounds, a key contribution of our work.

In the *Spectrograph* visualization, height is mapped to pitch with high frequency sounds appearing higher on the graph than low sounds. Color is mapped to intensity in the following order from loud to quiet: red, yellow, green, blue. The x-axis represents time and the visualization animates from right to left. Mechanical sounds are easily identifiable using this visualization technique as they often have regular pitch and amplitude patterns. Additionally, hardware requirements are minimal (a simple microphone) and setup is simple.

In the *Positional Ripples* visualization, the background displays an overhead map of a room. Sounds are depicted as rings. The center of the rings denotes the position of the
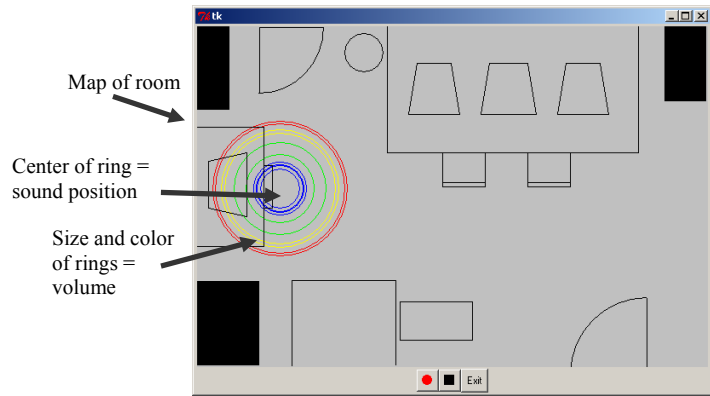
sound source in the room. The size of the rings represents the amplitude of the loudest pitch at a particular point in time. Each ring persists for three seconds before disappearing. This visualization carries higher equipment and set-up costs, and gives information about position missing from the previous visualization.

Monitoring and Notification

Monitoring and notification occur naturally in our awareness of background sound (an example is the "cocktail effect" which has been attributed to spatial disparity and pitch [1]). Our visualization exercises with hearing participants (Figure 3) confirmed the importance of position in sound identification. We built on these ideas in our visual displays of sound, highlighting pitch in the spectrograph visualization, and location in the ripples visualization.

A highly audible event should map to a highly visual event in support of notification. Similarly, when there is little or no sound in a room, we would like the visualization to remain calm. Previous work has suggested that animation is a more powerful visual signal than color change or shape change [2]. Also Maglio and Campbell have suggested the use of discrete motion over continuous motion to reduce distraction [9]. Therefore, in our designs, we devised a mapping where constant sounds, such as the hum of conversation in a room produces slow constant motion, but audio signals such as the start and stop of a phone ring map to a more dramatic animation as the sudden appearance and disappearance of bright visual elements. Additionally, the appearance of a sound source in a new position in the ripples visualization creates a strong visual signal as a mass of rings suddenly appears in a new location.

Implementation

The prototypes were implemented in Python. The SNACK v2.2a1 toolkit [12] was used to manipulate real-time sound data. The user interface widgets were implemented using Tkinter, a wrapper around Tcl/Tk for use with Python.

Figure 6: The experimental setup.

Figure 7: The primary task.

Both systems displayed sounds read from a live microphone attached to the local PC. However, we did not implement the determination of position for a sound source needed by the Positional Ripples visualization. Although techniques exist for determining the position of a sound source using microphone arrays [1, 10], it was not the focus of our research. Instead, we used Wizard of Oz methods to simulate positional detection.

EXPERIMENT

We ran a controlled experiment with both prototypes to examine the effect of different factors on a person's ability to detect notification sounds. Participants were asked to perform a primary task on a primary display, while trying to monitor the peripheral display for notification sounds. They were instructed to imagine that they were working on a computer while waiting for a phone call or for a visitor who would knock at the door. While waiting, sometimes the office would be quiet or noisy with background chatter. We chose a phone call and door knock based on suggestions by participants in our interviews, and because they represent common, highly noticeable sounds.

Participants and Apparatus

There were eight deaf participants, five male and three female. All were adult office workers. None had non-corrected visual impairments, such as color blindness. All were profoundly deaf. One regularly wore a hearing aid, but removed it for the session. All had computer experience through their work. Instructions for the study were available in PowerPoint slides, which the participants could browse through themselves. Participants had the choice of communicating with the researchers through an ASL interpreter or with pencil and paper.

