

Exploring Interactive Systems Using Peripheral Sounds

Saskia Bakker, Elise van den Hoven, and Berry Eggen

Eindhoven University of Technology, Industrial Design Department, P.O. Box 513, 5600MB
Eindhoven, The Netherlands

{s.bakker, e.v.d.hoven, j.h.eggen}@tue.nl

Abstract. Our everyday interaction in and with the physical world, has facilitated the development of auditory perception skills that enable us to selectively place one auditory channel in the center of our attention and simultaneously monitor others in the periphery. We search for ways to leverage these auditory perception skills in interactive systems. In this paper, we present three working demonstrators that use sound to subtly convey information to users in an open office. To qualitatively evaluate these demonstrators, each of them has been implemented in an office for three weeks. We have seen that such a period of time, sounds can start shifting from the center to the periphery of the attention. Furthermore, we found several issues to be addressed when designing such systems, which can inform future work in this area.

Keywords: Calm Technology, Periphery, Attention, Sound design, Interaction design.

1 Introduction

As a result of the upcoming of pervasive technologies, the role of computers in everyday life is rapidly changing. This development has led to a broad discussion on how digital technologies can fit into everyday life. Weiser [14] envisioned the computer of the future vanishing into the background. By this, he aimed not only at the technology ‘disappearing’ by being hidden in artifacts or surroundings, computers will also be perceived and interacted with in the background, so that “we are freed to use them without thinking and so to focus beyond them on new goals” [14, p. 3].

This vision refers to the way we naturally perceive information in the physical world. We do not have to consciously look out the window for an impression of the current weather, or count the papers on our colleague’s desk to know that he is busy. We know these things without specifically focusing our attention on them. Computing technology however, is generally designed to be in the center of the attention. Weiser and Brown introduced the term *calm technology*, “technology that engages both the center and the periphery of our attention and in fact moves back and forth between the two” [15, p. 8]. In other words, skills acquired through (inter)action in the physical world are used to interact with or perceive information from the computer. These skills, allowing us to perceive information in the periphery of our attention or focusing on it in the center, regard all our senses. However, as sound can be perceived without having to look at the source [6], we particularly see potential for the audio to

be used in calm technology. In this paper we describe an explorative study on design of systems that use audio as calm technology. We have created three such systems and placed each of them in a shared office environment for three weeks, in order to evaluate how users perceive audible information and how this perception changes over time. Focus group interviews with people working in this office have been used as input for a discussion on interactive auditory systems designed for the periphery.

2 Theoretical Background and Related Work

In almost any situation, multiple auditory channels will reach the ear. Although all these channels are heard simultaneously, we are able to focus our attention on one channels while ignoring others [3]. We can have a conversation, while at the same time we hear music, noise outside and someone closing a door. We are thus actively, but implicitly, selecting which channel to attend to [10]. This cognitive phenomenon is commonly called the *cocktail party effect* [3] and has been extensively studied in the area of cognitive psychology [2][12]. Theories of selective attention [10] describe a selective filter in the perceptual process, which selects one channel to attend to and blocks [2] or attenuates [12] others. In a normal situation, the channel you consciously choose to attend to is selected and the attention is focused on it. The other channels are not perceived in detail. However, the selection is not only influenced by conscious choice, but also by the *salience* [10] of the incoming stimuli as well as by a cognitive process known as *priming* [12]. A loud noise for example has such distinct physical properties that it becomes salient and therefore overrules conscious choice; when a sudden loud noise is heard, it is immediately attended to and thus selected. Priming is a process that makes the selection of certain stimuli more likely. One's own name is a common example of a primed stimulus. When someone mentions your name, even in a channel you were not attended to, this will likely attract your attention. The detection of the primed stimulus in an unattended channel causes the filter to select this channel over the attended channel. See Figure 1 for an overview of this theory.

