Final Report: Principles for Improving Interaction in Telephone-Based Interfaces (GR/L66373)

Stephen Brewster
Department of Computing Science,
University of Glasgow, Glasgow, G12 8QQ, UK
Email: stephen@dcs.gla.ac.uk   Project Website: www.dcs.gla.ac.uk/~stephen/research/telephone

BACKGROUND AND CONTEXT

Telephone-based interfaces (TBIs) are an increasingly important method for interacting with computer systems. The telephone is a ubiquitous device and is many people's primary method of entry into the information infrastructure. Access to an increasing number of services is being offered over the telephone (such as electronic banking and Web pages). The provision of these extra services may be rendered useless if usability issues are not considered [23]. The rapidly increasing use of mobile phones means that people access services at many different times and places. This is a huge growth area – there are approximately 25 million mobile phone subscribers in the UK and this is expected to rise to 45 million over the next five years [18]. Mobile telephones themselves now incorporate greater functionality (such as diaries, address books and WAP services). We therefore concentrated our research on the interfaces to mobile telephones. The Government's Foresight panel recognised mobile devices and their interfaces are a key area of focus in UK science and engineering [16]. The telephone itself allows only a limited form of interaction [27]. In TBIs (both for dial-up services and interfaces to mobile phones) the user interacts by navigating through a series of menus. A common problem is that users get lost in the menu hierarchy [26, 30, 31] of the device or service. One reason for this is that the interfaces are very limited – input is via a small number of keys and output via speech or a small screen.

There are two main approaches to overcome navigation problems. The first is to use speech recognition and a text-to-speech system. Here the user would speak a command and the system would reply in speech. This avoids many navigation problems but users have to remember all of the possible commands that they can say. Speech recognition is also poor in noisy environments and Williams *et al.* [29] claim that the length of spoken menu prompts is a significant problem with speech-based menu designs. We took a novel approach – using non-speech sounds to provide navigation cues to enhance the existing interfaces. Little use has been made of structured non-speech sound in TBI's (apart from the standard dial tone, engaged tone, etc.) [13]. Most guidelines for the design of TBI's [23, 27] include nothing about the use of non-speech sound. But sound has many advantages, for example it is good for communicating information quickly and is not language dependent. Our initial work, using structured sounds called *earcons* [5, 15], showed that sound could provide basic navigation cues in very simple menu hierarchies. The aim of this project was to develop and extend these ideas.

KEY ADVANCES & SUPPORTING METHODOLOGY

The key aim of the project was to show that we could use non-speech sounds to provide navigation cues in TBI menu hierarchies. We did this with a series of gradually more complex experiments that culminated in the testing of the menu hierarchy on a real Nokia phone. We performed rigorous experimental evaluations of our ideas at each stage to ensure our results were reliable.

Our approach was to provide both *semantic* and *contextual* cues [17, 25] in sound. The lack of screen space on mobiles means that it is very difficult to provide these kinds of cues graphically. For the simple menus used in our earlier studies timbre had supplied the key semantic cue and rhythm the key contextual one [5, 15]. This was sufficient for very simple menus but would not be enough to represent more realistic, complex menus.

Three main experiments were carried out and the results of each fed into better design of sounds for the next and also into our design principles for TBIs. The first two experiments allowed us to improve the design of the sounds beyond those done before the project started [5, 15]. The third forced us to extend our ideas to work with a real mobile phone and its menu structure. We also developed two new interaction techniques: one to control sound volume to avoid annoyance and the other to present continuous, time-varying data in sound. Finally, based on our experience, we produced a tool to help interface designers quickly and efficiently use the knowledge of earcon design that we gained.

