Design Patterns for Auditory Displays

C Frauenberger[†‡§], T Stockman[†], V Putz[§] & R Höldrich[§]

[†] *Department of Computer Science, Queen Mary University of London, Mile End Road, London E1 4NS, UK*

Email: *frauenberger@dcs.qmul.ac.uk, tonys@dcs.qmul.ac.uk*

URL: *http://www.dcs.qmul.ac.uk*

[‡] *Signal Processing and Speech Communication Laboratory, Graz University of Technology, Inffeldgasse 12, A-8010 Graz, Austria*

URL: *http://spsc.tugraz.at*

[§] *Institute of Electronic Music and Acoustics, University of Music and Dramatic Arts Graz, Inffeldgasse 10/3, A-8010 Graz, Austria*

Email: *v.putz@gmx.at, hoeldrich@iem.at*

URL: *http://iem.at*

This paper proposes the use of patterns in the design process for auditory displays and /or interfaces realized in other modalities. We introduce a meta-domain in which user interfaces can be designed using these patterns without determining their means of realization. The mode-independent description of such interfaces can then be used to create the real interface maintaining the strengths of the different interaction channels. While this work is focused on how this approach can be applied on auditory displays, we keep in mind that the approach shall be applicable on other interaction modalities equally. The development of a set of mode independent interaction patterns is shown along with descriptions of their representations in the auditory domain. A real world application was chosen to evaluate the approach; Microsoft Explorer was analysed, described through the mode independent interaction patterns

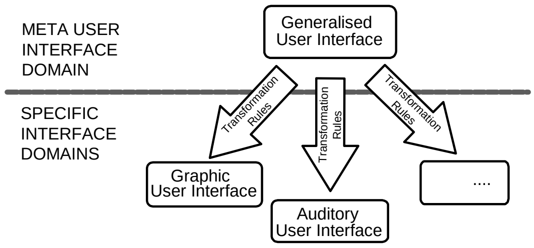

Figure 1: The meta domain as common ground for different representations of a user interface.

and transformed into the auditory domain making extensive use of 3D audio rendering techniques. The result, a file manager created in a virtual audio environment, was evaluated with sighted persons as well as visually impaired and blind participants showing the feasibility and usability of the approach.

Keywords: auditory displays, design patterns, virtual audio environments, assitive technology.

1 Introduction

The auditory mode in human-computer interaction is increasingly recognized as a strong and reliable channel of communication. However, the field is still very experimental, lacking widely accepted design principles and little is known about what makes good auditory displays. This paper proposes an approach towards establishing such design principles and guidelines for the auditory domain by adapting well known design methods from GUIs to other modalities. Doing so, however, we keep in mind that human senses are significantly different and interaction and information processing follows different rules — *"Many people get their auditory representations wrong because they think visually"* (anonymous paper reviewer). Therefore, we approach this problem field by introducing a common ground for user interfaces realized in different modalities: a mode-independent description of user interfaces that would determine the functionality without their means of realization. In such a meta-domain designers can use well known concepts like design patterns while their realization can maintain the particular strengths of each interaction channel. Figure 1 illustrates the approach.

We aim at developing design patterns that can be used to create user interfaces in the meta-domain and at providing tools and guidelines to create the real interfaces in the auditory domain. Although our work is focused on the audio mode, we develop concepts keeping in mind that they shall be usable with every other interaction channel equally. The design of multi-modal interface may also benefit from this approach, because it allows to split a user interface between modalities according to user interaction tasks.

The following sections describe in detail the motivation behind the approach and the state-of-the-art in auditory design and HCI design. Subsequently, section 2 describes the approach of using mode independent interaction patterns for the design of user interfaces, showing the way the patterns were developed and describing an example with an auditory realization of a pattern. Section 3 describes the evaluations conducted using this design method including the details of the test design, the evaluation and the analysis. Finally, we conclude this paper and state ideas for future work.

1.1 Motivation

Audio has been part of user interfaces since the very early days of computers and sounds were first deployed to indicate errors or warnings while processing a request. It was straight forward that sound is an appropriate way to alert users and the simple beeps used in such interfaces were about the same technical level as the graphical output at that time. Since then, the development of graphical user interfaces has been substantial, greatly outstripping the development effort and understanding of auditory interfaces. Currently audio is radically under used in human-computer interfaces, being still predominantly used to signal the occurrence of an error, or to draw the user's attention to the occurrence of a specific event. Sound remains largely unused as a fully integrated element of multimedia interfaces, and is little used in support of mainstream task performance.

