Fernström, M., Brazil, E. and Bannon, L. (2005) ' HCI Design and Interactive Sonification for Fingers and Ears', *IEEE Multimedia*, 12(2), p.36-44.

HCI Design and Interactive Sonification for Fingers and Ears

Mikael Fernström, Eoin Brazil, and Liam Bannon
University of Limerick

We examine the use of auditory display for ubiquitous computing to extend the boundaries of human–computer interaction (HCI). Our design process is based on listening tests, gathering free-text identification responses from participants. The responses and their classifications indicate how accurately sounds are identified and help us identify possible metaphors and mappings of sound to human action and/or system status.

With the emergence of ubiquitous and wearable computers, we need to explore alternatives to visual displays, as users sometimes need to focus their visual attention on the surrounding environment rather than on the computer. For example, using a GUI on a handheld computer while walking is extremely difficult, and indeed dangerous when standing on scaffolding on a building site high above ground.

On the other hand, we can use other kinds of equipment in such situations, such as walkie-talkies, mobile phones, and various forms of electronic instruments (for example, Geiger counters, metal detectors, and personal entertainment systems). Such devices usually have a few fixed buttons or controls (which can be attached to the user's clothing) that the user can operate with fingers. The user then learns to use the device over time with practice.

In this article, we report on the auditory displays that we devised and tested with users. We focus on two issues that we feel are important when designing interaction with auditory display:

- A method to inform the choice of suitable sounds for auditory icons, based on the analysis of listening tests.

- The design of soft buttons using auditory display to provide users with a pseudohaptic experience.

For background information on other researchers who are exploring sound, as well as basic information on some of the issues we've considered, please see the sidebar, "Considerations for Designing Auditory Displays."

Exploring what people hear

While it's assumed that everyday sounds have inherent meaning, learned from our everyday activities, hearing such sounds in isolation without context can be quite confusing. The sound of a single isolated footstep can, for example, be heard as a book being dropped on a table. Interestingly, this problem is somewhat similar to how linguistic homonyms work (words of the same spelling or sound can have different meanings depending on context).[1]

To further develop our understanding of people's perception of auditory events, we conducted listening tests, an approach also used by other researchers.[2,3] We made high-quality (44.1-kilohertz, 16-bit) recordings of 104 everyday sounds (durations between 0.4 and 18.2 seconds) and had 14 postgraduate students listen to the recorded sounds in random order using headphones, responding in free-text format to what each sound was. In most cases the descriptions they gave were quite rich. For example, the following responses (for three different recordings) describe the events quite accurately:

- "A person walking on a carpet with their hands in their pockets hence the clanging of keys or coins, taking five steps and turning to retrace their footsteps."

- "A metal spoon in stirring motion in an empty ceramic cup, tapping the cup as if to displace the liquid from the spoon and then placing the spoon onto a table."

- "Breaking of a cup (not really a glass sound more ceramic I think)."

Several ways exist to analyze the responses from such listening tests. The most obvious way would be to count how many responses could be deemed correct for each sound by linking the sound to the participants' reported understanding of the objects and actions involved in producing the sound. A somewhat more interesting measure is Ballas' method of causal uncertainty.[4] Ballas et al.[5] found that identifi-

1070-986X/05/$20.00 © 2005 IEEE

Considerations for Designing Auditory Displays

In novel human–computer interaction (HCI) paradigms,[1-3] such as ubiquitous, pervasive, wearable, and disappearing computing, interactive sonification might offer useful alternatives to the otherwise dominant visual displays, freeing up our eyes to see the surrounding world or do what small visual displays don't do so well. In brief, using sound for display in interaction design is useful for attracting attention to events or locations, for nonvisual communication in general (including speech), alarms, notification, and feedback. Sound is less useful for continuous display of objects, for absolute readings (most people perceive auditory dimensions such as pitch, loudness, and timbre as being relative), and for fine-detail spatial display. Sound is also problematic in noisy or noise-sensitive environments.