Using non-speech sounds for navigation in TBIs

Initial work leading up to the project had shown that earcons could provide successful navigation cues in very simple menu hierarchies [5, 15]. The hierarchical earcons from these early experiments (and those proposed by Blattner *et al.* [4]) had a problem: it was very difficult to create a hierarchy deeper than 4 levels as once timbre, pitch, rhythm and tempo had been used there was nothing left to manipulate to create more levels. Therefore our first experiment looked at the use of a 'decimal' notation for creating the sounds that would allow hierarchies to go deeper (which is very common, particularly in mobile phone menus which can be very deep). Here the position of the sound in the hierarchy was based around a chapter 1/section 1.1/sub-section 1.1.1 metaphor (see Figure 1). This allowed the possibility of going deeper than four levels (in fact, it would allow any depth). Earcons were created for each of the numbers 0-9 and '.' these could then be concatenated to produce a sound for any position in the hierarchy. The hierarchy in Figure 1 was the same as I had used in previous experiments so it was possible to make a direct comparison with 'traditional' hierarchical earcons (see [6] for a full discussion of these earcons). An experiment was conducted to assess the usability of this new design. The decimal sounds proved very successful with recall rates at an average of 97% correct, significantly better than for standard hierarchical earcons [6-8]. Users could recall the sounds well and predict where a sound that they had not heard before would place them in the hierarchy significantly better than with our standard hierarchical earcons. However, its drawback was that to create earcons for deeper menu items meant concatenation of more sounds, so the earcons got longer and longer. This was a problem, as the sounds had to be short (one of our key principles – see below) as users move between menu items very quickly. It also easily overloaded the users' short-term memory – important if users are trying to perform other tasks at the same time.

To solve this problem we attempted to expand the design space of standard hierarchical earcons. Leplâtre (the PhD student on the project) had a background in music so I was keen for him to use his skills to make more sophisticated sounds. Previous designs were based around low-level psychoacoustical parameters and sounded very mechanical. Leplâtre used work from musical composition (developing the ideas of Alty [1, 2]) to use higher-level musical constructs such as melodic structure and harmonic progression [20]. Our previous work had demonstrated that timbre was a key cue for the design of hierarchical earcons [5, 15]. We therefore decided to expand our knowledge of the other parameters of sound by running a second experiment where timbre was fixed (a piano) and we created a series of different melodic structures to see if we could generate some other useful parameters for the design of the sounds (see the project website for a Java applet that demonstrates these sounds).

The sounds were designed as follows. The root of each family (L1 in Figure 1) was represented by a short *motive* or musical snippet [4] (simplest for node 1 and most complex for node 4). The cues for level 2 were represented by the level 1 motive followed by a group of notes or chords: for the first node of this level (e.g. 1.1) one note/chord was used, for the second 2 notes/chords, and 3 for the third. For level 3, families 1 and 2 the cue was a chord following the sound of the parent node. For families 3 and 4 the cue was a high register melody played in parallel with the sound of the node above (full musical scores can be found in [19] or [35]). The sounds ranged from simple (family 1) to more complex (family 4) and from playing serially (families 1 and 2) to parallel (3 and 4).

An experiment was conducted to assess the effectiveness of this new sound design method (again using the same hierarchy). The results showed that, on average, 86% of the sounds were recalled correctly. This was not significantly different to the results in experiment 1 [6] but was achieved without the use of timbre – the most important discriminating factor. We also achieved this recall rate with a much lower subjective workload than in the decimal sound study. This meant that the higher-level musical structures used had a great effect on recall [19]. They could be combined with timbre differences to build a larger set of parameters to sonify larger hierarchies (see below for a summary of our principles).

Realistic evaluation of the sonified TBI

From the previous two studies we gained a much greater understanding of how to design earcons to represent hierarchies. In our third experiment we put this knowledge into a full-scale test of a real mobile phone menu structure based around the menu of a Nokia 6110. We built a simulator for this phone (see the project Web site) so that our testing was done on a desktop prototype that allowed us to easily use and control sounds. It was important that we grounded our research with a more realistic experiment that evaluated a real hierarchy of the type which would occur in our target devices. Until this point our largest hierarchy had only contained 35 nodes – far fewer than a real mobile phone (but more like a standard dial-up TBI). The Nokia 6110 had a hierarchy of over 150 nodes. No one had tried to sonify a