The most important reasons to change this in the context of modern user interface design can be summarized by [Kramer 1994; Frauenberger et al. 2004a]:

- Increased complexity of tasks.
- Miniaturization of devices.
- Mobility of the user.
- Naturalness of interfaces.
- Accessibility for disabled users.

Many requirements for an improved user interface resulting from the points above can be achieved by auditory displays, whether sound is being used as an element of a multimedia interface or on its own as a single medium. To prevent user interfaces from being overloaded due to the complexity of tasks, multi-modal approaches can be utilized to exploit human sensory capabilities to their optimum. Auditory displays can be designed as virtual or augmented reality and do not require much physical space. Sound is also highly portable and integrable in mobile devices. Since audio is an important part of our every-day environment, the integration of audio clearly improves the naturalness in human-computer interaction. Finally, visually impaired users rely on non-visual modes to interact with the increasing number of information technologies integrated in our lives.

Past research has shown that audio is capable of contributing much more than beeps to user interfaces. There are, however, significant problems in understanding human methods of acoustic communication and how complex tasks in human-computer interaction can be rendered in the auditory domain. Many prototypes have been developed and evaluated and very satisfying solutions were created for specific problem statements. But as audio displays become more widespread, user interface designers must be provided with robust, but customizable patterns and design principles.

1.2 State of the Art

The increasing computational power available for digital signal processing made increasingly complex simulations of acoustical environments possible. The term *Virtual Audio Environment* describes the simulation of acoustical scenes with sound reproduction techniques. The goal is to create natural environments which are customizable and controllable in real time, very much as visual virtual environments were developed [Lokki 2002]. Acoustical rendering of objects (sound sources), the environment and the listener can be realized using a number of different techniques like Ambisonics [Bamford 1995], Wave Field Synthesis [Verheijen 1998] or Vector Based Amplitude Panning [Pulkki 1997]. These techniques were just recently utilized to realize auditory displays[Strauss et al. 2003; Frauenberger et al. 2004b].

While these techniques may provide accurate acoustical rendering there are still many questions about the human capability of auditory perception. A problem still subject of investigation is sound source segregation in virtual environments when rendering sound sources concurrently [McGookin & Brewster 2003], another one is providing robust orientation cues for navigation [Gröhn et al. 2003]. Further challenges result from complex psychoacoustic effects like informational masking which are not fully understood in the context of auditory displays [Oh & Lufti 1999].

A variety of auditory displays were developed for specific problem domains (e.g. [Kobayashi & Schmandt 1997; Schmandt 1998; Walker et al. 2001]) and some efforts were taken towards a structured approach to more generic solutions. Early proposals include the *Mercator* project, the first framework targeting customary Unix desktops [Edwards et al. 1993; Mynatt 1995]. Another proposal was *Y-Windows* also following the idea of building alternative, audio rendering engines (servers) for existing clients requesting their user interface representation [Kaltenbrunner 2002]. However, both approaches implied that graphical concepts were translated into the auditory domain and had therefore their limitations. A first attempt to break up with this and introducing a mode independent meta-domain was made in [Frauenberger et al. 2004a] and subsequently led to the proposal stated by this paper.

Attempts to use common HCI engineering methods in the design of auditory displays include the investigation of audio metaphors [Mynatt & Edwards 1995] and other structural approaches to include sound into human-computer interaction (Earcons [Blattner et al. 1989; Brewster 1994]). Recently, the proposal of using patterns in sonification highlights the advantages of such methods in re-usable designs [Barrass 2003; Adcock & Barrass 2004].

However, auditory representations of user interfaces are in need of profound heuristics to assess user satisfaction similar to those in the graphical domain [Nielsen 1993]. Examples of where work is needed to identify heuristics to guide the process of auditory display design include minimizing the problems incurred due to the transient nature of sound, quantifying how the effectiveness of interactions can be improved through learning and providing guidelines for how sound can best be integrated with other media [Kramer et al. 2005]. The design pattern approach is one of the most widespread and popular techniques for user interface design in the graphical domain[van Welie & Trætteberg 2000; Trætteberg 2000] and we believe that it is also suitable for the task given. The subsequent section will elaborate on the use of this technique in the context of mode-independent design.