Designing interactive sonifications for HCI requires that we address numerous issues. We have to consider where and how sonification is appropriate. As designers, we also need to take into account the users' capabilities while carrying out tasks in real environments, and consider that surrounding noise levels might mask the system's sounds. If the sound will enhance the interaction, we need to explore ways of creating and testing auditory metaphors. We should also investigate to what extent the use of sound contributes to the users' performance and subjective quality of use.

To be able to design with sound, we need a high-level understanding of what and how we hear (for example, see Gaver's research[4,5]). While an extensive catalogue of studies exists on the perception of musical sounds and speech, researchers know relatively little about other kinds of nonspeech sounds—in particular everyday sounds (such as footsteps, creaking doors, water filling a container, bouncing, and breaking).

In the new HCI paradigms, we can explore new concepts of interaction and human activity. In previous work on auditory interfaces, ranging from Gaver's Sonic Finder[6] to Brewster's hierarchical earcons,[7] human action has to a large extent been thought of in a discrete way—like kicking a football, where a user action starts a process that then completes without any further user control. This view might be appropriate when typing, clicking, or flicking switches. An alternative view is action as a continuous flow, such as a pen stroke, where we continuously move a pencil on a surface, relying on our learned gesture through proprioception, as well as haptic, visual, and auditory feedback. This latter view is becoming important, now that several input devices (such as pens, digitizers, and cameras) are capable of detecting complex human actions.

Still, at the core of our design space, a fundamental problem is how to classify and select suitable sounds for a particular interaction design. Depending on our intended users, tasks, and context, initially a broad continuum exists in this design space, ranging from concrete to abstract displays (that is, from auditory icons to earcons).[8,9] If we're designing for casual everyday use, we probably need to consider concrete forms. If we're designing for highly specialized domains (such as cockpit or process-control applications) where our users will be selected and trained for high-performance requirements, we might need to focus on psychoacoustic issues such as detection accuracy or time and perceived urgency. In the latter case the design space can be more abstract.[10]

References

1. E. Bergman, *Information Appliances and Beyond*, Morgan Kaufmann, 2000, p. 385.
2. B. Shneiderman, *Leonardo's Laptop: Human Needs and the New Computing Technologies*, MIT Press, 2002.
3. H. Rheingold, *Smart Mobs: The Next Social Revolution*, Perseus Books, 2003.
4. W.W. Gaver, "How Do We Hear in the World? Explorations of Ecological Acoustics," *Ecological Psychology*, vol. 5, no. 4, 1993, pp. 285-313.
5. W.W. Gaver, "What in the World Do We Hear? An Ecological Approach to Auditory Source Perception," *Ecological Psychology*, vol. 5, no. 1, 1993, pp. 1-29.
6. W.W. Gaver, "The Sonic Finder: An Interface that Uses Auditory Icons," *Proc. Assoc. for Computing and Machinery Computer–Human Interaction* (ACM CHI), ACM Press, 1989, pp. 67-94.
7. S.A. Brewster, P.C. Wright, and A.D.N. Edwards, "A Detailed Investigation into the Effectiveness of Earcons," *Auditory Display: Sonification, Audification, and Auditory interfaces*, G. Kramer, ed., Addison-Wesley, 1994, pp. 471-498.
8. S.A. Brewster, "Non-Speech Auditory Output," *The Human–Computer Interaction Handbook*, J. Jacko and A. Sears, eds., Lawrence Erlbaum Associates, 2002, pp. 220-239.
9. W.W. Gaver, "Auditory Interfaces," *Handbook of Human–Computer Interaction*, M.G. Helander, T.K. Landauer, and P. Prabhu, eds., 2nd ed., Elsevier Science, 1997.
10. R.D. Patterson and A.J. Datta, "Extending the Domain of Auditory Warning Sounds: Creative Use of High Frequencies and Temporal Asymmetry," *Human Factors in Auditory Warnings*, N.A. Stanton and J. Edworthy, eds., Ashgate, 1999, pp. 73-90.