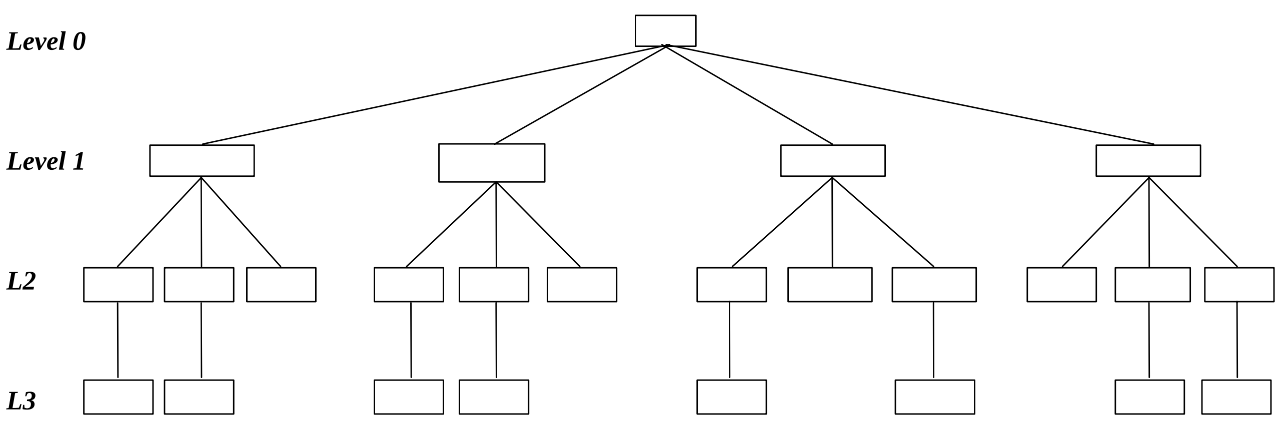

Figure 1: The hierarchy used in the first two experiments in the project (based on the hierarchy from Brewster [5]). The numbers by each node show the position in the hierarchy. L0 – L3 represent the depth of the level within the hierarchy.

hierarchy of anything like this size before. We needed to refine our principles (see below for a summary) to cope with designing such a large hierarchy and ensure that people could use our sounds for navigation.

The design of earcons used in all of our previous studies encoded in the sound the path from the root to the current node. Users could then listen to this and hear their location. This meant that the sounds got more complex the deeper one went into the hierarchy. It proved to be difficult to follow such a strategy when designing sounds for a menu with up to 10 levels. For this experiment we tried a different approach: we reversed the way that we constructed the sounds. They were most complex at the top level and became simpler as the user moved down through the hierarchy.

In our original designs the earcons had provided absolute positioning within the menu structure as the tasks performed had not been based around those done with a real phone. Now all we needed was relative position as users had to move from one menu to another to navigate, as they would in a real phone. This allowed us to simplify the design problems somewhat – we could make use of changes between sounds to provide information. For large hierarchies it is also important to know what is beneath the current node in helping to decide if it is the right path to follow (a key contextual cue) so we also included this in our sounds.

The sounds were designed as follows. Items on the first main level of the menu hierarchy (the equivalent on the phone of L1 in Figure 1) were visually displayed differently to all of the others – they had an icon associated with them. We followed this (to aid in semantic navigation) by giving each menu item a different timbre (the main grouping factor we had found from our earlier studies). A harmonic progression involving an alternation of chords and arpeggios was used from the first menu item at this level to the last to keep the items sounding homogeneous when a user moved through the items (providing a contextual cue). For example, the menu item for Call Divert can be seen at the top of Figure 2. The new approach we took at this level was to create each sound based on the size of the hierarchy beneath it. For a sub-tree with little below it the menu sound would be very simple. For a menu with larger number of items beneath it we constructed a complex sound (see Figure 2). This gave users some information about what was beneath them in the hierarchy.

At the next level down timbre was inherited from the level above (keeping the semantic cues consistent). The complex sound from the parent node was used as a harmonic reference for a progression of simpler chords. Figure 2 illustrates how the middle level sounds involved a descending chord progression (from left to right) in the tonality of the parent motive. The same process was used to represent the items at subsequent levels. The sounds for the bottom level in Figure 2 are the notes that make up the parent at the middle level.

One problem that often occurs in menu navigation is looping: when a user gets to the last item another key press loops him/her back to the first one. This is often not noticed, as visual feedback is limited, so that users scroll too far. With our design of sounds this error is obvious as the sound for the first item in the menu does not follow on from the last in harmonic structure and so stands out very saliently.

To test our design an experiment was conducted, based on users performing a range of typical mobile phone manipulation tasks, e.g. setting profiles, call diverts and the clock, on a phone that they had not used before. One group used a phone sonified with our sounds and the other used a standard one (which just played a tone when keys were pressed). Our hypothesis was that the navigation sounds should allow users to build a better mental model of the menu structure of the phone and which would translate into faster performance and fewer errors. Users with sonified menus should also get better more quickly than users of the standard phone. The results showed that the number of key-presses needed to perform tasks started at a similar level in the two groups but by the end of the experiment the sonified group took an average of 28% fewer. Similarly, the number of tasks completed was initially similar but by the end, significantly more were completed correctly with the sonified phone. Subjective workload ratings showed no differences between our sonified phone and the standard one (avoiding the problems of experiment 1). For full details see [22].