The study took place in a quiet room resembling an office. Two 17-inch LCD displays were connected to a single PC running the study software. Participants sat approximately

60 cm from the primary display, with the secondary display off to the side (Figure 6).

Procedure and Design

The study had a 2x2x2 within subjects factorial design with a dual attention task based on past experiments on performance using peripheral displays [2, 1, 15]. We manipulated three independent variables: prototype [Spectrograph, Ripples], background noise level [Quiet, Noisy] and notification sound [Door Knock, Phone Ring]. The background noise was speech, produced by a random selection from five different speech files. These occurred at random time intervals between one and five seconds. Each speech file was about two seconds long, so there was often the case of concurrent speech generation. All sounds were recorded at approximately the same volume. The notification sounds (door knock and phone ring) were also recorded for the experiment. This allowed us to reduce confounds from unwanted noise in the environment.

The primary task (Figure 7) was an attention-demanding visual task where participants searched a screen of buttons, each containing a number between 0 and 9. Participants were asked to click on as many 0's as possible. Once clicked, a 0 changes into a 1. A new number field was randomly generated for each trial.

In a concurrent, secondary task, participants were told to report any occurrence of either notification sound by pressing the <ESC> key on the keyboard. A dialog box would then appear on the primary display that would ask which sound they had detected (Door Knock or Phone Ring) and their certainty on a 5 point scale (1=Unsure, 5=Very Sure). Example notification sounds are illustrated in Figure 8, as they would appear in each display

Sounds were presented to the participants in a series of blocks of trials. Before starting a block, a participant was given instructions and training for as long as it took to become comfortable with the task. A single notification sound was produced in a random interval between 5 seconds and 60 seconds after the start of each trial. The participant was given 30 seconds to respond to the notification sound. After 30 seconds, or if the participant correctly detected the notification, a new trial was started.

The participants began with two baseline blocks of 6 trials each. In the first, they performed the primary task without any sounds being played. In the second, they attempted to detect sounds without the visualization. After the baseline

Figure 8: The secondary task was to peripherally distinguish between the phone ring and the door knock. The same sounds represented on different displays are shown for comparison. A primary task was performed concurrently.

a. Phone ring, Ripples **b.** Phone Ring, Spectrograph **c.** Door Knock, Ripples **d.** Door Knock, Spectrograph
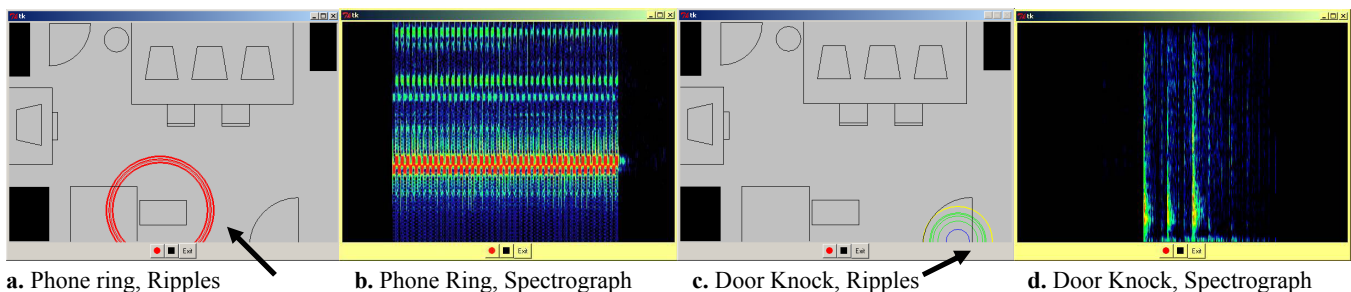
blocks, we ran two blocks of fourteen trials with each of the two prototypes. Two trials in every block of fourteen did not contain a notification (14%) to measure expectation. They were not included in the analyses presented here. The order in which the prototypes were presented was counterbalanced. Individual trials within a session were presented in a random order. Trials were participant initiated. There was a background questionnaire distributed at the start of the session and a post-experiment questionnaire distributed at the end. Data was collected by video, questionnaires, computer log files and interviews.