cation time for everyday nonspeech sounds was a function of the logarithm of the number of alternative interpretations of a sound. This led them to suggest using an information measure H to quantify the causal uncertainty of a sound, as Equation 1 shows:

$$H_{CU} = \left| \sum_{j}^{n} p_{ij} \log_2 p_{ij} \right| \quad (1)$$

H_{CU} is a measure of causal uncertainty for sound i, p_{ij} is the proportion of all responses for sound i

sorted into event category *j*, and *n* is the number of categories for responses to sound *i*. Applying this equation implies that if all participants in a listening test give the same response, the causal uncertainty is 0 (all participants agree). For example, with 14 participants if the responses are distributed 50/50 between two alternatives, the causal uncertainty is 1.0. If the distribution of responses is skewed—such as 13 of the 14 responses are the same but one response is different—the causal uncertainty is 0.37. If all 14 responses are different, the causal uncertainty is 3.8. From this we can see that calculating causal uncertainty according to Ballas' method gives a good measure of how easy it is for users to identify everyday sounds.

With our collected data (responses from 14 participants listening to 104 different sounds) the responses were sorted and categorized, as well as evaluated for correctness, by two of the authors and a research assistant. The reliability between the evaluators was significant (weakest $r = 0.78$, $p < 0.0016$). From the responses, we extracted and categorized action and object segments of the texts, such as how the objects/materials interacted and what objects/materials were used. We found that in general 32 percent of the sounds were identified correctly, while for action segments it was 38 percent and for object segments it was 25 percent.

The collected data set with all responses, categorizations, and measurements of causal uncertainty can also be used for suggesting the possible use of sounds in interaction design, somewhat similar to Barrass' method of collecting stories about when sound is useful in everyday life.[6] From a designer's point of view it's interesting to note that the responses from the listening tests contain information about how people describe everyday sounds as well as measurements of causal uncertainty.

Being able to parametrically control sound models in real time can also, potentially, help make sonifications less annoying.

Sounding objects

With the results of the listening tests, we can begin to suggest possible auditory displays and metaphors for interaction design. Based on Barrass' TaDa approach,[6] we can do a task and data analysis that lets us select sounds that can communicate the dimensions and directions that give users adequate feedback about their actions in relation to the system as well as the system's status and events. We then need to create ways so that the system can produce the selected sounds and finally evaluate the resulting design with users.[7] If we were to just play sound files as feedback to user actions, it would always sound the same and never (or seldom) be expressive (for example, to be mapped to the user's effort or the size of the data objects involved).

This was one of the issues addressed by the European Union's Sounding Object project (see http://www.soundobject.org), where new methods for physically inspired modeling of sounds for sound synthesis were explored. Our work was initially largely informed by ecological acoustics and Gaver's work on auditory icons.[8]

We also worked toward *cartoonification* of sound models—that is, simplifying the models while retaining perceptual invariants. We implemented the models in Pure Data (commonly known as PD; see http://www.puredata.org) and tested them in a number of ways, ranging from perceptual experiments to artistic performance. Compared to ordinary sound files, sound objects can provide "live" sound models that we can parametrically control in real time with reasonable computational power.

Being able to parametrically control sound models in real time can also, potentially, help make sonifications less annoying. With prerecorded sound files, sounds used in an auditory interface always sound exactly the same. In contrast, with sound objects and parametric control we can vary properties of the sounds—for example, mapping the size of objects or the effort of actions—so that small objects or actions make small sounds and large objects or actions make large sounds.

Revisiting the overall results from the Sounding Object project,[9] it's interesting to note that all the sound models developed throughout the project point toward an epistemology that differs from Gaver's trichotomy of primitives of solids, liquids, and gases. An alternative emerging view indicates that the primitive classes might be better understood if we think of the world of sound-producing events as composed of impacts, frictions, and deformations. The simulated material properties of the objects involved in such interactions are controllable through parameters passed on to our sound models. The analysis of listening tests, as previously described, also suggests that actions are better identified than objects. This might suggest that interaction design using auditory display should focus on mapping human activity to actions rather than objects.

