Guidelines for the Use of Sound in Multimedia Human-Computer Interfaces

Stephen A. Brewster & Philip D. Gray
Department of Computing Science, University of Glasgow

PART1: INTRODUCTION

There is a small but growing body of evidence which indicates that the addition of non-speech sound to human-computer interfaces can significantly improve performance and increase usability [1, 19, 20, 25]. Our own work [14, 16] has shown that adding sound to a graphical interface can significantly reduce both the time taken to complete tasks and the time taken to recover from errors.

Sound is a very important means of communication in the everyday world and the benefits it offers should be exploited at the computer interface. Sonically-enhanced graphical interfaces allow more natural communication between the computer and the user. They allow the user to employ more appropriate senses to solve a problem, rather than using vision alone to solve all problems. However, because this area is still in its infancy, there is little systematic research to demonstrate the best ways of combining these different media [7]. This means sounds are often added in *ad hoc* and ineffective ways by individual designers [4, 29]. Arons & Mynatt [3] describe one effect of this: "…the lack of design guidelines that are common for the creation of graphical interfaces has plagued interface designers who want to effectively build on previous research in auditory interfaces".

In spite of this lack of research, industry's interest in sound and multimedia has increased, with most computer manufacturers now including sound producing hardware in their machines. The UK Technology Foresight Programme also identified multimedia and sound as specific technology opportunities available to the UK. It also made it clear that user needs, and in particular ease of interaction, were priorities. This research will study multimedia human-computer interfaces and investigate the possibilities for improving interaction using sound. The knowledge gained from this research will be formalised in a set of guidelines that will overcome the problems suggested by Arons & Mynatt.

Some systems, such as Microsoft Windows95, allow users (rather than designers) to add sounds to some interface events. Again this is an *ad hoc* use of sound; There is no guarantee that the sounds will be anything more than gimmicks and so are unlikely to improve usability. Sonic feedback may be added in ineffective ways, without care for the annoyance that it can cause. It is therefore important that the use of sounds is properly investigated, so we can ensure that problems of poor graphic interface design, which occurred before graphical interface guidelines became available, are not repeated in sound.

The innovative task in this proposal is to produce a set of guidelines, and a toolkit based on it, to facilitate the use of sound so that designers can improve the usability of their multimedia interfaces. The guidelines and toolkit will offer five advantages:

- Simplify the design of sonically-enhanced interfaces;
- Allow designers who are not sound experts to create sonically-enhanced interfaces;
- Ensure, by experimental evaluation, that any sounds added are effective and improve usability;
- Ensure sounds are used in a clear, coherent and consistent way across the human-computer interface.
- Improve the usability of interfaces for partially-sighted users (e.g. users who have problems due to disability, protective clothing or working environment).

The lack of research on the use of sound means that we must experimentally evaluate the most effective places to use it before the guidelines can be created. Therefore, we are applying for one research assistant and one research student for the project. We will do the evaluations from the bottom up by investigating graphical widgets (interface components such as buttons and menus) to solve their usability problems with sound. The results gained from these experiments will form the basis of the guidelines. The next stage will be to enhance an existing graphical toolkit, e.g. Java's Foundation Classes, to include sound. This will make it easy for designers to create interfaces that embody the guidelines.

In addition, we will also investigate designs for new widgets that fully integrate both graphics and sound (rather than just using sound to correct problems with existing widgets). They will be designed from scratch to take advantage of both forms of output. These widgets will form the basis for more usable interfaces in the future.

Our previous research provides a solid foundation for this project [14-16]. This unique approach has attracted international attention and has been published at the major HCI conferences [11, 14, 16]. Two papers have been published describing the proposed general structure of a sonically-enhanced toolkit [8, 9] and a more detailed description accepted elsewhere [10]. Our progress will be measured against the following milestones:

- First year: Develop a preliminary set of guidelines from experimental results. Demonstrate design of the toolkit of sonically-enhanced widgets. Produce report.
- Second year: Produce final set of guidelines. Build the complete toolkit and make it available to designers. Produce designs for new widgets.
- Third year: Demonstrate effectiveness of guidelines/ toolkit by evaluation with interface designers. Demonstrate with experimental results the new fully-integrated widgets. Produce final report. Host workshop to disseminate knowledge gained.

By using the guidelines and the toolkit auditory interface development will be simplified and designers will no longer have to add sounds in *ad hoc* ways.

BACKGROUND

The use of structured non-speech sounds at the human-computer interface is a new research field. In the short time it has been under investigation some significant benefits have been identified [7, 20, 25]. One reason for this is that sound combined with graphics can significantly improve usability by taking advantage our natural ability to share tasks across sensory modalities. The impact of this project will be to establish guidelines for the effective use of sound so that interface designers can easily include sound into their interfaces. This will allow other researchers and industry to use the advantages of sound in new interface designs. The research is novel because currently no other group is systematically studying the use of non-speech sound to improve graphical human-computer interfaces.