Sound is used in many interactive systems, such as for alerts, notifications, warnings, status indication, or as feedback to support interaction [13]. Although most of these sounds are designed to be in the center of the attention, some systems have been designed for the periphery. 'Audio Aura' [9] for example, uses background auditory cues to provide office workers with relevant information such as the number of unread emails or the availability of colleagues. 'Birds Whispering' [4] uses ongoing bird-sounds to subtly communicate information about the activity in the office. The 'AmbientRoom' [8] makes stock information audible through ongoing soundscapes of bird and rain sounds. Hermann et al. [7] used auditory weather reports broadcasted on the radio to provide users with awareness of the upcoming local weather. Schloss and Stammen [11] present a three art installations that make information about the current weather conditions audible in public indoor spaces.

Although all these examples aim at background monitoring of information through sound, hardly any have experimented with over a longer period of time. However, a learning period will often be needed to get used to audio of such systems. In order to find issues to address when designing peripheral audio, evaluating them over a period of time is therefore crucial.




Fig. 1. An overview of selective attention theory: normal situation where the attention is focused on a conversation (left), a situation where a (salient) loud noise is heard (middle) and a situation where the listener's name (primed stimulus) is heard on the radio (right)

3 Demonstrators and Experiments

In this paper, we present three working demonstrators, named AudioResponse, EntranceSounds and RainForecasts, each of which uses sound to convey information, intended to allow users to attune to it, but also ignore it if desired. The demonstrators, which should be considered research tools, implement a diverse range of sounds and types of information, enabling us to compare different functionalities and sound designs among each other. We have evaluated these demonstrators in three separate experiments to gain new insights in how informative sounds can play a role in interactive systems as well as to find issues to address when designing such systems and their sounds and mappings, particularly regarding longer term use.

All three experiments took place in an open office in which 12 researchers work, including the first author, of which nine actively participated in the experiments. The participants (4 female, 5 male, age 23 to 32) have diverse cultural backgrounds and none have extensive knowledge in audio related topics. Other than people entering, talking and working, PCs, lights humming, doors opening and closing, there are no significant sounds already present, e.g. there is no music playing. Footsteps are softened because of carpet on the floor. Each demonstrator ran in this office for three weeks continuously. At the start of each experiment, the participants were explained about the working of the demonstrator. See Figure 2 for an impression of the location.




Fig. 2. Lay-out of the open office used in the experiments, indicating the desks of our nine active participants in grey, and the location of the demonstrators

To evaluate the use of sound in these demonstrators, we were mainly interested in the experiences of the participants. Therefore, we gathered qualitative data to inform a discussion on this topic. All comments made by either participants or by visitors were carefully noted during the experiments. Furthermore, after each period of three weeks, a group interview was conducted with 4 to 5 participants. In this section, we will describe the demonstrators and the analysis of the experiments and interviews.

3.1 AudioResponse

Design. The AudioResponse is a simple interactive system that plays an ongoing soundscape of piano tones with semi-randomized pitch. The AudioResponse system constantly monitors the loudness (in decibel) of the sound registered by a microphone, located in the center of the room (Figure 2). This loudness determines the amplitude of the piano tones; the higher the registered loudness, the larger the amplitude of the tones. The information communicated through this sound can provide awareness of the loudness of sounds the participants and their surroundings produce. This may provide awareness of ‘what is going on around you’ in a broad sense.

Results. From reactions during the experiment and the group interview, we extracted 44 quotes regarding experiences with the AudioResponse interactive system. To analyze these data, the quotes were clustered topic-wise by the first author.

When the AudioResponse experiment ran in the office, four participants indicated that it made them aware of the loudness of certain everyday sounds. For example, a door in the hallway triggered a loud sonic response from the system, while the sound of this door was normally not experienced as very loud. Furthermore, two participants found the system useful as it warned them when they were too loud, which also caused them to attempt working more quietly to avoid triggering the system.