Example: Auditory soft buttons

The idea of software-defined buttons—*soft buttons*—emerged from research on direct manipulation and GUIs (from early work in Xerox PARC—Palo Alto Research Center[10]) and is now an important part of all GUI widget libraries. Most GUIs use soft buttons extensively, ranging from desktop personal computers and laptops (see, for example, Figure 1), to personal digital assistants (see Figure 2). With soft buttons the designer—and sometimes also the user—can easily modify, add, or remove a software application's interactive controls.

The ways that users can activate soft buttons vary. On desktop computers the most common way is to move a pointing device, such as a mouse, that in turn indirectly moves a visible cursor on screen into the rectangle surrounding the soft button. The user then activates it by clicking with the pointing device. Visual soft buttons are often animated to improve feedback to the user—for example, when the user clicks, the visual button displayed temporarily changes its appearance so that the graphical symbol looks like it's moving inwards, into the display surface. On other kinds of computers, such as handheld computers, the user can point directly to a visual soft button, either with a handheld stylus or simply with a finger.

User interface widgets such as these make the design of GUIs highly malleable and flexible, as the designer can display and represent highly complex underlying functionality with simple graphical symbols that, ideally, look like concepts or entities in the user's task domain. Because it's usually the software rather than the hardware that defines such widgets, the same physical display surface can be used for different widgets at different times, supporting the varying needs for the user to carry out different tasks. These features make soft buttons attractive components of user interfaces, both for designers and users. Along these lines, our interest is in using forms of display other than vision to create similar affordances, in this particular case through auditory feedback mimicking what it would sound like to touch differently structured surfaces.

Figure 1. Visual soft buttons, such as in this Microsoft GUI example, are currently one of the most common widgets used by interaction designers.

Figure 2. With soft buttons, the same display surface can easily be used for different button layouts and functionality—all defined by software. This example shows two different calculator applications on a Palm Pilot. To the left, a simple standard calculator; to the right, a scientific calculator.

Pseudohaptic soft buttons using auditory display

In three experiments we investigated the use of auditory display to create a pseudohaptic experience of soft buttons.

Pilot 1: Real haptics

First, we conducted a pilot experiment, based on Gibson's[11] cookie-cutter study. With four different paper shapes glued on paper, a participant found it easy to feel the shapes and then draw an image of them, picking up the shape of the objects through haptic perception and visualizing the shapes.

Pilot 2: Pseudohaptics using auditory display

Based on the first pilot experiment, we designed a second study of a soft-button proto-

Figure 3. Using visual soft buttons on a Xybernaut arm-mounted touch screen.

Figure 4. With interactive sonification, we don't need to look at the touch screen. The user can hear an interactive sonification of soft buttons when touching the touch screen, now worn on the belt.

Table 1. Mapping between actions and sound.

Action	Sound	Function
No touch	N/A	N/A
Touch area outside button	N/A	N/A
Enter button area	Tick	N/A
Move finger on button	Friction sound	N/A
Exit button area	Tack	N/A
Lift finger off button	Tock	Select/activate function

type, testing the idea of having soft buttons displayed by audio instead of graphics. We asked three users to make drawings of three different soft-button layouts. We used a wearable computer and touch-sensitive display subsystem from Xybernaut (http://www.xybernaut.de; see Figure 3), which is normally worn on the arm. In this experiment, however, we affixed the subsystem to the user's belt in a position so that the users could comfortably rest their hand on the unit, with fingers free to access the touch area (see Figure 4). The size of the active touch area was 120 × 90 mm. We only used the touch detection of the device, not the visual display.