We are following up this experiment with a fourth and final evaluation. Seppo Helle at Nokia has put a set of

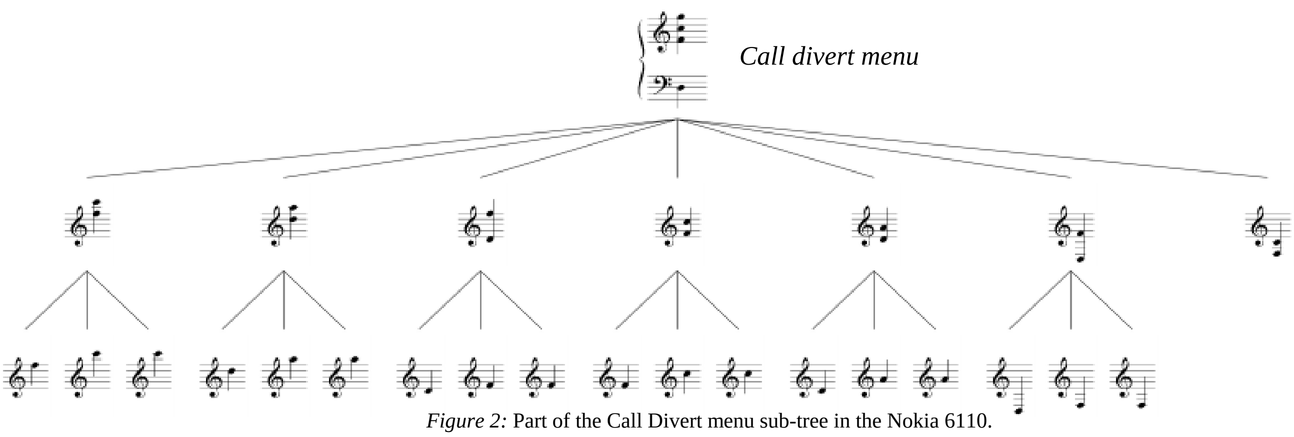

Figure 2: Part of the Call Divert menu sub-tree in the Nokia 6110.

sounds that we designed into some real phones to allow a more realistic field trial of our ideas. The sounds are much simpler than those in experiment 3 as the Nokia 8210 phone only has a monophonic square-wave tone generator with a very limited pitch range. The sounds were developed with our design tool (see below). We concentrated on the melodic, harmonic and rhythmic aspects of sound for our design (reusing ideas from our second experiment [19]). Our principles were refined when creating these sounds to deal with this worst-case design scenario. The sounds are limited because of the hardware but it will be a very effective way of testing users reactions to the sounds in real usage (this is only a limit with current hardware – new phones that are coming out now incorporate MP3 players and so would have the capabilities to play the more complex sounds from experiment 3). Due to the problems we had with industrial collaborators, Nokia is currently conducting the evaluation in Finland. Users have been given phones to use in their everyday lives and we are collecting a range of measures from time taken to navigate to user preference for the sounds (using a standard questionnaire from Nokia). A paper has been submitted to ICAD2001 in which we discuss the sound design and will present the final results [32].

Earcon design tool

After discussions with sound designers, Nokia and from our own work, we decided that a tool was needed to help create hierarchical earcons. Current tools (such as sequencers) do not support the creation of hierarchical earcons. A change in one part of the hierarchy of sounds is not reflected in the sounds below it. This is particularly important if the hierarchy has many levels and nodes (the Nokia 6110 had 150 sounds in its menu hierarchy); if a change is made at the top of the hierarchy then changes must be made by hand to all of the sounds below. This makes a standard sequencer an impractical tool for building complex hierarchies. We realised that we needed to create a tool to simplify this process. Our tool works in a similar way to a standard sequencer but allows much more structuring of the sounds ([21] and [33]). It takes in a description of a hierarchy and allows a designer to create sounds for it in a structured way. This tool enforces our design principles so that sounds created are as effective as possible. We have used the tool to create two sets of sounds for the telephone simulator: one using rich sounds and one using simple, square wave sounds (which were used by Nokia for including into a real telephone for final evaluation). The tool was evaluated with sound and interface designers to ensure that it was effective [28].