All participants agreed that the information conveyed by the system was not relevant. However, some participants found the system ‘fun’ at certain moments, as it triggered laughing or conversation. Others experienced it as being annoying, because already disturbing sounds were enhanced to be even more disturbing. Regarding the piano sounds used in this design, the randomness of the pitch of tones caused some confusion, as some participants expected the pitch to be linked to certain information.

3.2 EntranceSounds

Design. The EntranceSounds is an interactive system located at the main entrance door of the open office (see Figure 2). A motion sensor located above the entrance registers if someone passes through the door (see Figure 3). Whenever a person is detected, a short piano chord is played. The pitch of the root of this chord indicates the number of people detected in the last hour. For example, if someone enters at 11.32h, the number of people registered between 10.32h and 11.32h is represented. Low pitch means that few people have passed and high pitch means that many people have passed. Since the EntranceSounds system does not register the direction in which people pass through the door, entering or leaving the room is not distinguished.

This system provides information about how busy the office was in the last hour, but also informs the people working in the office that someone is entering or leaving. The door used was always open during the experiment. As the office floor is covered with carpet, one will normally hardly hear someone entering or leaving.




Fig. 3. Picture of the EntranceSounds system, located at the main entrance of the open office

Results. The 39 quotes gathered regarding the EntranceSounds system were processed in a way similar to the approach used in the AudioResponse experiment. With the EntranceSounds system, there are clearly two types of users; people that enter or leave the room and thus trigger the sounds to be played (direct users) and people working in the open office hearing the sounds triggered by others (indirect users). Most indirect users noted that the system mainly informed them that someone is entering or leaving. This was not always experienced as relevant information, as many people can directly see the door, or are just not interested in this information unless the person entering comes to visit them. However, these participants also noted that it was easy to ignore the system and that they even experienced moments where someone had come to visit them, while they had not noticed the sound.

The information conveyed through the pitch of the chord was generally considered most useful by direct users; it made them realize that it was a busy hour in the office, or that “they had not been active enough”. When knowing the routines in this office, the information turned out to be rather useful in some cases. For example: one participant came in at 10.00h one morning and noticed that the sound was higher than expected, while the office was empty. This informed her that her colleagues must have gone for a coffee break. To indirect users however, the information about the number of people having entered or left did not turn out to be relevant; none of them felt the need to be informed of this each time someone passed the door. However, the sound was not experienced as annoying or disturbing by any of the participants.

All participants could clearly recognize the pitch changes when multiple people passed the door together. However, small differences were not noticed when the time in between two chords was bigger (say 5 minutes or more). This also became apparent from an experience of one of the participants who entered the office when the sound was much higher than normal, due to an event in the room. When this participant entered the office, she noted “this is not my sound, normally it always gives me the same sound, but now it is totally different”. Apparently, she usually did not notice pitch differences, even though small differences must have been present.

3.3 RainForecasts

Design. The RainForecasts system provides audible information about the short term rain forecasts for the city in which the experiment was held. Every half hour, the system sonifies the rain forecast for 30 minutes in the future. This data is extracted from

a real-time online weather forecast [1] in terms of an 8 point scale (0 meaning no precipitation and 7 meaning heavy thunderstorm). This value is represented by a specific auditory icon [5] (see Table 1), played in the center of the room. The sounds were selected to resemble the natural occurrence of each level precipitation, but also to be recognizable as such while presented out of context. This last consideration motivated our choice for drop sounds rather than recordings of actual rain, as short samples of such recordings played at an unexpected moment sound like white noise.

The RainForecasts system differs from the systems that sonify weather information mentioned previously, by providing short term forecasts rather than real-time information [11], or forecasts for 24 hours [7]. This information may in a different way be relevant to users as it may influence their short term planning.