2 Interaction Patterns

The concept of interaction patterns is well known in user interface design, but is very much focused on graphical interfaces and there are very different approaches to it available. Despite the fact that patterns are often created from the perspective of the programmer or designer, they also tend to be very specific to a certain problem domain. A more generalized approach that aimed at user-centred pattern design led to a set of patterns addressing the most common user requirements[van Welie & Trætteberg 2000]. These patterns are task-related and address different types of user requirements (e.g. the need to be able to locate specific commands and to find out how to activate them), based upon the most fundamental usability principles [Norman 1988]:

Availability: The required parts of the application need to be available at the right time and should imply correct usage. Availability[1] concerns the mapping between intended actions of the user and the operations actually required.

Affordances provide strong clues to the operation of things (e.g. knobs are for turning, buttons for pushing). When affordances are effectively used within an interface, the user knows what to do with no further instruction needed.

Constraints minimize the number of possible actions and give additional information about the correct usage of the interface elements.

Natural mapping: If the relationship between the controlling elements of an application and their results are natural for the user, it simplifies the learning of the application and assists recall. Natural mapping depends on physical analogies and cultural standards and is therefore subjective to different user groups.

Conceptual models: By interacting with an application, the user builds up a conceptual model of it. If this model is equivalent to the task model of the application, it allows the user to predict the effects of his actions.

Feedback: Information about the result of his actions is sent back to the user and enables immediate control of the input.

In order to re-formulate the chosen set of patterns from [van Welie & Trætteberg 2000] we slightly changed the terminology used: The *representation medium* means

[1]The term availability was chosen to replace visibility to avoid visual associations.

the domain or the combination of the domains in which the user interface will be realized. Within this representation medium there are *representation areas* defined which provide the boundary for *objects* of the user interface. These objects may result from one or more *interaction patterns* transformed into the representation medium.

2.1 Structure

While developing the patterns, we recognized that certain tasks or parts of patterns recurred in other patterns too. This led to the concept of atoms and contextual attributes; similar to a vocabulary for instantiating designs. A set of atoms was developed from which patterns may draw when addressing a particular set of user requirements. This also implies consistent representation of similar elementary units throughout the whole interface although atoms are not sufficient to solve any interaction problem. The name indicates that this level should be the smallest piece in any pattern to avoid over exaggerating modularity.

In order Not to end with a totally unrelated patchwork of small pieces of a user interface, each atom provides contextual attributes. These attributes need to be set by the parent pattern in order to indicate their context. In the graphical domain this would, for example, mean that certain elements like buttons or text fields are *in the same window* sharing the same frame and background colour. The following contextual attributes were identified for our set of atoms:

Similarity: Atoms in the same pattern share properties like timbre, rhythm or type of voice in their acoustical representation.

Proximity: Atoms in the same pattern are grouped together based on the available dimensions of the representation area (space or pitch ranges).

Homogeneity: The same types of atoms should be placed adjacently in a pattern on the basis of the available dimensions of the representation area (space or pitch ranges).

It is important to state that not only the patterns and the atoms undergo the transformation process in order to form a real user interface, but also the contextual attributes must be mapped into the different representation media. Their realization in the auditory domain will differ considerably from the visual domain.

2.2 The Patterns

Each pattern consists of a description of the related *user-problem*, a listing of further *conditions* which must be fulfilled to solve the problem, a possible *solution* for the problem and a listing of the *attributes* which must be mapped. Figure 2 shows how those descriptions were made (example taken from [Putz 2004]).

Similar descriptions were made for all patterns from [van Welie & Trætteberg 2000] (Command area, Wizard, Contextual menus, Mode cursor, Setting attributes, Link in the real world, List browser, Continuous filter, Preview, Navigating categories, Container navigation, Unambiguous format, Focus / Selection, Grouping layout, Progress, Hinting, Message / Warning, Shield, Preferences, Favourites).