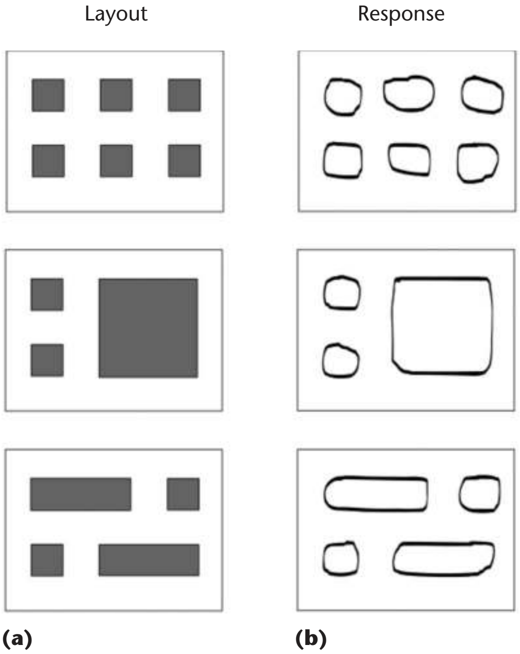

Figure 5. (a) Soft-button layouts and (b) example of user-response drawing.

We created three different layouts with soft buttons (see Figure 5a). When a user moved a finger on a button area, a simple friction-like sound was produced. To emphasize the boundaries of each button, click sounds were produced when entering or exiting a button area (see Table 1 for the mapping between actions and sounds). The sounds were heard in mono, using a simple headphone in one ear.

We recruited three participants. Each participant spent approximately 10 minutes getting familiar with the design and making three drawings. We found that the participants were able to feel their way around the touch device and make quite accurate drawings of the soft-button layout, as we show in Figure 5b. (This example is from one user, although all three made similar drawings.) This indicated that this kind of auditory display of soft buttons lets users have a pseudo-

Figure 6. Tactex touch tablet.

haptic experience, giving them a mental model of the device's layout.

Pilot 3: Haptics and pseudohaptics

The focus in our final experiment was the user detection of and interaction with soft buttons using an auditory display. We refined our experiment by first selecting a different touch device—a Tactex touch tablet called MTC Express (see Figure 6)—that differs from the previously used Xybernaut device in that the MTC Express device doesn't have a visual display.

The interactive area of the Tactex device matches the size of a human hand quite well. The active touch area is 145 × 95 mm. The point of contact where the hand rests is a spatial and haptic reference point for finger movements on the device. We also redesigned our software to minimize any latency and implemented six different soft-button layouts (see Figure 7).

Procedure. We recruited 10 participants among our postgraduates. We tested all six layouts both as paper shapes on cardboard and with the Tactex touch tablet connected to a Windows PC. Each stimulus was tested twice, resulting in 24 drawings per participant (12 haptic, 12 auditory/pseudohaptic). The order between stimuli was randomized. Users had headphones to listen to the sounds (in mono) while interacting with the system. They were allowed to use either their left or right hand to explore the layouts and to draw their understanding of the layouts.

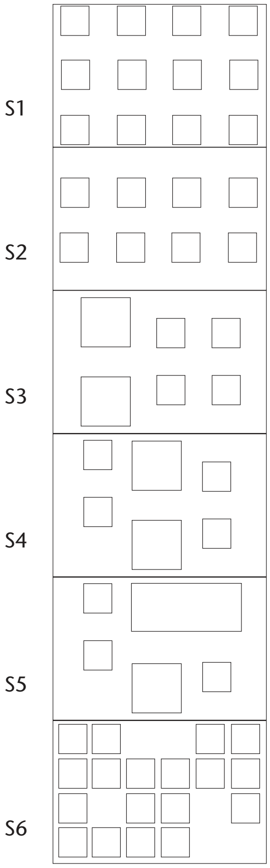

Figure 7. Six different soft-button layouts used in the final experiment. S1 and S2 are simple symmetrical layouts with buttons of the same size; S3–S5 have buttons of different sizes and also introduce different degrees of asymmetry; S6 challenges the user as buttons are packed closely together.