One motivation for this work is that users' eyes cannot do everything. The visual system has a small area of high focus. If users are looking at one part of the display then they cannot be looking at another part at the same time [7]. Many interface errors occur because users fail to interact with a widget correctly. Their visual attention is focused on the information they are interested in and not the widget [10]. This problem is even worse for partially-sighted users whose area of focus is reduced. It is suggested here that some feedback from the widgets should be presented in sound. Sound has certain advantages: It can be heard from all around, it is good at getting our attention whilst we are looking at something else and it does not disrupt our visual attention. Presenting feedback in this way would allow users to continue looking at the information they require but to hear information that they would otherwise not see (or would not see unless they moved their visual attention away from the area of interest, so interrupting the task they are trying to perform). One important research problem is to find the situations where the errors occur and to design sounds to overcome them. Sound and graphics can then be used together to exploit the advantages of each to overcome the feedback problems.

This proposal will build upon the results of researchers working in related areas:

Meera Blattner, Lawrence Livermore National Laboratory, initially developed the idea of structured audio messages called *earcons* [5, 6]. These are abstract, musical sounds built up from smaller units that can be manipulated to build complex structures [12]. Blattner proposed methods for constructing earcons but never evaluated them. As part of his PhD research, Brewster performed a detailed investigation of earcons with very promising results [15]. Earcons will constitute the majority of the auditory feedback used in the project. Professor Blattner and Stephen Brewster are currently working together on a book about non-speech audio.

Figure 1 shows, as an example, a simple hierarchy of earcons based on one possible family of errors. Each earcon is a node on a tree and inherits the properties of the earcon above it. The different levels are created by manipulating the parameters of sound (for example, rhythm, pitch, timbre). In the diagram the top level of the tree is a neutral earcon. It has a flute timbre. The structure of the earcon from level one is inherited by level two and

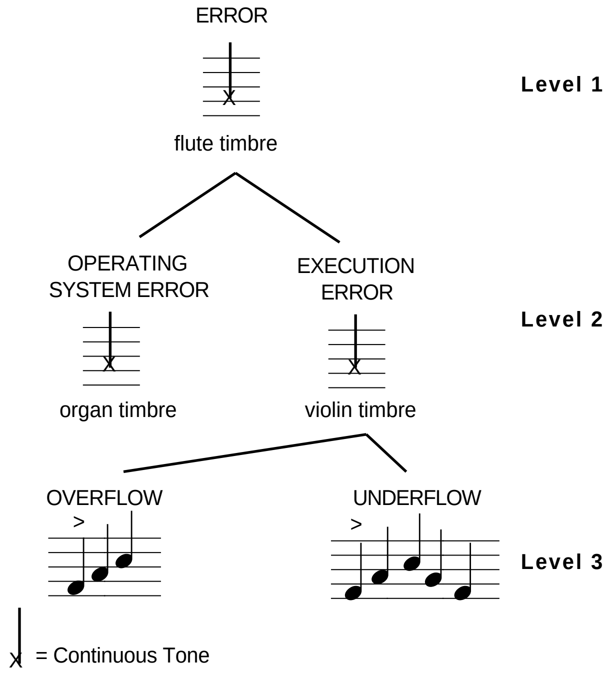

Figure 1: A hierarchy of earcons representing errors.

then changed. At level two there is still a continuous flute sound but new timbres are added to play alongside it. At level three a rhythm is added to the earcon from level two to create a sound for a particular error. This rhythm is based on the timbre from the level above. In the case of the overflow error there would be a continuous flute sound with a three note rhythm played on an organ accompanying it.

Using earcons, this hierarchy is easily extensible. For example, to add another major category of errors all that is needed is a new timbre. To create a new type of execution error only a new rhythm is needed and it can be added to the existing hierarchy. Therefore earcons provide a very flexible system for representing structured information.

Alistair Edwards, University of York, has done much work in the field of improving interfaces for blind and partially-sighted people [21, 22]. For example, he successfully developed a completely auditory interface for a wordprocessor [21]. In recent research, Brewster worked with him and his colleagues on a system for presenting mathematics to blind people [31]. This system used earcons to provide a 'glance' at the overall structure of an algebraic expression. This gave some idea of the complexity of the expression before users started to browse it with synthetic speech. The glance allowed users to choose the appropriate browsing strategy so that they were not overwhelmed with synthetic speech. Many of the principles developed from this work and from the design of other systems for blind and partially-sighted people will be applied to the design of the toolkit described here.

Bill Gaver, Royal College of Art, has developed an alternative method for presenting information in sound called *auditory icons*. He used these to improve interfaces for collaborative working [24] and process control systems [25]. Auditory icons are based on everyday sounds that have an intuitive link to an operation or action in the interface. Some of the strengths of auditory icons (such as their naturalness) will be used in the design of the sounds for this project.

James Alty, University of Loughborough, has used music to enhance the display of algorithms and sorting routines [1]. Results of his experiments on algorithm sonification were successful and the design considerations he came up with will be used in this project. Alty's previous work also dealt with some problems of combining different sensory modalities [2]. He worked on a multimedia process control system that tried to use the appropriate sensory modality to present information to the operator. One of the problems in this work was that it was not possible to specify when to use different modalities; not enough was known about when to use the different senses. Some of the problems he found will be addressed by the set of guidelines to be produced as part of the project.