Command Area	
Problem	**The user needs to know where to find the possible commands and how to activate them.**
Conditions	• Immediate access to all available functions increases interaction speed but consumes a large amount of the available representation area. • The main working area should be kept as large as possible. • Concurrent representation of many objects increases the cognitive load. • Some functions are used more often then other functions. • Some functions need additional parameters to be set by the user before they can be executed.
Solutions	**Put the shortcuts to the possible commands in a specific recognisable area** A part of the representation area is reserved for shortcuts to the functions of the application. This area should be distinguishable from other working areas because of contextual attributes. If the number of shortcuts is large, they should be conceptually grouped. The command area can be subdivided to give access to a group of commands. The command area and the included shortcuts should be accessible as direct as possible, especially frequently used shortcuts. Feedback information about the availability of the shortcuts can be given. The number of represented functions should not be too large to limit the necessary representational area and the cognitive load.
Attributes	• Atoms for the representation of the shortcuts (Links, Triggering Elements, Selections). • Separation of groups, when conceptual grouping is required
Graphical Example: Command Area of MATLAB 6.5	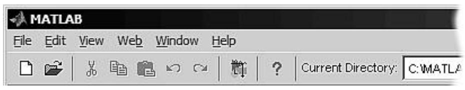

Figure 2: An example of a pattern description.

Besides the example of a visual realization given in the description, auditory representations were developed for the patterns needed in the prototype.

To illustrate such an auditory representation, the *Command Area* pattern is used. It is addressed to solve the following interaction problem: The user needs to know where to find the possible commands and how to activate them. When solving this problem, several further conditions have to be taken into account:

1. Immediate access to all available functions of the application increases interaction speed but consumes a large amount of the available representation area.

2. The amount of the remaining representational media space for the main working area should be kept as large as possible.

3. Concurrent representation of many objects increases the cognitive load.

4. Some functions are used more often then other functions.

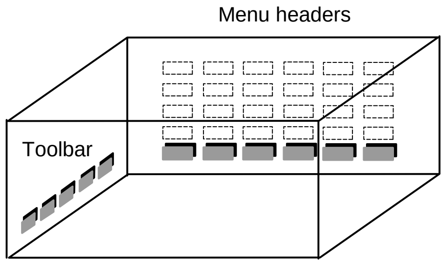

Figure 3: The layout of the auditory representation of the Command Array pattern.

5. Some functions need additional parameters to be set by the user before they can be executed.

The solution describes the creation of a recognizable area in the representation medium which contains shortcuts to important functions of the applications and a menu. The shortcuts exploit affordances to indicate the usage and use metaphors (natural mapping) to save representation space. Atoms used by the pattern are the link atom and the triggering element atom. The atoms will be connected to the pattern by similarity attributes and triggering elements will be grouped together, if necessary, in meaningful sub-groups (e.g. for a menu).

Transforming this interaction pattern to the auditory domain was performed for our evaluation test described in Section 3. The two main parts — the menu headers and the toolbar — are situated alongside two different walls of the virtual room as shown in Figure 3.

The triggering elements of the toolbar are presented with short rhythmic patterns (Earcons). The toolbar elements are connected by proximity and homogeneity with regard to space and similarity with regard to the sound design. The menu headers are presented as combination of speech and instrumental sound presented concurrently. The name of each menu item is spoken accompanied by an unobtrusive sound with rising pitch, starting with lowest pitch at the first header item and increasing in pitch while moving towards the last one. The menu headers are connected by proximity and homogeneity with regard to space and pitch and by similarity with regard to the compound sound design (speech and instrumental tone). The sub menu items belonging to the menu headers are placed above the headers within the virtual room also providing compound sounds.

3 Evaluation

For evaluation purposes a real world application was chosen and analysed. Its user interface was described through the mode independent interaction patterns and then transformed to the auditory domain. As one of the key applications on a computer desktop a file manager was chosen — the Microsoft Windows Explorer.

File	Edit	View	Favourites	?
New Folder	Undo	Statusbar	www.iem.at	Information
Delete	Cut	List	www.orf.at	Support
Rename	Copy	Details	www.google.com	
Properties	Paste	Sort		
Close	Select All	Change To		
		Reload		

Figure 4: Implemented functions of the auditory Explorer version.

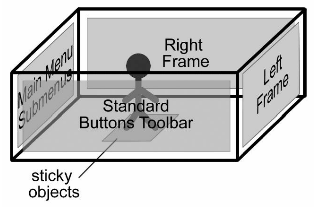

Figure 5: The virtual audio environment for the explorer.

3.1 *Prototype Design*

For testing, the complexity of the application was reduced, rarely used functions were not implemented. The remaining menu structure is shown in Figure 4. It still covers the basic functionality of a file manager.

From the standard buttons toolbar, five shortcuts were implemented: Undo, Delete, Cut, Copy, Paste. Through the context menu the functions Cut, Copy, Delete, Rename and Properties were available.