Results and Analysis

We compared the two visualizations in terms of individual preference, correct detection of signal sounds, distraction from the primary task, and learning. All eight participants expressed a preference for the Positional Ripples visualization in the post-experiment questionnaire.

Correct Detection

To compare how the visualizations support notification of events, we performed a preliminary analysis of the percentage of correct detections of notification sounds.
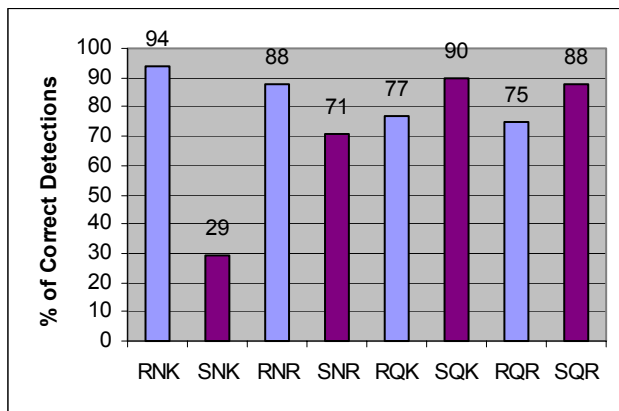

Figure 9: Percentage of Correct Detections for all combinations of factors. RNK = Ripples / Noisy / Knock. SQR = Spectrograph / Quiet / Ring.

We graphed the percentage of correct detections across all experimental conditions (Figure 9). The highest detection rate appeared in the case with the Ripples prototype with a Noisy background and with Knock as the target sound. The worst case occurs with the Spectrograph also with the Noisy background with Knock as the target sound.

Another observation is that the Spectrograph performs better than the Ripples during the quiet conditions with an average mean of 89% compared to 76% for Ripples during the quiet conditions. This difference was shown to be significant with a two-tailed paired samples t-test ($p < 0.05$). However, the case is reversed during noisy conditions with Ripples performing significantly better ($p < 0.05$) at 91% and the Spectrograph at 50%.

The Ripples had a significantly higher overall detection rate of 83% compared with 69% for the Spectrograph. A two tailed related pairs t-test showed that this difference was

significant ($p < 0.05$). In the t-test calculation for each participant, we used the average of the 24 trials run with each prototype (excluding the four no signal trials).

Distraction

Distraction was measured by counting the rate at which a participant could select 0's in the primary task. Note that our results measure distraction due to "noise" during the noisy conditions. They do not measure distraction due to the notification sounds themselves since the participants had to interact with the interface when they detected a notification sound and this interaction time would skew the results. In a given trial, we measured the zero selection rate starting at the beginning of a trial and ending at the moment when the notification sound began to play. Therefore any degradation in performance that occurred after the notification sound had been played was not included in our analysis.

The performance of the primary task was compared across three conditions: monitoring of the Spectrograph prototype, monitoring of the Ripples prototype, and absence of a secondary task as a control. A one-way repeated measures ANOVA did not reveal any significant differences between any of the three conditions ($F_{obs} = 1.152$, $p > 0.05$). Using our definition of distraction, this means that there was no detected difference between distraction without a secondary task and distraction with it. Also notable is that there was no detected difference between distraction when using either the Spectrograph or Ripples prototypes. In these analyses, there were unequal sample sizes because two participants did not complete the baseline block so we removed them from this analysis.

Analysis of the questionnaire suggested that the rings visualization was perceived as less distracting than the Spectrograph. Although this difference was not significant ($p > 0.05$), visual inspection of the data showed that seven of the eight participants reported the Ripples visualization as less distracting.

Learning

Results from the questionnaire reported that the Ripples visualization was easier to learn than the Spectrograph. Each individual rated the Ripples visualization higher than the Spectrograph visualization with means of 5.43 and 3.57 respectively on a 7-point scale (1=Difficult to learn, 7=Easy to learn) respectively. This difference was significant ($p < 0.05$). There does not appear to be large learning effects during the trials. There were two blocks of trials per system and the order of the blocks was not a significant factor towards the percentage of correct detection according to a 2-way repeated measures ANOVA ($F_{obs} = 1.795$, $p > 0.05$).

Discussion

The low value for SNK in Figure 9 can be explained by qualitative observations. The Spectrograph did not perform well for temporally adjacent sounds. When a knock was played right after noise, participants had difficulty visually distinguishing the door knock from the background noise.