There has been a little work on the development of toolkits using sound in interfaces for blind users [27, 30]. Savidis & Stephanidis [30] built a toolkit that allowed visual and non-visual representations of the same interface. This meant that sighted and non-sighted users could share the same application by using different interfaces. This work did not address the benefits of adding sound to the visual interface but the toolkit developed will be analysed to examine how and where sounds were added. We will build upon these ideas in our toolkit.

There is much work in the area of alarms and warnings, but as yet little of this has been applied to user-interface design. In particular, Roy Patterson [28] and Judy Edworthy [23] have put forward guidelines for the design of alarms. These include ways of controlling the annoyance and urgency of auditory feedback. These guidelines will be applied to the design of the sounds used for the project to make sure that they are attention grabbing when necessary.

Brewster has performed the most in-depth studies of the use of structured audio messages called earcons [7, 15]. He also produced a very simple set of guidelines for designers to enable them to create effective earcons [17]. These guidelines were an initial attempt at making the use of sound easier, but they were only the first step. They helped in the design of the earcons themselves but not in how to use them to improve interaction.

We have experimentally tested the use of earcons in buttons, scrollbars and menus and found that, in each case, usability was significantly improved: The time taken to recover from errors and the time taken to complete tasks were reduced along with reductions in subjective workload [10, 11, 14, 16]. There were also significant increases in user preference for the sonically-enhanced widgets without any increase in annoyance. These initial results have been very promising. We now want to carry this further with this current project.

The other related strand of our work has been the use of sound to improve the usability of telephone-based interfaces and interfaces for blind and disabled people [12, 13]. We have experimentally tested earcons for navigation cues in non-visual interfaces. They were very successful with participants being able to work out their location in a hierarchy of information easily. Brewster currently has EPSRC project GR/L66373 in collaboration with Telecom Sciences Ltd. investigating this area. Results from this other project will feed into this research to help us design better earcons.

PROGRAMME AND METHODOLOGY

As described, there are some experimentally-proven advantages to using sound at the interface. However, most previous examples were *ad hoc* solutions because the use of sound to improve human-computer interfaces is a new research area. Therefore, there was little to build upon. One major problem holding up the acceptance and use of sound by industry and academia is that there are no guidelines to aid a designer in how to use sound effectively [3]. We will therefore undertake the first systematic evaluation of the interface to find out how sound can be used effectively.

Some systems do allow sounds to be added to certain interface events. These are often added in an *ad hoc* way by the user of the system, not the designer. There is no guarantee that the sounds will be anything more than gimmicks and are therefore unlikely to improve usability. As more systems that include sound become available, the more conflicting uses of sound there will be. It is therefore important that this project systematically investigates the use of sound before there is a cacophony of sound, just as there was garish use of colour before guidelines for graphical interface design became available [26].

The innovative task in this proposal is to produce a set of guidelines, and a toolkit embodying the guidelines, so that designers can improve the usability of their multimedia interfaces by the addition of effective audio feedback. There are five main aims for this research:

- *To simplify the design of sonically-enhanced interfaces*. Currently it is difficult to create sonically-enhanced applications. A set of guidelines would make the inclusion of sound much simpler. A toolkit built on top of the guidelines would also greatly simplify implementation. Myers [26] suggests that the development of graphical interfaces was greatly simplified by the introduction of graphical toolkits. Our toolkit will do the same for sonically-enhanced interfaces.
- *To allow designers who are not sound experts to create sonically-enhanced interfaces*. Interface designers are often unskilled in sound design. A set of guidelines and toolkit would reduce the need for knowledge of sound design.
- *To ensure, by experimental evaluation, that the sounds added are effective and improve usability.* The sounds added will not be gimmicks. Detailed investigations of usability problems will show where sounds can be beneficial.
- *To ensure sounds are used in a clear, coherent and consistent way across the interface.* This consistency will avoid the problems of each application having its own sounds that mean different things in other applications. In graphical interfaces, widgets look consistent across different applications, e.g. a scrollbar looks the same in any application. By using the guidelines designers can ensure widgets sound consistent across different applications.

- *To improve the usability of graphical interfaces for partially-sighted users* (e.g., users who have problems due to disability, injury or protective clothing). Such users often have difficulty with high-resolution graphical interfaces where text and graphics may be very small. By testing our designs on partially-sighted users we will ensure the widgets are usable by them.

There are four main stages to the work described in the proposal: 1. design and test sonic enhancements of standard widgets; 2. from this work, create the guidelines for where to use sound and then build a toolkit of sonically-enhanced widgets based on the guidelines; 3. design and test new widgets that use sound and graphics as equal partners; 4. implement systems based on the toolkit to test its effectiveness. Reports will be produced after each of these stages and can be used as milestones to measure our progress through the project.

First year

The first year will focus on investigations of how to sonically enhance standard widgets in order to produce the guidelines, and the design of the toolkit. The research student (RS) will design and implement experiments to test the sonically-enhanced widgets. Some aspects of the first stage have been met by our previous work. Brewster and colleagues have built and tested sonically-enhanced buttons, scrollbars and menus [11, 14, 16]. Results from these experiments were very promising. However, these widgets still need further development to include solutions to problems identified after experimental analysis. Some other widgets that must be evaluated are: Tool palettes, dialogue and alert boxes, valuator bars, windows, and drag and drop. We have considerable experience in designing and conducting experiments so will help the RS design the experiments quickly and effectively. Each investigation will take approximately three months (based on the time taken for our previous experiments).