The functions described of the Explorer application was analysed and interaction patterns were identified. The *Container Navigation* pattern was used to describe the two main frames of the Explorer. For the folder tree in the left frame the *tree structure* atom was used and the *list* atom described the right content frame. The *Command Area* pattern described the menu structure and the tool bar area. Finally, the *Contextual Menu* pattern solved the availability of the context menu and the *Message* pattern was used for all pop-up windows at their occurrence. Figure 5 shows the basic layout of the virtual audio environment into which the patterns were transformed.

The container navigation pattern was transformed into the auditory domain as two different areas in the virtual environment, the walls to the front and to the right. The list at the front wall was stretching out to the top using speech for the name of the items with different voices to indicate the type (folder or file). The tree structure was laid out on a grid on the wall to the right with the left-bottom corner being the root and the right-top corner the last hierarchical level. Unfolding and folding a node in the tree was also implemented. In both areas the user was able to select items and get a context menu. The contextual pattern was realized as sticky objects following the user wherever she moves. The content of the contextual menu pattern was again solved by triggering elements. The same concept was used to realize pop-up windows — sticky objects remaining to the front of the user, but with different background sound. The menu and toolbar was realized as shown in the example described above.

All sounds were audible when the user moved into their range. This means that there was silence in the starting position, but moving towards a wall meant that one could hear the 5 menu items at once at different levels depending on the distance.

Interaction with the prototype was done by joystick and keyboard. To avoid confusion while navigating through the virtual environment no relative movements are supported. Bringing the joystick to the starting position means moving to the centre of the room facing the front wall. Moving up, along the z-axis, in the environment was implemented using the throttle handle of the joystick. The localization of different sound sources was improved by using a head tracker.

3.2 Implementation

The prototype was implemented using Pure Data (PD) by Miller Puckette and the binaural Ambsonics library extension for PD developed at the Institute of Electronic Music and Acoustics Graz [Noisternig et al. 2003b,a]. This extension allows the simulation of virtual audio environments with binaural Ambisonics of 3rd order and room acoustics including reflections and late reverberation. Multiple sound sources may be placed in such an environment and will be rendered efficiently for headphones in real time allowing interaction with joysticks, head-trackers or other feedback devices.

The content of the hard disk was faked in order to keep complexity low and not to introduce additional problems for the prototype interacting with the operating system.

3.3 Test Design

The test was performed by a group of 15 test participants divided into two groups. Group S were seven students of the Graz University of Technology and one person already holding a masters degree. All of them were between 20 and 27 years old and had good experience with computers and Windows. Group B consisted of four persons who are totally blind and three persons with visual disabilities. The use of visual screens was only feasible for them using additional magnification software. Six participants in Group B hold the ECDL (European Computer Driving License) and use a computer in their work being very experienced with Windows software. One member had little experience with computers, but was attending the course for receiving the ECDL. In average Group B was a little older.

After instruction, the participants got 15minutes of training time with the application, participants were given a list of 7 tasks to perform. The tasks involved finding out how many files are in a specific folder, finding the size of files, copying, moving and creating files or folders.

Throughout the test, different types of data were collected. On the one hand, the hierarchic structure of files and folders after the test is stored in a text file. Apart from that, two further lists report the whole test sequence, one list containing the movement of the joystick within the virtual room (x,y,z-coordinates, rotation around

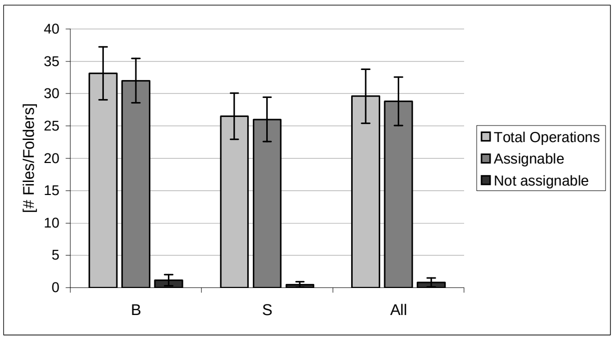

Figure 6: Number of file/folder handling operations

the z-axis) in a resolution of 50ms, the other list reporting any action performed by the participants with a time index, so that both lists can be combined. With these two lists, the whole test performance of the participants can be reproduced and visualized. The list of reported events can also be used to compute the quantity of different events. Apart from that, the whole performance of the participants was attended by the test administrator via headphones who additionally took notes.