Table 1. Sound used in the RainForecasts system, indicating different levels of precipitation

<i>Level of precipitation</i>	<i>Precipitation in mm per hour</i>	<i>Auditory icon</i>
0	0	Bird sounds
1	< 1	Three rain drops
2	< 2	Four rain drops
3	< 5	Six rain drops
4	< 10	Eight rain drops
5	< 50	Mild thunder sound
6	< 100	Medium thunder sound
7	> 100	Heavy thunder sound

Results. The 56 quotes gathered regarding the RainForecasts system were processed in the same way as the previously described experiments. All participants agreed that they could easily recognize the rain forecast based on the sounds produced by the system. However, some participants felt that it was difficult to distinguish the different levels of rain when the sounds were played with 30 minute intervals. All participants agreed that the sound indicating ‘no precipitation’ was not distracting at all, whereas one participant indicated that the rain-drop sounds were more distracting as they seemed louder than the no-rain sounds. The participants indicated that they noticed the sounds less often towards the end of the three weeks of the experiment. Most participants mentioned that they did not hear it when working concentratedly.

From the group interview, it became evident that the information conveyed by the RainForecasts system was not of equal relevance for all participants. Some were just not interested in the weather, while one participant, who traveled by bicycle, even based the time of going home on the weather. The information was therefore very relevant for this latter participant, who noted that it was exactly the information she needed: “The internet provides a lot of information, which makes it hard to find the specific information I need”. Another participant however, wanted more detailed information (e.g. temperature) and preferred using the internet to look up forecasts.

The participants also indicated that the system provided them with information other than the rain forecasts, namely the time. The system makes a sound every half hour, which often resulted in reactions such as “did another half hour pass already? I must have been very focused!” Furthermore, some participants mentioned that the sound caught their attention more often at noon, which is the time of their usual lunch

break. This indicates that the sound is more noticeable, or moves to the center of the attention, when the conveyed information (time in this case) is more relevant.

4 Discussion

In this paper, we have described the design and evaluation of three working demonstrators that use sound to subtly communicate different kinds of information. In this section, we will discuss the insights we gained regarding *types of information* and *sound design* suitable for such systems and the *perception of audible information*.

4.1 Types of Information Conveyed by the Demonstrators

When comparing the three systems, all participants agreed that the RainForecasts system was most useful as they found the conveyed information most relevant. The AudioResponse system was considered least useful as the information provided was of no direct relevance to the participants. For this reason, some participants also experienced the AudioResponse system as being disturbing, while the other two systems did not disturb them. Interestingly, the volume of the AudioResponse sounds was not higher than that of the other sounds, and the used piano tones were similar to those used in the EntranceSounds. This indicates that the relevance of the information is related to the extent to which the sound representing it is experienced as disturbing.

Although the relevance of the audible information seems to be of importance, we have also seen that it is difficult to predict which information is relevant at which moment. For example, when one participant heard that many people had passed the door, she knew that her colleagues had gone for a coffee break. Another participant noticed the sounds of the RainForecasts system more clearly at 12.00h than at other times, as this indicated lunch time. The information that users take from such systems, is thus not always related to the information that is intended to be communicated. When and what information is relevant highly depends on the context as well as on the interests, state of mind and knowledge of the user. As multiple users are provided with the information at the same time, it can be relevant in one way to one user, in another way to another user and not at all to a third user. Given this difficulty in selecting relevant information, an iterative approach seems to be valuable for the design of such systems. Only if demonstrators are evaluated in an everyday context and with the intended users, one can assess when and how they will be useful.

When we look at the RainForecasts system, the conveyed information seemed relevant to many of the participants. However, some noted that the system did not provide enough information regarding the weather forecasts. For these participants, the information may have been too relevant to be conveyed in such 'limited' form. Although more sophisticated sound design could partly solve this (e.g. [7]), it points out an interesting issue regarding the choice of information to be made audible. This information should be relevant, but in case it is that relevant that users require more detail, the interactive system should provide easy access to a layer of detailed information. This way, general information can be monitored in the periphery via audio, and details can be examined in the center of the attention when desired. The layer of details could be displayed through audio or by other means such as visual display.