The study of standard graphical widgets will take the following form. The widget must first be analysed to find out any potential usability problems. An analysis technique that Brewster developed as part of his PhD study can be used by the RS for this [7]. It considers interactions in terms of event, status and mode information. It identifies areas where the feedback from the widget will not provide information that the users need. The technique will be refined and improved by the RS. The next stage is to design the auditory feedback necessary to solve the usability problems. An experimental design must then be produced and any necessary software written. The sonically-enhanced widgets must then be evaluated to discover if usability has been improved. The experiments will be based on a testing framework that we have designed and evaluated [7]. The design of the experiments will follow that shown in Figure 1. A two-condition, within-subjects repeated-measures design will be used to test all of the widgets. This allows each participant to be used twice and their results compared. In one of the conditions the standard graphical widget will be tested, in the other condition the sonically-enhanced widget. This design has been successful in testing previous widgets [14, 16]. Participants will be drawn from both Computing Science and other areas to get a full range of users. We will also make sure we test a range of partially-sighted users so their particular problems are addressed. Quantitative measures of time taken to complete tasks and time to recover from errors will be taken. These will be accompanied by qualitative measures of subjective workload, overall preference and annoyance. We will also video participants to collect qualitative data. Combining quantitative and qualitative data gives a full measure of the usability of a system. One of the main concerns potential users of auditory interfaces have is annoyance due to sound pollution. This will be explicitly measured to make sure it is not a problem. Results from the experiments will be statistically tested to ensure usability has been improved. If any problems are found the widgets will be re-designed to overcome them. The results of these experiments will be a set of sonically-enhanced widgets that demonstrably improve usability. A report will be produced describing the experiments completed and guidelines established from the results.

Participants	Condition 1		Condition 2	
First half of participants	Sonically-enhanced Widget Train & Test	Workload Test	Visual Widget Train & Test	Workload Test
Second half of participants	Visual Widget Train & Test		Sonically-enhanced Widget Train & Test	

Figure 1: Format of the experiments to test the sonically-enhanced widgets.

The design of the sounds and the experiments will give the RS a good background necessary for he skills he/she will need for the rest of the PhD. We will provide guidance on the experiments but the RS will have the flexibility to use his/her own creativity to find novel solutions.

During the first year, the research assistant (RA), Brewster and Gray will begin to design and build the overall structure of the toolkit (based on our previous work and the widgets previously tested). The general structure proposed [8, 9] must be refined so that the individual widgets can be plugged into it when they have been tested. It must be flexible enough to handle the different sound requirements of the different widgets and also allow sounds to be attached to the events that are required. The RA will work together with the RS to ensure that the toolkit structure satisfies the widgets' need for sounds.

The toolkit will be implemented as an enhancement to the Java Foundation Classes (Java's interface toolkit). The code for these are available free for academic use. This approach has the advantage that the toolkit will work on many different computing platforms and therefore be available to many different users. This will also aid in the dissemination of the work when complete. We will build on the standard Java widget classes and add our own code to generate the sounds necessary. Developers will continue to use the widgets in the same way as before (by making calls to the toolkit via functions) but they will now make sounds as well as display graphics.

One possible structure of the sounds is as follows: Each application will have its own timbre as a base for all of its sounds [8, 9]. All widgets within an application will use this and modify it by changing the rhythm, pitch, etc. Figure 2 shows such a hierarchy. At level 1, the three

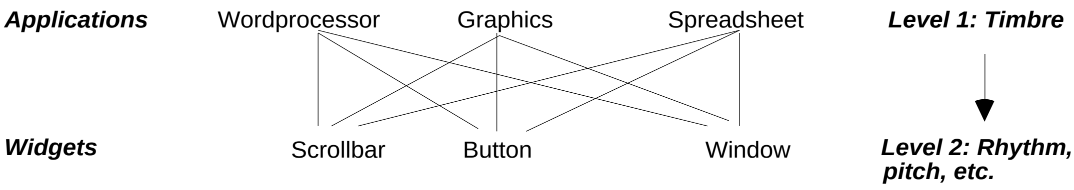

Figure 2: A hierarchy of sonically-enhanced widgets.

applications all have different timbres. These are inherited by level 2 and modified with pitch, rhythm, etc. These modifications are constant across applications so that widgets in different applications sound consistent. Other designs will also be evaluated.

We will use a layered structure like toolkits for the X Windows system [26] which will allow us to put basic audio functions in an 'intrinsics' library. These will then be available to other toolkits (including the new widgets to be designed by the RS in the second year). This will also allow us to investigate the possibility of an audio 'server' that can provide sounds needed by machines that do not have the necessary sound generating hardware themselves. This would work across a local network (for example, our local ATM network using Java's Remote Method Invocation tools) with sound producing hardware at a central location servicing clients requests for sound. A report will be produced describing the toolkit structure.

Once the overall structure has been defined, the RA will work together with the RS to help develop and implement the experimental designs that will be needed to test the sonically-enhanced widgets. This will speed up the process of testing the widgets.