After the test, the participants had to answer two questionnaires. One concerning the individual background of the participants, the other trying to catch the subjective impression of the participants after the test.

3.4 Analysis

The three-dimensional layout of the virtual room with the different meanings of the four surrounding walls and the two-staged movement (ground-plane movement towards the walls, vertical movement for selection) proved to be sufficient to host the elements of a real-world application. On the ground floor, at least 20 items (5 on each wall) can be placed, not to mention the potential with regard to vertical placement. The thematic grouping of elements on the different walls was easy to memorize for the test users. The usability of the mappings of particular interaction patterns is different. While the menu structure was easy to use for most test participants, the representation of the folder hierarchy is in need for improvement: Hardly any test user had a clear overview of the file and folder structure. According to the participants the static grid layout was confusing because they lost track of the absolute position while navigating through the tree structure. This was fostered by the fact that hardly anyone could re-construct the file structure correctly after the test.

Remarkable is the percentage of user operations on items which are not related to the task, which were performed by mistake. Figure 6 shows that Group B (the blind user group) required slightly more file handling operations to fulfil the tasks than Group S (normal sighted user group). Hardly any file or folder was selected by mistake, the percentage of selected files that can be assigned to one of the tasks lies between 97% and 98% for all groups. The same was observed with menu selections.

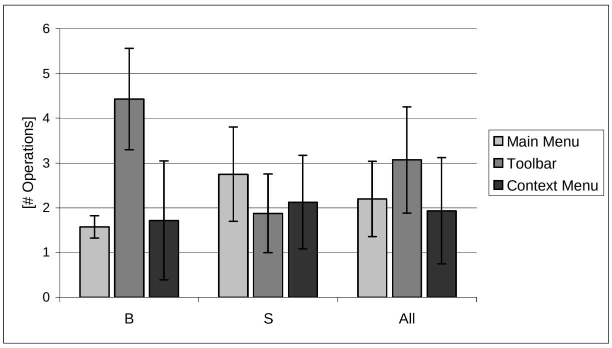

Figure 7: Usage of different menus

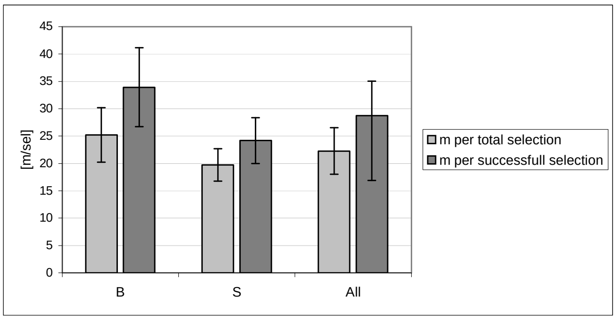

Figure 8: Covered distance per selection

Group B needed slightly more menu operations, but no single user chose a menu item not related to the task he was performing.

Throughout the whole test, the participants were free to choose how to fulfil the tasks: With the aid of the main menu, the standard buttons toolbar or the context menu. Figure 7 shows, which option was preferred by the two groups. Group B had a clear preference for the toolbar, whilst Group S had a slight preference for the main menu. The usage of the context menu is similar for the group averages. The standard deviation is high for all values.

The virtual distance covered by the participants for one successful selection of an item shows that it lies within the dimension of the virtual room (10m x 14m and 12m high = 26m diagonal). Again it is slightly more for Group B, but also lies within the standard deviation (Figure 8). These measures are comparable with mouse movements in graphical interfaces although no comparable data was collected for the graphical Explorer application.

In general, the subjective questionnaire showed that the users liked the system and felt comfortable to solve the problems given. The sound design was chosen to be as non-obtrusive as possible and the silent centre position was appreciated.

The main goal of the application was to have equal usability for blind and sighted users. With regard to the main performance measurements (total test time, necessary file/folder handling operations, necessary menu handling operations, working speed and the percentage of correct results), the two groups of users have reached similar results, although there is a high standard deviation for some values.

4 Conclusions

This paper proposed a set of mode independent interaction patterns for designing user interfaces in different interaction domains. The approach was chosen to overcome the difficulties of translating existing user interfaces into other domains like GUIs into auditory displays. The idea behind the patterns and their structure were explained followed by an evaluation of a real world application. The existing graphical user interface of this application was analysed, described through the proposed set of interaction patterns and transformed into the auditory domain.