When audio is used however, it would be advisable to implement a different sonic character than the peripheral sounds. Using similar sound may confuse indirect users, who are listening to but not directly interacting with the system.

4.2 Sound Design

The presented demonstrators implemented three different sound designs and mappings. The AudioResponse used an *ongoing soundscape* of piano tones, the EntranceSounds played short *auditory cues* when users passed through the door and the RainForecasts conveyed information through *auditory icons* every half hour.

The pitch changes realized in the AudioResponse system were random and did therefore not convey any information, which caused confusion. The pitch differences in the EntranceSounds system however, revealed the number of people detected in the last hour. This has shown to be valuable at certain distinct moments. However, experiences with the EntranceSounds system have also shown that smaller pitch differences were not recognized, particularly when two tones were played with some time in between. The same issue was seen in the RainForecasts system, where participants found it hard to distinguish sounds indicating different levels of rain. If two sounds are not played successively, the differences between the sounds should therefore be clear enough to be perceived and remembered.

The sounds we used in the demonstrators presented in this paper were intended to be unobtrusive, so that they can be perceived in the periphery of the attention. The participants in the RainForecasts experiment noted that the rain-sounds were much more distracting, and thus in the center of the attention, than the sound indicating no rain. When comparing these two sounds, the rain-sounds were impact sounds and the no-rain sound was not. Using the same style of sounds may support them being perceived in the periphery and shift to the center of the attention only when required. Furthermore, impact sounds may be less suitable for these kind of applications.

4.3 Perception of Audible Information

All three experiments ran for three weeks successively, which enables evaluation of how the perception of audible information changed over time. In each experiment, the sounds were perceived in the center of the attention at the start. This means that the participants consciously heard them and often also reacted to them by looking at the demonstrator. However, in both the EntranceSounds and the RainForecasts experiments, we saw that the sounds shifted more to the periphery towards the end of the three weeks. This may indicate that getting used to the sounds can support the process of them moving back and forth between periphery and center of the attention, which is the intention of calm technology. The sounds of the AudioResponse indicating above average loudness did not shift to the periphery but were always in the center of the attention. This may be explained by the fact that many participants linked the information to themselves being loud, which annoyed them.

As mentioned before, we have seen that the information that participants took from the sounds often different from what was intended by the design. For example concluding that colleagues went for coffee based on the pitch of the chord in the EntranceSounds system. These kind of events occurred more often toward the end of the

experiment period, even though the participants were informed about the meaning of the sounds at the start of the experiments. This may indicate that a learning period is needed to get used to the direct meaning of the sounds, before users can interpret them to gain additional information. This also emphasizes the need for iterations involving longer term use of prototypes when designing such systems.

In the results described in this paper, we see that the systems were most useful when the information conveyed by the sounds differed from what the user expected. This happened for example with the EntranceSounds system when the colleagues of the direct user had already gone for a coffee break. In such cases, the sounds were clearly in the center of the participants' attention and were experienced as relevant. However, this only occurred in a small number of cases. In all other cases, the information conveyed by the sounds were as expected and did therefore not add to the knowledge of the participants. In fact, it is likely that in over 95% of the cases that a sound was played, no new information was conveyed. Though this may appear to be useless, it is exactly the intention of our designs. When comparing this to sounds in our physical environment, we see the same thing; when driving a car, the engine will sound as usual in most cases. Only in case of a problem, the sounds will be different. This conveys new and relevant information and immediately shifts to the center of the attention. When the sounds are only relevant in a low number of cases however, it is crucial to design them such that they only shift to the foreground when required.

As we have seen, at times the sounds were in the periphery and at other times they were in the center of the attention. Relating this to selective attention theory [2][12], we see that most cases when sounds shifted from the periphery to the center related to *saliency*. For example, when the rain sounds were experienced as louder than the no-rain sounds. However, sometimes participants attuned to the system for other reasons. For example, when the RainForecasts sound at 12.00h attracted the attention more than the sounds at other moments, as it indicated lunch time. This could have been the result of *priming*; the participant likely knew in the back of her mind that it was almost time for lunch, so stimuli indicating time may have been *primed*.