Second year

The RS will continue testing any remaining standard widgets. Brewster and Gray will aid in the design of the experiments and the RA will aid in the implementation so that the research progresses quickly. Once the tests of the standard widgets are complete, the RS and Brewster will produce a report describing guidelines for the integration of sound into standard graphical widgets. We will also incorporate our guidelines into a standard set of HCI guidelines such as Apple's Human Interface Design Guidelines. The RS will then begin the next phase of the research: Designing the new widgets. Current widgets rely almost entirely on graphical output. One reason for this is that most of them were designed when only graphics was available. One novel aspect of this project will be the design of new widgets that make greater use of sound. These will not simply be existing widgets with sound to correct usability problems but they will be new designs that fully integrate sound and graphics as equal partners from the beginning. These new widgets will improve interaction in desktop computer interfaces because of their increased usability. They will also increase usability in systems with small screens (e.g. notebooks and PDA's) because they will not require as much screen space as feedback will be presented in sound. They will also improve systems for people with partial sight or those who work in situations where screens cannot easily be used, again as feedback will be presented in sound.

The RS and Brewster will work on designs for these new widgets throughout the second year. The RS will have gained the necessary skills in earcon design and have seen many of the interaction problems from work in the first year. To design the new widgets the analysis technique Brewster developed as part his thesis will be combined with User Action Notation and task analysis. We will look at each widget in turn, analysing and describing the tasks it performs and the information it represents so that we can come up with new designs that integrate sound and graphics to display the information. This work will also pave the way for future projects that will integrate other techniques, such as force feedback, into graphical widgets.

The RA and Gray will build the sonically-enhanced widgets into the toolkit structure designed in the first year. This will be the major stage of implementation and will take up the whole of the second year of the project. The widgets will be designed from the start so that they will work together when combined but there may still be difficulties that must be ironed out. Creating the toolkit will involve work to fit the individual widgets together and to avoid any sonic conflicts, for example when one widget requires a sound to be playing but another requires silence.

As mentioned, we will also begin to look at the possibilities of creating an audio server to provide audio services to machines without the necessary sound hardware. Gray and the RA will test alternative designs to see if this is possible, leading the way for a future grant application to develop this further.

Third year

In the final year the RS will begin to experimentally evaluate his/her new widgets designed in the previous year. The evaluations will be similar to the experiments to test the sonically-enhanced standard widgets from the first year. Brewster and Gray will help the RS with the design of these experiments. We will use the same experimental methods and collect the same types of data as before. Results from these experiments will help us improve the design of the new widgets and will be published as a report. Given time, we will begin to investigate the combination of these new widgets into a toolkit.

The RA will spend the final year doing a full evaluation of the guidelines and the toolkit with interface designers. Sun Microsystems (the creators of Java) have suggested that we might test our guidelines on their engineers (with the current rapid rate of change in this area they did not want to commit to funding at this stage). We will also test them on interface designers working in the Department of Computing Science. We envisage working with designers to enhance products such as JMaker or the HotJava web browser with our toolkit. This will allow us to evaluate the use of the toolkit and guidelines by designers in a realistic context. Any problems found with the toolkit or guidelines will be corrected. These sonically-enhanced products will then be exemplars of what can be done with sound and will be made available over the Internet. We will also be able to test the hypothesis that the toolkit is usable without knowledge of the guidelines on which it is based.

We will work with the designers to conduct field trials to evaluate the effectiveness of the sonically-enhanced versions of their products. We will do this with experiments similar in design to those performed in the first year: a sonically-enhanced application will be tested against a purely visual one. Results should show increased usability for the sonically enhanced applications. Results will be published as a report.

As the final part of this project we will hold a workshop to allow us to communicate our findings to academia and industry. The work will then all be brought together and published as a final report.

RELEVANCE TO BENEFICIARIES

We see three main beneficiaries of this work:

Interface developers will benefit because it will be easier for them to create sonically-enhanced interfaces. They will be able to include sounds simply and effectively by using the guidelines. The toolkit based on the guidelines will make it even easier to incorporate sound so that they will not have to spend a large amount of development time writing code to include sound. Using the toolkit will allow them to improve the usability of their interfaces and this will give interface developers a significant advantage when selling their products.

End-users will benefit because of the increased richness of their interaction. Research has shown that using sound can improve performance and usability. However, because it is difficult for developers to create sonically-enhanced interfaces, users have not been able to take advantage of this. This research will allow the creation of more systems that use sound. This will improve the quality of life of anyone who uses a computer with a standard graphical interface.

Partially-sighted users will also benefit from the use of sounds. High resolution screens can lead to very small text and widgets. For example, on current displays the icons may be less than 1cm high. It can therefore be difficult for people with poor eyesight to use them. Sounds will aid in the identification of widgets, positioning the mouse over them and giving feedback from interacting with them.

DISSEMINATION AND EXPLOITATION

The results of each stage of the work will be presented at national and international conferences so that industry and academia can see the work. Conferences will be targeted that attract delegates from both academia and industry such as BCS HCI, ACM CHI and ACM UIST. The work will also be published in a more detailed form in journals.