The evaluation showed that the approach is feasible and promising for establishing design principles for auditory displays. Both groups of test participants, blind and normal sighted persons, were able to use the application with equal efficiency. All participants reported that they felt comfortable and could imagine to work with a similar system in their real working environments.

Problems remained with some of the representations where either sound design or the idea of representation was weak. However, having the user interface built from patterns and atoms, it is possible to isolate the problems and improve the transformation of the single item instead of re-designing the whole interface.

Future work needs to look closer into the mechanisms of acoustic communication considering psychoacoustics, learning effects, mental load and cross modal interactions. The patterns need to be revised and improved in their acoustical representation and more flexible frameworks must be developed in order to ease the integration of auditory displays into customary computing systems.

5 Acknowledgements

We would like to thank all who volunteered to participate in the evaluation. Especially the people from the ISIS group (Integration, Service, Information and Schooling) in Graz / Austria for the close collaboration and for giving us an idea about the needs the visually impaired have when it comes to computers. ISIS is a project run by BFI (Berufsförderungsinstitut) Styria.

Support by the SonEnvir (http://sonenvir.at) project sponsored by the Zukunftsfonds Steiermark is gratefully acknowledged.

References

Adcock, M. & Barrass, S. [2004], Cultivating Design Patterns for Auditory Displays, *in* S. Barrass & P. Vickers (eds.), *Proceedings of the International Conference on Auditory Display (ICAD'04)*, International Community for Auditory Displays. Proceedings will be online via http://www.icad.org/.

Bamford, J. S. [1995], An Analysis of Ambisonic Sound Systems of First and Second Order, Master's thesis, University of Waterloo. http://audiolab.uwaterloo.ca/~jeffb/thesis/thesis.html.

Barrass, S. [2003], Sonification Design Patterns, *in* E. Brazil & B. Shinn-Cunningham (eds.), *Proceedings of the International Conference on Auditory Display (ICAD'03)*, Boston University Publications for International Community for Auditory Displays, pp.170–5. http://www.icad.org/websiteV2.0/Conferences/ICAD2003/paper/42 Barrass.pdf.

Blattner, M. M., Sumikawa, D. A. & Greenberg, R. M. [1989], Earcons and Icons: Their Structure and Common Design Principles, *Human–Computer Interaction* **4**(1), 11–44.

Brewster, S. A. [1994], Providing a Structured Method for Integrating Non-speech Audio into Human–Computer Interfaces, PhD thesis, University of York, England, UK.

Edwards, W. K., Mynatt, E. D. & Rodriguez, T. [1993], The Mercator Project: A Nonvisual Interface to the X Window System, *The X Resource* **7**(summer).

Frauenberger, C., Höldrich, R. & de Campo, A. [2004a], A Generic, Semantically Based Design Approach for Spatial Auditory Computer Displays, *in* S. Barrass & P. Vickers (eds.), *Proceedings of the International Conference on Auditory Display (ICAD'04)*, International Community for Auditory Displays. Proceedings will be online via http://www.icad.org/.

Frauenberger, C., Putz, V. & Höldrich, R. [2004b], Spatial Auditory Displays —A Study on the Use of Virtual Audio Environments as Interfaces for Users with Visual Disabilities, *in* G. Evangelista (ed.), *Proceedings of the Seventh International Conference on Digital Audio Effects (DAFx04)*, *****PUBLISHER*****, pp.384–9.

Gröhn, M., Lokki, T. & Takala, T. [2003], Comparison of Auditory, Visual and Audio-Visual Navigation in a 3D Space, *in* E. Brazil & B. Shinn-Cunningham (eds.), *Proceedings of the International Conference on Auditory Display (ICAD'03)*, Boston University Publications for International Community for Auditory Displays, pp.200–3. http://www.icad.org/websiteV2.0/Conferences/ICAD2003/paper/49 Grohn.pdf.

Kaltenbrunner, M. [2002], Y-Windows: Proposal for a Standard AUI Environment, *in* R. Nakatsy & H. Kawahara (eds.), *Proceedings of the International Conference on Auditory Display (ICAD'02)*, Advanced Telecommunications Research Institute (ATR) for International Community for Auditory Displays, pp.154–8. http://www.icad.org/websiteV2.0/Conferences/ICAD2002/proceedings/74_Kaltenbrunner.pdf.