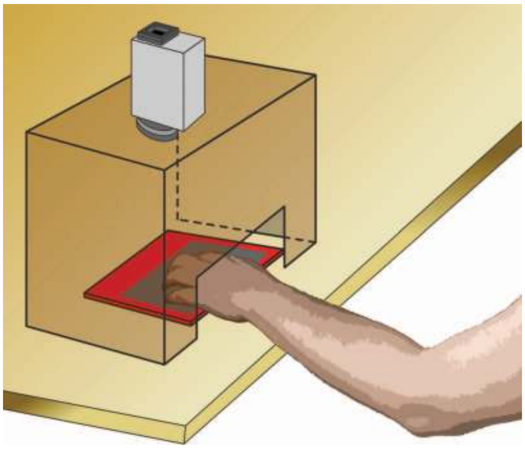

Figure 8. Experimental setup of the touch area. We wanted to make sure that the users didn't pick up visual clues from seeing where they moved their hand while exploring the haptic and pseudohaptic soft buttons. A video camera was positioned on top of the box so that we could record the user's hand movements.

To prevent our participants from seeing where they were moving their fingers on the Tactex tablet, we covered it with a cardboard box with a cut-out for the user's hand to reach the active touch area. On top of the box, a video camera was fitted for recording the participants' hand and finger movements (see Figure 8). The same box was used for both haptic and pseudohaptic stimuli. The participants were given 3 minutes of explo-

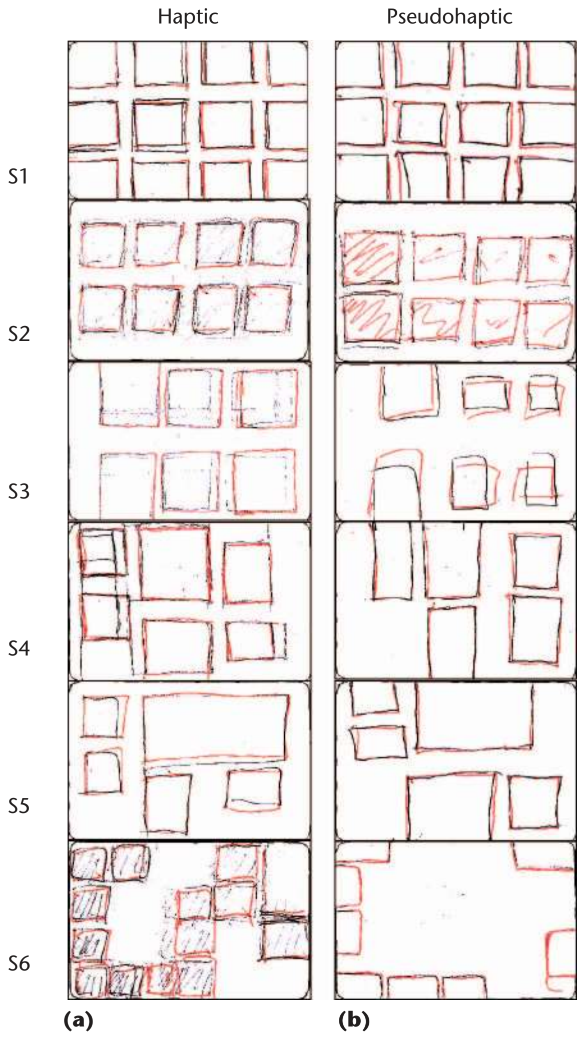

Figure 9. Examples of the drawings that the users made when exploring the button layouts (a) with haptic buttons and (b) with pseudohaptic soft-buttons user auditory display.

ration time per stimulus. During their exploration they used a black pen to sketch, and at the end of each 3-minute period they used a red pen to mark their detected layout on a blank sheet of paper.