One other feature of the tool is to provide basic automatic generation of hierarchical earcons [33]. Given a hierarchy it will automatically compose a set of sounds to represent it – we calculate the size of each of the subtrees and construct the sounds needed based on the method described in experiment 3. This is only done in a limited (but functional) way at present. It allows designers who are not sound experts to create sets of sounds, or for sound designers to get an initial set of sounds which they could then modify (so speeding up the sound design process). This needs further development and results from the area of algorithmic composition must be applied to make the results more effective, but currently there is no other earcon design tool that provides this support for design.

Principles for using non-speech sounds in TBIs

We have developed a set of principles and guidelines from our series of designs and experiments that designers can use to help them design interfaces. Here is a brief summary of some of the main points (for full details of these see [33] and [6, 22]). *Use earcons to:*

1. Provide navigation cues when screen space is limited or not available;
2. Increase the perceptual differences between main menu items/sub-trees/families;
3. Cue changes of level within a menu structure;
4. Provide information on the relative position of the current menu item within its menu;
5. Provide information on the size of the branch of which the current node is the root;
6. Keep the sonification homogeneous;

More practical guidelines:

7. Sounds must be short. 1-2.5 seconds is the maximum length a sound can play otherwise users may skip passed it and miss the information is provides;
8. Load as much of the content into the beginning of the sound as users may skip on to the next before the current one has finished;
9. Timbre should be used to create semantically distinct groups at the top level of the menu structure;
10. Tempo changes alone are not distinct enough to provide good cues;
11. Low quality sounds can be used to present the cues but recall rates will be lower;
12. Allowing users some ‘active learning’ time will significantly improve recall;

Compound earcons:

13. Compound earcons can provide very good cues for shallow, wide menus. Avoid them for deep menus as they can cause problems with short-term memory;

SoundGraphs:

14. Sounds can be used to present continuous, time-varying data when screen space is not available;

New interaction techniques

As part of the project we also investigated two other interaction problems. The first was automatic volume control for auditory interfaces [9] (this work was done in conjunction with project GR/L79212). We tried to address a problem that comes up in both of these projects: auditory feedback played at an inappropriate volume. Sounds are sometimes played too quietly so that they cannot be heard or (more frequently) too loudly so that they are annoying (for example loud sounds played by a mobile phone in a quiet environment). By using the microphone of a device we could monitor the ambient sound around it and then automatically control the volume of the sound coming out of the speakers appropriately. We built a simulator to test our design and it proved to be successful at keeping output volume at the correct level [9] but feedback under control. This method could be used in future mobile device interfaces to ensure that audio feedback is at the correct level. It is particularly important to get the volume of the sounds right as using non-speech sounds for navigation means that mobile devices will make more sounds.

The second area investigated was using sound to present information that was not hierarchical but continuous and time varying. One limitation of many current mobile devices is that they present mainly static information – address books, to do lists, etc. Presenting dynamic information would make devices more useful – users could access a whole new range of services when on the move. For example, using a mobile device on a wireless network it would be possible for stock market traders to keep track of stock trends when away from the trading floor. The problem is that getting access to such information is currently very difficult to do on mobiles because visual displays are so small that standard line graphs are not effective. Our novel approach was to sonify the data based on the idea of *SoundGraphs* [24] – a technique originally developed for people with visual disabilities. We built a simulator that allowed us to sonify stock market data [14]. The price of a share was mapped to the pitch of a note so that a change in pitch over time represented the trend of the market. The results showed that earcons were a more effective way of presenting price trends than a standard visual line graph [14].

PROJECT PLAN REVIEW

I broadly kept to the plan as specified in the original proposal. The main change was that our industrial partner, Telecom Sciences Ltd, went into receivership just as the project started. Therefore our contacts with them stopped and we did not receive the equipment or engineer support that they had promised. This was a problem as I had put a focus on evaluating our ideas with them throughout the project. After some searching I made contact with George Clelland of IBM Hursley Park. He loaned us an IBM voicemail system and network simulator to work with. IBM also trained a Masters student in how to operate and manage the equipment. After several months and several experiments with this kit it proved unsuitable for our work. I later made contact with Seppo Helle from Nokia Mobile Phones. Helle is interested in the use of sound in mobile phones. This proved to be a fruitful contact as Nokia provided several mobile phones for us to use and access to their audio engineers. This changed the focus of the project slightly; we were no longer working on dial–up TBIs, but on interfaces to mobile phones themselves. We had discussed the latter in the original project proposal so the change in direction was slight. Most mobile phones have very complex menu structures (much more complex than a dial-up TBI so provided more of a challenge). This did not have a great effect on the work we did during the project, as many of the problems were the same. In fact, due to the rapid growth in interest in mobile devices of all kinds this has proved to be a very timely and fruitful area for research.