The results of the research will be made available on the World Wide Web so that they will be available to all who want them. This an increasingly important method for scientists to communicate ideas and we will take full advantage of it. It is particularly important that this work is presented in an interactive form because the sounds cannot easily be understood by reading about them. Therefore the web is a useful dissemination tool as we can provide interactive demonstrations, including the sounds from the toolkit with the papers that discuss them.

We will provide the toolkit, guidelines and a set of sounds that interface designers can use in their systems. These sounds be in standard formats and therefore usable in many different computing environments and will be download-able over the Internet. Other sounds that we create will be reserved for future commercial exploitation by industry.

We will host a workshop in the last year of the project to communicate our ideas to academia and industry. This will allow us to demonstrate our guidelines, toolkit and new widgets. A final report will be produced that brings together all of the knowledge gained.

JUSTIFICATION OF RESOURCES

The following resources are required for this proposal:

One Research Assistant post for three years is the minimum required to do the design, implementation and testing of the guidelines and toolkit. We envisage employing a computer scientist with knowledge of human-computer interaction and Java programming. He/she will be expected to gain a deep understanding of Java when creating the sonically-enhanced toolkit. We hope to recruit one of our undergraduates.

One Research Student is required to do the design and testing of the sonically-enhanced widgets and to develop the new fully-integrated widgets. The research will provide an ideal training ground for a PhD student. He/she will learn about the design, conduct and analysis of experiments, the use of sound and programming in Java. There will also be great scope for individual research for an able student when designing the experiments. We will recruit Mr Murray Crease, a first class honours graduate from our department last year. He has published a paper in the area and has spent a year in industry so has all of the skills necessary.

The project requires three Macintosh computers. The Macintosh is the standard environment used by sound professionals to design sounds and the best tools for the creation of sounds are available on this platform. They must be powerful machines with multimedia capabilities to be able to create and process sounds. These machines will also be used for the design and development of the toolkit, the experiments to evaluate our widget designs, the production of papers and the analysis of results. The Department can provide the Java development tools required and Unix access if we need to do any development in that environment. We are also requesting printing facilities and backup media to ensure that any data, code, sounds or papers are not lost.

We require good quality MIDI sound modules plus sound-output hardware for each of the machines. Together, these will be used to design and produce the sounds and then present them to experimental participants for evaluation. We are also requesting audio/video logging equipment so that we can record participants in our experiments to allow us to do a full analysis of any usability problems that arise. This equipment will also be used for experiments to be undertaken in EPSRC Grant GR/L66373.

Travel funds are necessary to ensure the adequate dissemination of our work. As the work is based on sound it is very difficult to describe without demonstrations. Presenting the work at conferences is the most suitable way of demonstrating our results to academia and industry. We have also requested funding for a workshop at the end of the project so that our knowledge can be passed on to industry and academia.

PART 2: PREVIOUS RESEARCH AND TRACK RECORD

Stephen Brewster

I been a lecturer in the Department of Computing Science at the University of Glasgow since October, 1995. From September, 1994 until October 1995 I was an ERCIM Research Fellow working at VTT in Finland and SINTEF in Norway. Before that, I worked on my PhD at the University of York. In 1992 I was awarded the Gibbs-Plessey prize for postgraduate research in computer science. I have been on the program committees of APCHI'96 and BCS HCI'95, '96 and '97 and was on the international committee for ACM CHI'95 and CHI'96. I was on the review committee for ICAD'94 and ICAD'96 and will be co-chair of ICAD'98 in Glasgow.

For my PhD research [7] I investigated the design and use of non-speech sounds in human-computer interfaces. I carried out a series of detailed experiments to discover the most effective ways of constructing earcons [15, 18]. These were the first to evaluate the design of earcons. I then used earcons to correct usability errors in standard graphical widgets. Results showed statistically significant improvements in time to recover from errors, time to complete tasks and reduced workload [14, 16]. I extended this research during my ERCIM Research Fellowship spent at two major European research centres, both of which were investigating the uses of sound. I worked on developing earcons as navigation aids in interfaces where graphical feedback is not available (for example, telephone-based interfaces). The use of structured sounds proved to be very successful. Listeners were able to identify their location with a high degree of accuracy [12]. The knowledge gained from these experiments will feed into the work described in this proposal.

Grant funded research

I have recently been awarded an EPSRC grant to investigate Principles for Improving Interaction in Telephone-Based Interfaces (GR/L66373). Telecom Sciences Ltd. are industrial partners in this research. This project focuses on the problems when interacting with telephone-based services such as phone-banking and voicemail. These systems are notoriously hard to use [12]. During this project principles will be developed to allow designers to use non-speech sounds to provide navigation cues to stop users becoming lost, a major problem in such systems. Initial results have been promising [12].

Philip Gray

Philip Gray has been actively involved since 1984 in research into user interface architectures and user interface development tools. As an investigator on the Druid Project (1987-90) he helped develop Chimera, one of the first dynamically reconfigurable user interface management systems. Subsequently, he was investigator on the Innovative Iconic Interfaces project (1989-93) where the Chimera linkage component formed the core of the Iconographer visualisation system. As the result of collaboration with Dr William Gaver, the Iconographer system also featured an early version of realistic auditory icons. The Representer System, a successor to Iconographer developed by Mr Gray, includes more sophisticated auditory icons which use the ENO sound server of Dr M Beaudouin-Lafon. Mr Gray was also a member of the Esprit-funded FADIVA Working Group (1994-96) which focused on research into advanced information visualisation systems.