However in the Ripples visualization, position disambiguated these cases. In our experiment, amplitude + position was a more powerful combination for disambiguating the signal sounds from the background noise than amplitude + pitch. Sometimes sounds that are easily distinguishable aurally were difficult to distinguish on the basis of amplitude, pitch and time alone.

The Spectrograph performed better than the Ripples visualization in the quiet conditions. Our observation from the experiment is that the Ripples visualization did not produce a large enough visual change to gain the attention of the participant. Thickening the rings to produce a more dramatic visual effect may help address this problem.

Two participants wanted to move the visualization closer to the center of their primary visual focus. Transparency and judicious use of screen borders could be used to create a single-screen design in the future.

QUALITATIVE OBSERVATIONS

Although the experiment uncovered results about the prototype's ability to support notification, there was still something missing from our data. We wanted to see how exploration with the prototypes allowed participants to discover sound information, which they could not detect previously. We had the opportunity to observe more informal interaction of one deaf participant with the Spectrograph prototype during a one-week field study and interaction of two participants with both prototypes during pilot sessions.

Field Study

We deployed the Spectrograph visualization for a week with a deaf graduate student. He had participated in the previous experiment. The visualization was deployed with an IBM Thinkpad laptop computer with a 13-inch screen. Sound input was received through the built in laptop microphone. He was instructed on how to adjust the sensitivity of the microphone and encouraged to adjust the height and location of the display to be most comfortable. The office in which it was deployed was fairly quiet. Although there were other offices nearby, he was the only occupant of his office at the time of the field study. The observations were self reported by the participant.

We learned that robustness is extremely important for a display that is meant to run all the time. Although the prototypes never crashed during the experiments, when forced to run continuously for a period of a week, crashes occurred making practical use difficult. Additionally, the use of a low-quality, built-in laptop microphone limited his ability to distinguish the sounds of the computer from the more distant sounds such as someone knocking on the door and a kettle whistling in the next room.

Despite this, our participant was able to detect a number of sounds including: speech, mobile phone calls, chair movement, typing, mouse movement, page turning, papers rustling, footsteps of people entering the office and a university truck which turned around outside several times.

When alone, he learned most of the sounds by inference.

> "If I see a truck outside and see the display show a sound pattern that tracks the truck's movement closely, I infer that there must be a linkage"

However, there was also paired training that developed around our prototype. We found that hearing coworkers at the office would spontaneously experiment with the prototype and explore sounds visually with the deaf participant.

> "When I first set up the display, I had to explain how it worked to my fellow officemates in the other room. They were fascinated and wanted to learn more about the sounds that could be detected – they liked watching the patterns of the sounds. So we were testing the display by knocking on the office door… speaking, setting the mobile phone to ring. The display performed excellently on these counts – it could show the sounds."

Pilot Study

Participants during our pilot sessions had time for informal interaction with the prototype. Also, during these sessions, the prototypes read environmental information live. This provided us with the flexibility to explore ambient sound with our participants.

One participant who had been deaf since birth noticed a sound signature in the Spectrograph visualization, "What's that?" he asked. When we pointed it out that it was the sound of his chair creaking, he burst out laughing and did a dance in his chair to see it creak. He had worked in that office for years, but had never known his chair creaked.

The joy that our participants gained from interaction with our prototype was not collected by any of our quantitative methods, but was best summarized in the words of one of our participants, "This is great! … I'm learning to hear again after 30 years!"

CONCLUSIONS AND FUTURE WORK

We have developed two assistive visual displays for the deaf to provide them with awareness of background sounds. Based on our fieldwork, we focused on supporting awareness of non-speech audio sounds in a work environment. Unlike past work, our designs support both monitoring and notification of sounds, support discovery of new sounds, and do not require *a priori* knowledge of sounds to be detected.

The Spectrograph display, which shows amplitude and pitch information, is an inexpensive design that could work with little initial investment or overhead. It met our goal of discovery even when deployed with a built-in laptop microphone, and supported detection at over 70% for three of our four experimental conditions in a lab setting. The Positional Ripples display, which requires a larger initial and ongoing investment, provides information about position but does not show pitch. It was not tested in the field, but performed at approximately 90% in a laboratory setting. Neither display distracted our participants measurably from their primary task.