The fact that we did not get Nokia involved as an industrial partner until later in the project meant that we are only now doing the final real-world evaluation of our design (a paper describing this work has been submitted to ICAD2001 [32]).

RESEARCH IMPACT AND BENEFITS TO SOCIETY

During the time of the project there has been an explosion in the use of mobile devices such as mobile phones and the UK and Europe are leading research and development in this area. The project proved to be very timely as there are many interaction problems when navigating around the complex menu structures that are commonly used in such devices and we have provided a novel solution. We have developed new interaction techniques and shown that they can significantly improve usability for users of mobile phones. This is an important benefit as such devices are very common. One advantage of our ideas is that, in a basic form, they can be implemented using current technology so that other companies in the UK could use them.

This work has a strong benefit for both developers of mobile phone interfaces and users of such interfaces. Developers can now use non-speech sounds to aid in navigation through their interfaces knowing that these can make significant improvements to usability. There is also a tool for them to use to speed up the develop of the sounds needed. This also provides some automatic generation tools for designers who do not have experience in sound. For users this work means that future TBIs will be easier to use and they will be able to make better use of the functionality available in the phones. It will also be possible for them to use a range of different services (such as stock market monitoring) that are not possible to do with current devices.

During the project we have had good collaboration with Nokia. They have been to visit my group several times to see our research and discuss ideas. Leplâtre spent a week visiting the audio researchers at Nokia. He gave a presentation on the tools we have developed and then did extensive interviews and walkthroughs of our software so that we could get feedback from our industrial partner (details of this evaluation will appear in Leplâtre's thesis [33]). Nokia gave us access to their sound designers as we were keen to make sure our tools would support the kinds of problems faced by designers in the real-world.

As a result of our collaboration with Nokia (and contacts made at a workshop we hosted as part of the project) I now have a visiting researcher, Dr Antti Pirhonen from the University of Jyvaskyla in Finland, working with me for one year. His research project (partly funded by Nokia) is investigating evaluation techniques for mobile devices with non-speech audio interfaces. He is working with Leplâtre to design the sounds for his experimental prototypes.

The earcons and principles developed here have been taken on into EPSRC grant GR/M44866. This project is looking at presenting complex information (such as graphs, tables and 3D plots) to blind people. We will use the tool developed to create the earcons we need for this project. We are already using the hierarchical earcons and SoundGraphs developed.

I authored a detailed paper on the work in the project that was published in the ACM journal Transactions on CHI. This is one of the key journals in the field and so brought the work to a wide audience. Directly from this publication I was contacted by the authors of a book on speech-based TBI design [3] to include the details of our work. This was the first book in the area and therefore a key one for our research to feature in.

Leplâtre has gained a wide range of skills during the project. He developed skills in general interface design for mobile devices, sound/earcon design, programming, experimental usability evaluation and project management (of masters and undergraduate student projects). All of these are important for him because of the range of research and development into the design of mobile devices that is going on in the UK and Europe at present. He how has the skills to make an impact in this area.

EXPLANATION OF EXPENDITURE

Funds were spent according to the project plan. I did redirect a limited amount of money from other categories in the grant to make up the project student's living expenses so that they matched the increment that was paid by the EPSRC to other PhD students. This was agreed in advance with the EPSRC.

Due to confusion in the assignment of funds to categories in the original proposal for the project (this being my first proposal) I put items that should have come under consumables into the equipment category. During the project I therefore overspent on the consumables category but paid for it from the equipment one. Overall, the spending on the two categories matched what I had put into the original project proposal.

We did not get any of the agreed contributions from Telecom Sciences but IBM and Nokia provided equipment that allowed us to carry on the research.

FURTHER RESEARCH/DISSEMINATION ACTIVITIES

The results from the project have been published widely (we have published 11 papers in conferences, journals and workshops. Leplâtre's thesis [33] will be submitted in spring 2001). We have particularly focussed on the ICAD series of conferences at this is the main place to present work on auditory interfaces. These conferences get a good mixture of academic and industrial participation so were a good place for us to present and discuss our ideas.