Grant funded research

Dynamically Reconfigurable User Interface Design (DRUID) (SERC Grant GR/D 80124), 1987-1990, with Dr A Kilgour.

An Environment for Creating Innovative Iconic Interfaces (SERC Grant GR/F 67129), 1989-1993, with Dr S W Draper (Psychology Dept, Glasgow University).

Rubicon (Digital Equipment Corporation). 1991-1994.

Temporal Aspects of Usability (JCI in Cognitive Science and HCI Grant G9201233), 1992-1995, with Dr D England, Dr C Johnson, Dr S Draper, Mr P O'Donnell.

Foundations of Advanced Information Visualisation (FADIVA) Working Group (Esprit Grant 8422), 1994-1996, with Dr R Cooper. Collaborating with GMD (Germany) & University of Rome.

User Interface Design for Mechanized Theorem Proving (EPSRC Grant GR/K25038), 1994-1997; with Dr T Melham, Dr M Thomas.

Systematic Generic Support for User Interfaces to Databases (EPSRC Grant GR/L 02692), 1996 - 1999, with Dr R Cooper. Collaborating with University of Manchester and Napier University.

Local Expertise

There is much local expertise at Glasgow that can be used to provide feedback and support to this proposal. We are both part of the internationally-renowned GIST Group. This multidisciplinary group contains skills in all areas of human-computer interface design from computing science to psychology. Members will be able to provide feedback on all areas of the research. Dr Steve Draper, Department of Psychology, has much experience in the design and evaluation of interfaces. His skills will be beneficial when designing experiments to test the sonically-enhanced interfaces. Professor Derek Mcaulay, Department of Computing Science, has much expertise in sound technology. His existing EPSRC MINIM and SHEFC NetMuse projects aim to investigate the provision of high quality interactive distributed music services. Professor Mcaulay will be able to provide valuable advice on the problems of presenting sounds via computers. Professor Malcom Atkinson is currently running the PJava project in association with Sun Microsystems. He and his researchers will be able to give us valuable help on the Java language.

Beneficiaries

Our work has improved the usability of computer interfaces for a wide range of users. For example, Brewster's research has shown significant benefits for sighted users using graphical user interfaces because it allows them to work faster and recover from errors more quickly. It also has potential benefits for partially-sighted users who will be able to use the enhanced graphical interfaces more easily.

The other strands of our work have improved the use of single-switch scanning systems for users with severe physical disabilities. We have also created a new method by which blind users can access hierarchies of information and this also has benefits for users who interact over the telephone or where there is a very restricted channel.