The Positional Ripples display provided valuable information about the importance of position to participants in the context of our laboratory study. However, its dependence on a working microphone array for identifying location makes it prohibitively expensive for the everyday user. We are currently working on outfitting a computer laboratory with a network of sensors that detect sound amplitude for other ubiquitous computing research, and we plan to take advantage of this to further test the Positional Ripples display. Perhaps in future ubiquitous computing environments, such technology will be standard. In the meantime, our experiments show that the less expensive Spectrograph is feasible today.

Although our field deployment of the Spectrograph display was successful, especially in supporting discovery of new sounds, room for improvement remains. The built-in microphone made it difficult to distinguish the sounds of the computer from the more distant sounds that interested our participant more, such as the door knocking and kettle whistling. A better quality wireless microphone could provide the flexibility to monitor a specific device such as the kettle, or be placed further from the user to capture more sounds outside of his visual field. In the future, we would like to study long-term use of this system in the field. We also plan to improve the extent to which our visualizations are embedded in the user's visual environment and explore further visual and tactile designs.

ACKNOWLEDGMENTS

REFERENCES

1. Arons, B. A review of the cocktail effect. *Journal of the American Voice I/O Society* **12** (Jul. 1992), pp. 35-50.

2. Bartram, L., *et al.* Moving Icons: Detection and distraction. In *Proceedings of Interact 2001*, pp. 157-166, 2001.

3. Coyne K. P. and Nielsen J. How to conduct usability evaluations for accessibility: Methodology guidelines for testing websites and intranets with users who use assistive technology. Nielsen Norman Group, Oct, 2001, http://www.NNgroup.com/reports/accessibility/testing.

4. Dantzich, M., *et al.* Scope: Providing awareness of multiple notifications at a glance. In *Proceedings of AVI 2002*, 2002.

5. Dutta, S., *et al.* Evaluating benefits and distractions of animated secondary displays for attention-centric primary tasks. Regular research report in *Proc. of the 2002 International Conference on Imaging Science, Systems, and Technology*, 2002.

6. Elliott, R., *et al.* The development of language processing support for the ViSiCAST project. In *Proceedings of ASSETS 2000*. pp. 101-108, 2000.

7. Ellis D.P.W. The Weft: A representation for periodic sounds. In *Proceedings of Int. Conference on Acoustic, Speech & Sig. Proc. ICASSP-97*. pp. 1307-1310, 1997.

8. Friedman B., *et al.* Users' conceptions of web security: A comparative study. In *Extended Abstracts of CHI 2002*. pp. 746-747. 2002.

9. Maglio, P.P. and Campbell, C.S. Tradeoffs in displaying peripheral information. In *Proceedings of ACM CHI 2000 Human Factors in Computing Systems*. pp. 241-248. 2000.

10. Mamykina, O. *et al.* Time aura: Interfaces for pacing. In *Proceedings of* ACM *CHI 2001*. pp. 144-151. 2001.

11. Ono Sokki Co., Sound and Sensors II. 2002. http://www.onosokki.co.jp/English/hp_e/whats_new/SV_rpt/SV_4/main.htm

12. Sjolander, Kare. Snack Toolkit Homepage. 2002. http://www.speech.kth.se/snack/

13. Starner, T., *et al.* A wearable computer based american sign language recognizer. In *Proceedings of ISWC'97, IEEE.* p. 84 ff. 1997.

14. Stevens, R. D. and Edwards, A. D. N. An Approach to the evaluation of assistive technology. In *Proceedings of ASSETS '96*. pp. 64-71. 1996.

15. Somervell J., *et al.* An evaluation of information visualization in attention-limited environments." In *Proceedings of VisSym 2002*, 2002.

16. US Census Bureau, Employment, Earnings and Disability – Tables. 2002. http://www.census.gov/hhes/www/disable/emperndistbl.pdf

17. Weiser, M. and Brown, J.S., Designing calm technology. 1995. http://www.ubiq.com/hypertext/weiser/calmtech/calmtech.htm

18. Wisneski, C., *et al.* Ambient displays: turning architectural space into an interface between people and digital information," In *Proceedings of the First International Workshop on Cooperative Buildings*. p. 22 ff. 1998.