The prototypes developed have been put on the Web for distribution (see http://www.dcs.gla.ac.uk/~stephen/research/telephone). We have details of the project, principles developed, our audio design tools and a selection of the sounds we have produced. These have been downloaded by a wide range of researchers from other universities and companies (for example, David Walker from Philips UK has used the Nokia 6110 simulator and sound design tool).

In the first year of the project we presented our initial findings at the First International Workshop on Human-Computer Interaction with Mobile Devices held at Glasgow [13]. This was one of the first meetings to consider the usability aspects of mobile devices and we were able to present our research to a group of people who had not considered non-speech sound as a possible solution to menu navigation problems. It was from this meeting that I started to make contacts with Nokia which later turned into a more significant collaboration on the project.

In 1998 I co-chaired the ICAD'98 conference in Glasgow ([12], http://www.dcs.gla.ac.uk/icad98). Initial results from our work were presented here [19]. This allowed us to discuss the ideas with many academic and industrial researchers with an interest in sound (companies from Europe, USA and Japan attended the conference).

I was the co-chair of the Second International Workshop on Human-Computer Interaction with Mobile Devices. This workshop was held as part of the IFIP INTERACT'99 conference in Edinburgh ([10, 11] and see http://www.dcs.gla.ac.uk/mobile99). I used the workshop funding requested for the project to finance this meeting. The workshop was successful – we had over 75 attendees, with many different companies represented, for example Nokia, Siemens and Erickson from Europe and Tegic from the USA. The proceedings of the workshop were published in a special issue of Personal Technologies [11] which I co-edited. This work allowed us to discuss the ideas from the project with other key academics in the area and also to build industrial links with companies (in particular Nokia).

I presented details of the research at an invited seminar at Sun Microsystems in 1998. More informally, an article appeared in the Scotsman (22nd April, 1998) describing the project, which exposed the ideas to a much more general audience.

As this was my first fast-track project I wanted to use it to allow me to develop my connections with industry. The connections I now have with Nokia are strong and we are continuing to work in this area (and also in another joint project on 3D sound) now that this project has finished. Nokia are keen to fund a follow-on project at Glasgow to develop the tools we have produced. We are currently finalising the plans for this, which will start in spring 2001 (when Leplâtre has submitted his thesis) and will employ Leplâtre for at least one more year to continue this research.

REFERENCES

1. Alty, J.L., Vickers, P. and Rigas, D. Using music as a communication medium. In *Extended Abstracts of ACM CHI'97* (Atlanta, GA) ACM Pressy, 1997, pp. 30-31.

2. Alty, J.L. Can we use music in human-computer interaction? In *Proceedings of BCS HCI'95* (Huddersfield, UK) Cambridge University Press, 1995, pp. 409-423.

3. Balentine, B., Morgan, D. and Meisel, W. *How to Build a Speech Recognition Application*. Enterprise Integration Group, 1999.

4. Blattner, M., Sumikawa, D. and Greenberg, R. Earcons and icons: Their structure and common design principles. *Human Computer Interaction* 4, 1 (1989), 11-44.

5. Brewster, S.A. Navigating telephone-based interfaces with earcons. In *Proceedings of BCS HCI'97* (Bristol, UK) Springer, 1997, pp. 39-56.

6. Brewster, S.A. Using Non-Speech Sounds to Provide Navigation Cues. *ACM Transactions on Computer-Human Interaction* 5, 3 (1998), 224-259.

7. Brewster, S.A. Using Earcons to Provide Navigation Cues in Telephone-Based Interfaces. *ACM Interactions* 6, 2 (1999), 9-10.

8. Brewster, S.A., Capriotti, A. and Hall, C. Using compound earcons to represent hierarchies. *HCI Letters* 1, 1 (1998), 6-8.

9. Brewster, S.A., Crossan, A. and Crease, M. Automatic volume control for auditory interfaces. In *Volume II Proceedings of BCS HCI2000* (Sunderland, UK) British Computer Society, 2000, pp. 17-18.

10. Brewster, S.A. and Dunlop, M.D. Second Workshop on Human-Computer Interaction with Mobile Devices. In *Volume II of the Proceedings of INTERACT '99* (Edinburgh, UK) British Computer Society, 1999, pp. 133.