5 Conclusions

In this paper, we have described three interactive demonstrators that use audio to convey information in the periphery of the user's attention. We evaluated each demonstrator in a three week experiment. As a result of our studies, we have found that the participants did perceive some of the sounds used in our designs in the periphery of their attention, though getting used to the systems was required to achieve this. Furthermore, the kinds of information that participants picked up from the sounds differed depending on the context, interests, knowledge of the user, as well as on their experience with the system. As this is difficult to predict, we propose that an iterative approach, in which systems are experimented with for a period of time, is most suitable when designing such systems. This will also ensure the relevance and unobtrusiveness of the design. Although the unobtrusiveness of sound may differ per participant and depends on the relevance of the information, using a set of sounds that is consistent in style will support it being perceived in the periphery of the attention.

This paper adds to existing work by describing longer term evaluations with systems using unobtrusive sounds, which provides new insights in how sound may play a role in calm technology and what issues to address when such sounds are aimed to be perceived in the periphery of the attention.

References

1. Buienradar.nl, <http://www.buienradar.nl/> (Last accessed 15-06-2010)
2. Broadbent, D.E.: Perception and Communication. Pergamon Press, London (1958)
3. Cherry, E.C.: Some experiments on the recognition of speech, with one and with two ears. *J. Acoust. Soc. Am.* 25(5), 975–979 (1953)
4. Eggen, B., Mensvoort, K.: Making Sense of What Is Going on ‘Around’: Designing Environmental Awareness Information Displays. In: Markopoulos, P., De Ruyter, B., Mackay, W. (eds.) *Awareness Systems, Advances in Theory, Methodology and Design*, pp. 99–124. Springer, London (2009)
5. Gaver, W.W.: The SonicFinder: an interface that uses auditory icons. *Human-Computer Interaction* 4(1), 67–94 (1989)
6. Gaver, W.W.: What in the World Do We Hear? An Ecological Approach to Auditory Event Perception. *Ecological Psychology* 5(1), 1–29 (1993)
7. Hermann, T., Drees, J.M., Ritter, H.: Broadcasting auditory weather reports - a pilot project. In: *Proceedings of the International Conference on Auditory Display*, pp. 208–211 (2003)
8. Ishii, H., Wisneski, C., Brave, S., Dahley, A., Gorbet, M., Ullmer, B., Yarin, P.: ambien-tROOM: integrating ambient media with architectural space. In: *CHI 1998 conference summary on Human factors in computing systems*, pp. 173–174. ACM, New York (1998)
9. Mynatt, E.D., Back, M., Want, R., Baer, M., Ellis, J.B.: Designing audio aura. In: *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 566–573. ACM Press, New York (1998)
10. Pashler, H.E.: *The Psychology of Attention*. MIT Press, Cambridge (1998)
11. Schloss, W.A., Stammen, D.: Ambient media in public spaces. In: *Proceeding of the international workshop on Semantic Ambient Media Experiences*, pp. 17–20. ACM, New York (2008)
12. Treisman, A.M.: Verbal Cues, Language, and Meaning in Selective Attention. *Am. J. Psychol.* 77(2), 206–219 (1964)
13. Walker, B.N., Nees, M.A.: Theory of Sonification. In: Hermann, T., Hunt, A., Neuhoff, J. (eds.) *Handbook of Sonification*. Academic Press, New York (in press), <http://sonify.psych.gatech.edu/~mnees/>
14. Weiser, M.: The computer for the 21st century. *SIGMOBILE Mob. Comput. Commun. Rev.* 3(3), 3–11 (1999)
15. Weiser, M., Brown, J.S.: The Coming Age of Calm Technology. In: Denning, P.J., Metcalfe, R.M. (eds.) *Beyond Calculation: the next fifty years of computing*, pp. 75–85. Springer, New York (1997)