REFERENCES

1. Alty, J.L. Can we use music in human-computer interaction? In *Proceedings of HCI'95* (Huddersfield, UK) Cambridge University Press, 1995, pp. 409-423.
2. Alty, J.L. and McCartney, C.D.C. Design of a multi-media presentation system for a process control environment. In *Eurographics multimedia workshop, Session 8: Systems* (Stockholm), 1991.
3. Arons, B. and Mynatt, E. The future of speech and audio in the interface. *SIGCHI Bulletin* 26, 4 (1994), 44-48.
4. Barfield, W., Rosenberg, C. and Levasseur, G. The use of icons, earcons and commands in the design of an online hierarchical menu. *IEEE Transactions on Professional Communication* 34, 2 (1991), 101-108.
5. Blattner, M., Papp, A. and Glinert, E. Sonic enhancements of two-dimensional graphic displays. In *Proceedings of ICAD'92* (Santa Fe Institute, Santa Fe) Addison-Wesley, 1992, pp. 447-470.
6. Blattner, M., Sumikawa, D. and Greenberg, R. Earcons and icons: Their structure and common design principles. *Human Computer Interaction* 4, 1 (1989), 11-44.
7. Brewster, S.A. *Providing a structured method for integrating non-speech audio into human-computer interfaces*. PhD Thesis, University of York, UK, 1994.
8. Brewster, S.A. The development of a sonically-enhanced widget set. In *Proceedings of EWHCI'95* (Moscow, Russia) International Centre for Scientific and Technical Information, 1995, pp. 126-129.
9. Brewster, S.A. A sonically-enhanced interface toolkit. In *Proceedings of ICAD'96* (Xerox PARC) Xerox and Santa Fe Institute, 1996, pp. 47-50.
10. Brewster, S.A. Using Non-Speech Sound to Overcome Information Overload. *Accepted for publication in Displays*, Special Issue on Multimedia Displays (1997).
11. Brewster, S.A. and Crease, M.G. Making Menus Musical. In *To appear in Proceedings of IFIP Interact'97* (Sydney, Australia) Chapman & Hall, 1997.
12. 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.
13. Brewster, S.A., Raty, V.-P. and Kortekangas, A. Enhancing scanning input with non-speech sounds. In *Proceedings of ACM Assets'96* (Vancouver, Canada) ACM Press, 1996, pp. 10-14.
14. Brewster, S.A., Wright, P.C., Dix, A.J. and Edwards, A.D.N. The sonic enhancement of graphical buttons. In *Proceedings of IFIP Interact'95* (Lillehammer, Norway) Chapman & Hall, 1995, pp. 43-48.
15. Brewster, S.A., Wright, P.C. and Edwards, A.D.N. An evaluation of earcons for use in auditory human-computer interfaces. In *Proceedings of ACM/IFIP INTERCHI'93* (Amsterdam) ACM Press, Addison-Wesley, 1993, pp. 222-227.
16. Brewster, S.A., Wright, P.C. and Edwards, A.D.N. The design and evaluation of an auditory-enhanced scrollbar. In *Proceedings of ACM CHI'94* (Boston, MA) ACM Press, Addison-Wesley, 1994, pp. 173-179.
17. Brewster, S.A., Wright, P.C. and Edwards, A.D.N. Experimentally derived guidelines for the creation of earcons. In *Adjunct Proceedings of BCS HCI'95* (Huddersfield, UK), 1995, pp. 155-159.
18. Brewster, S.A., Wright, P.C. and Edwards, A.D.N. Parallel earcons: Reducing the length of audio messages. *International Journal of Human-Computer Studies* 43, 2 (1995), 153-175.
19. Brown, M.L., Newsome, S.L. and Glinert, E.P. An experiment into the use of auditory cues to reduce visual workload. In *Proceedings of ACM CHI'89* (Austin, Texas) ACM Press, Addison-Wesley, 1989, pp. 339-346.
20. Cohen, J. Monitoring background activities. In *Proceedings of ICAD'92* (Santa Fe Institute, Santa Fe) Addison-Wesley, 1992, pp. 499-532.
21. Edwards, A.D.N. Soundtrack: An auditory interface for blind users. *Human Computer Interaction* 4, 1 (1989), 45-66.
22. Edwards, A.D.N., Ed. *Extra-Ordinary Human-Computer Interaction*. Cambridge University Press, Cambridge, UK, 1995.
23. Edworthy, J., Loxley, S. and Dennis, I. Improving auditory warning design: Relationships between warning sound parameters and perceived urgency. *Human Factors* 33, 2 (1991), 205-231.
24. Gaver, W. and Smith, R. Auditory icons in large-scale collaborative environments. In Proceedings of *IFIP Interact'90* (Cambridge, UK) Elsevier Science Publishers B.V. (North Holland), 1990, pp. 735-740.
25. Gaver, W., Smith, R. and O'Shea, T. Effective sounds in complex systems: The ARKola simulation. In *Proceedings of ACM CHI'91* (New Orleans) ACM Press, Addison-Wesley, 1991, pp. 85-90.
26. Myers, B.A. State of the art in user interface software tools. In *Advances in Human-Computer Interaction*, Hartson, H.R. and Hix, D. (Eds.), Ablex Publishing, 1991
27. Mynatt, E.D. Auditory presentation of graphical user interfaces. In *Proceedings of ICAD'92* (Santa Fe Institute, Santa Fe) Addison-Wesley, 1992, pp. 533-555.
28. Patterson, R.D. *Guidelines for auditory warning systems on civil aircraft*. Civil Aviation Authority, London, 1982, CAA Paper, 82017.
29. Portigal, S. *Auralization of document structure*. MSc. Thesis, The University of Guelph, Canada, 1994.
30. Savidis, A. and Stephanidis, C. Developing dual user interfaces for integrating blind and sighted users: The HOMER UIMS. In *Proceedings of ACM CHI'95* (Denver, CO) ACM Press, Addison-Wesley, 1995, pp. 106-113.
31. Stevens, R.D., Brewster, S.A., Wright, P.C. and Edwards, A.D.N. Providing an audio glance at algebra for blind readers. In *Proceedings of ICAD'94* (Santa Fe Institute, Santa Fe) Santa Fe Institute, 1994, pp. 21-30.

PART 3: DIAGRAMMATIC PROJECT PLAN

	Stephen Brewster / Philip Gray	Research Assistant	Research Student
Year 1	Help RA with the overall design of the toolkit Supervise RS and help design the widget experiments	Design and develop the basic structure of the toolkit With RS, design and implement widget experiments	Design, implementation and conduct of widget experiments Develop guidelines First year report
Year 1	Write paper describing the detailed structure of the toolkit. Write paper describing results of widget experiments and preliminary guidelines.		
Year 2	Support work of the RS and RA Work on the detailed design of the toolkit, ironing-out any problems. Help in design of new widgets	Implement and test the remaining standard widgets Take the individual widgets and build them into the toolkit	Using experimental skills from first year, design and implement new widgets that fully integrate graphics and sound
Year 2	Write paper describing the detailed guidelines and implementation of the toolkit.		
Year 3	Design sonically-enhanced applications that will use the toolkit with RA. Support RS writing up thesis	Evaluate toolkit and guidelines with interface designers. Evaluate sonically-enhanced applications.	Complete the design and evaluation of new widgets Write up thesis
Year 3	Write papers on evaluation of the toolkit and new widgets.		
Year 3	Work on final report and plan career development for RA.		
	Submit Final Report		Submit Thesis