11. Brewster, S.A. and Dunlop, M.D. Human computer interaction with mobile devices. *Personal Technologies* 4, 2 & 3 (2000), 71-72.

12. Brewster, S.A. and Edwards, A.D.N., Eds. *Proceedings of ICAD'98*. British Computer Society, Swindon, 1998.

13. Brewster, S.A., Leplâtre, G. and Crease, M.G. Using non-speech sounds in mobile computing devices. In *Proceedings of the First Workshop on Human Computer Interaction for Mobile Devices* (Glasgow, UK) Department of Computing Science, University of Glasgow, 1998, pp. 26-29.

14. Brewster, S.A. and Murray, R. Presenting dynamic information on mobile computers. *Personal Technologies* 4, 4 (2000), 209-212.

15. Brewster, S.A., Raty, V.-P. and Kortekangas, A. Earcons as a method of providing navigational cues in a menu hierarchy. In *Proceedings of BCS HCI'96* (London, UK) Springer, 1996, pp. 169-183.

16. Foresight. *ITEC Group Technologies Document*. Foresight ICM Panel, 1999, Report.

17. Howe, A. A model of the acquisition of menu knowledge by exploration. In *Proceedings of AM CHI'94* (Boston, MA) ACM Press, Addison-Wesley, 1994, pp. 445-451.

18. IEGoMP. *Mobile phones and health*. Independent Expert Group on Mobile Phones, 2000. http://www.iegmp.org.uk/

19. Leplâtre, G. and Brewster, S.A. An investigation of music to provide navigation cues. In *Proceedings of ICAD'98* (Glasgow, UK) British Computer Society, 1998.

20. Leplâtre, G. and Brewster, S.A. Perspectives on the design of musical auditory interfaces. In *Focus conference on Anticipation, Cognition and Music* (Liege, Belgium) International Journal of Computing Anticipatory Systems, 1998.

21. Leplâtre, G. and Brewster, S.A. Improving the Design of Telephone-Based Interfaces. In *Volume II of the Proceedings of INTERACT '99* (Edinburgh, UK) British Computer Society, 1999, pp. 45-46.

22. Leplâtre, G. and Brewster, S.A. Designing Non-Speech Sounds to Support Navigation in Mobile Phone Menus. In *Proceedings of ICAD2000* (Atlanta, USA) ICAD, 2000, pp. 190-199.

23. Maguire, M. A human-factors study of telephone developments and convergence. *Contemporary Ergonomics* (1996), 446-451.

24. Mansur, D.L., Blattner, M. and Joy, K. Sound-Graphs: A numerical data analysis method for the blind. *Journal of Medical Systems* 9 (1985), 163-174.

25. Norman, K. *The psychology of menu selection*. Ablex Publishing, Norwood, NJ, USA, 1991.

26. Rosson, M.B. Using synthetic speech for remote access to information. *Behaviour Research Methods, Instruments and Computers* 17, 2 (1985), 250-252.

27. Schumacher, R.M., Hardzinski, M.L. and Schwartz, A.L. Increasing the usability of interactive voice response systems. *Human Factors* 37, 2 (1995), 251-264.

28. Wilkie, K. *Sequence and Motif Editing Environment. The implementation of an interactive GUI for an audio language.* MSc Thesis, University of Glasgow, 2000.

29. Williams, D., Cheepen, C. and Gilbert, N. Experiments in how automated systems should talk to users. In *Proceedings of BCS HCI'98* (Sheffield, UK) British Computer Society, 1998, pp. 44-45.

30. Wolf, C., Koved, L. and Kunzinger, E. Ubiquitous Mail: Speech and graphical interfaces to an inte-

grated voice/email mailbox. In *Proceedings of Interact'95* (Lillehammer, Norway) Chapman & Hall, 1995, pp. 247-252.

31. Yankelovich, N., Levow, G. and Marx, M. Designing SpeechActs: Issues in speech user interfaces. In *Proceedings of ACM CHI'95* (Denver, Colorado) ACM Press, Addison-Wesley, 1995, pp. 369-376.

Items from the project not yet published

32. Helle, S., Leplâtre, G., Laine, P. and Marila, J. (2001). Menu sonification in a mobile phone – a prototype study. *Submitted to ICAD2001* (Helsinki, Finland).

33. Leplatre, G. (2001). *The design and evaluation of non-speech sounds to support navigation in restricted display devices.* PhD thesis, Glasgow University, UK. To be submitted Spring 2001.