Results. Figure 9 details 12 typical examples out of 240 drawings, all of which are available at http://www.idc.ul.ie/mikael/softbuttons/. As Figure 9 shows, the participants were good at detecting the layouts, both in the haptic and pseudohaptic (using auditory display) conditions. With the more complex layouts, the number of errors increased—or on occasion they ran out of time (particularly in the pseudohaptic condition). In the debriefing sessions, participants reported the haptic and pseudohaptic conditions to be almost as easy (or difficult, for the more complex layouts).

Discussion. Our results indicate that using an auditory display to create a pseudohaptic experience, based on sound object models synthesized in real time, is almost as efficient and accurate as a haptic display. The participants needed more time to detect more complex layouts using auditory display, but because they didn't know anything about the layouts in advance, we assume that if we allowed them to familiarize themselves for a longer period of time with a particular set of layouts, the difference would become smaller. The same applies to many other human activities, such as shifting gears in a car or typing. As our actions become automatic, we need fewer clues regarding the success of our actions.

Our findings indicate that this kind of auditory display of soft buttons lets users have a pseudohaptic experience that supports the development of a mental model of the device's layout. A similar pseudohaptic approach was investigated by Müller-Tomefelde,[12] who in one of his demonstrations communicated differences in surface texture through friction-like sounds in a pen-based digitizer application. In the commercial world, Apple Computer's Ink application for handwriting input with a digitizer tablet also attempts to enhance the user experience through friction-like sounds as feedback to the user's pen strokes with a stylus.

Future research

In this article, our approach has been that designs should be based on the results from listening tests with possible metaphors being extracted from users' descriptions of everyday sounds. The listening tests can also provide guidance in our understanding of how users interpret combinations of auditory icons.

More studies are needed on what people hear when they listen to everyday sounds to increase our understanding of the perceptual and cognitive processes involved. In particular, studies of the effects of combinations of different auditory icons in sequence or in parallel are lacking.

We've found that the PD environment and Sounding Objects project are both highly productive approaches for prototyping sound designs for interactive sonification. However, for fully integrated applications we need to seriously consider if we can more closely integrate a set of sonification primitives with operating systems. This can in turn result in the development of toolkits for developers, similar to what's available for GUIs today. A need also exists to educate and support interaction designers so that they can open up their creative thinking toward interactive sonification, and realize that it's possible to provide continuous feedback in real time for gesture-based devices.

All components in HCI also have aesthetic properties. It's probably possible to design sonifications that are psychoacoustically correct and quite efficient but unpleasant to listen to (just as it's possible to design visual interfaces that users find unpleasant to view). As Somers[13] has suggested, we need to draw upon the knowledge and ideas of Foley artists (sound design for film, radio, and television) as well as lessons learned from various theories of acousmatic music. **MM**

Acknowledgments

We would like to express our thanks to all the postgraduate students at the University of Limerick who participated in our experiments. We thank Xybernaut for giving us access to their wearable computers and Tactex for giving us access to their device driver source code. We're indebted to all our colleagues in the Sounding Object project for an ongoing and engaging dialogue. Finally, we thank our colleagues Annette Aboulafia, Thomas Waldmann, and Michael Cooke for their comments on earlier versions of this article; Stephen Shirley for software development for the Tactex device; and the anonymous reviewers for their creative and encouraging comments. This work was partly financed by the European Union Future and Emerging Technologies program; the Disappearing Computer/Sounding Object project IST-2000-25287; the Irish Research Council for Science, Engineering and Technology/Multimodal Systems project SC-02-302; and the Science Foundation Ireland: Shared Worlds Project, 2003–2007.

The research in this article was originally published in the *Proceedings of the International Workshop on Interactive Sonification* in January 2004. The article has been revised and updated to fit with *IEEE MultiMedia's* standards.

References

1. J.A. Ballas and R.T. Mullins, "Effects of Context on the Identification of Everyday Sounds," *Human Performance*, vol. 4, no.3, 1991, pp. 199-219.
2. N.J. Vanderveer, "Ecological Acoustics: Human Perception of Environmental Sounds," *Dissertation Abstracts Int'l*, PhD dissertation, Cornell Univ., 40/09B, 4543 (Univ. Microfilms no. 8004002), 1979.
3. W.W. Gaver, *Everyday Listening and Auditory Icons*, PhD dissertation, Univ. of California, San Diego, 1988.
4. J.A. Ballas, "Common Factors in the Identification of an Assortment of Brief Everyday Sounds," *Experimental Psychology: Human Perception and Performance*, vol. 19, no. 2, 1993, pp. 250-267.
5. J.A. Ballas, M.J. Sliwinsky, and J.P. Harding, "Uncertainty and Response Time in Identifying Non-Speech Sounds," *J. Acoustical Soc. of Am.*, vol. 79, no. 47, 1986.
6. S. Barrass, "Auditory Information Design," PhD dissertation, Australian National Univ., 1997.
7. B. Walker and G. Kramer, "Mappings and Metaphors in Auditory Displays: An Experimental Assessment," *Proc. 3rd Int'l Conf. Auditory Display* (ICAD 96), Int'l Community for Auditory Display, 1996.
8. W.W. Gaver, "Using and Creating Auditory Icons," *Auditory Display: Sonification, Audification and Auditory interfaces*, G. Kramer, ed., Addison-Wesley, 1994, pp. 417-446.
9. D. Rocchesso, R. Bresin, and M. Fernström, "Sounding Objects," *IEEE MultiMedia*, vol. 10, no. 2, Apr.–June 2003, pp. 42-52.
10. A. Kay and A. Goldberg, "Personal Dynamic Media," *Computer*, vol. 10, 1977, pp. 31-44.
11. J.J. Gibson and S. Stanely, "The Accuracy of Form Discrimination with Three Types of Tactual Input: Passive, Moving, and Active," *Purple Perils*, Aug. 1961; http://www.trincoll.edu/depts/ecopsyc/perils/folder1/accuracy.html.
12. C. Müller-Tomefelde, "Sounds@Work—Akustishe Repräsentationen für die Mensch-Computer Interaktion in kooperativen und Hybriden Arbeitsumgebungen [Sounds@Work—Auditory Displays for Interaction in Cooperative and Hybrid Environments]," *Elektrotechnik und Informationstechnik*, PhD dissertation, Technical Univ. Darmstadt, 2003; http://elib.tu-darmstadt.de/dis/000313/.
13. E. Somers, "Abstract Sound Objects to Expand the Vocabulary of Sound Design for Visual and Theatrical Media," *Proc. Int'l Conf. Auditory Display* (ICAD 2000), Int'l Community for Auditory Display, 2000.

Mikael Fernström is the director of the master's program in Interactive Media at the University of Limerick, Ireland. His main research interests are interaction design, sound, music, and electronics. Fernström studied electronic engineering and telecommunications at the Kattegatt technical college in Halmstad, Sweden, and has an MSc in human–computer interaction from the University of Limerick.

Eoin Brazil is a research assistant in the Interaction Design Centre at the University of Limerick, Ireland. His main research interests are interaction design, sound, software visualization, and music retrieval. Brazil has a BSc and MSc in computer science from the University of Limerick.

Liam Bannon is the director of the Interaction Design Centre and professor of computer science in the Department of Computer Science and Information Systems at the University of Limerick. He's interested in improving the usability and desirability of computational artifacts. He also examines alternative conceptual and methodological frameworks for understanding human activities (and their mediation by various technologies). Bannon has an MSc in computer science from Trinity College in Dublin, Ireland, and a PhD in experimental psychology from the University of Western Ontario, Canada.

Readers may contact Mikael Fernström at the Interaction Design Centre, Computer Science Dept., Univ. of Limerick, Limerick, Ireland; mikael.fernstrom@ul.ie.

For further information on this or any other computing topic, please visit our Digital Library at http://www.computer.org/publications/dlib.