Kobayashi, M. & Schmandt, C. [1997], Dynamic Soundscape: Mapping Time to Space for Audio Browsing, *in* S. Pemberton (ed.), *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'97)*, ACM Press, pp.194–201.

Kramer, G. (ed.) [1994], *Auditory Display: Sonification, Audification, and Auditory Interfaces*, Vol. 18 of *Santa Fe Institute Studies in the Sciences of Complexity*, Addison–Wesley.

Kramer, G., Walker, B., Bonebright, T., Cook, P., Flowers, J., Miner, N. & Neuhoff, J. [2005], Sonification Report: Status of the Field and Research Agenda, http://www.icad.org/websiteV2.0/References/nsf.html.

Lokki, T. [2002], Creating Interactive Virtual Auditory Environments, *IEEE Computer Graphics and Applications, special issue "Virtual Worlds, Real Sounds* **22**(4), 49–57. Electronic publication http://www.computer.org/cga/.

McGookin, D. K. & Brewster, S. [2003], An Investigation into the Identification of Concurrently Presented Earcons, *in* E. Brazil & B. Shinn-Cunningham (eds.), *Proceedings of the International Conference on Auditory Display (ICAD'03)*, Boston University Publications for International Community for Auditory Displays, pp.42–6. http://www.icad.org/websiteV2.0/Conferences/ICAD2003/paper/10 McGookin.pdf.

Mynatt, E. [1995], Transforming Graphical Interfaces into Auditory Interfaces, PhD thesis, Georgia Institute of Technology.

Mynatt, E. D. & Edwards, W. K. [1995], Metaphors for Nonvisual Computing, *in* W. K. Edwards (ed.), *Extraordinary Human–Computer Interaction*, Cambridge University Press, pp.201–20.

Nielsen, J. [1993], *Usability Engineering*, Academic Press.

Noisternig, M., Höldrich, R., Musil, T. & Sontacchi, A. [2003a], A 3D Ambisonic based Binaural Sound Reproduction System, *in* *****EDITOR***** (ed.), *Proceedings of the International Conference on Multichannel Audio (AES 24)*, Audio Engineering Society, p.*****PAGES*****. *****Author says these are online proceedings but the AES Web site says it is a 350 page book*****.

Noisternig, M., Musil, T., Sontacchi, A. & Höldrich, R. [2003b], 3D Binaural Sound Reproduction using a Virtual Ambisonic Approach, *in Proceedings of the 2003 IEEE International Symposium on Virtual Environments, Human–Computer Interfaces and Measurement Systems (VECIMS'03)*, IEEE Publications, pp.174–8.

Norman, D. A. [1988], *The Psychology of Everyday Things*, Basic Books.

Oh, E. L. & Lufti, R. A. [1999], Informational Masking by Everyday Sounds, *Journal of the Acoustical Society of America* **6**(106), 3521–8.

Pulkki, V. [1997], Virtual Sound Source Positioning Using Vector Base Amplitude Panning, *Journal of the Audio Engineering Society* **45**(6), 456–66.

Putz, V. [2004], Spatial Auditory User Interfaces, Master's thesis, Institute of Electronic Music and Acoustics, University of Music and Dramatic Arts Graz. http://iem.at/projekte/dsp/spatial/dp_putz.

Schmandt, C. [1998], Audio Hallway: A Virtual Acoustic Environment for Browsing, *in* M. E. Atwood, C.-M. Karat, A. Lund, J. Coutaz & J. Karat (eds.), *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'98)*, ACM Press, pp.163–70.

Strauss, M., Höldrich, R. & Sontacchi, A. [2003], A Spatial Audio Interface for Desktop Applications, *in* *****EDITOR***** (ed.), *Proceedings of the AES 24th International Conference: Multichannel Audio, The New Reality*, Audio Engineering Society, p.*****PAGES*****. *****Author says these are online proceedings but the AES Web site says it is a 350 page book*****.

Trætteberg, H. [2000], Model Based Design Patterns, Unpublished position paper given at CHI'00.

van Welie, M. & Trætteberg, H. [2000], Interaction Patterns in User Interfaces, *in Proceedings of the Pattern Languages of Programs Conference (PLoP2000)*, number wucs-00-29 *in* "Washington University Technical Report", Washington University. http://jerry.cs.uiuc.edu/~plop/plop2k/proceedings/Welie/Welie.pdf.

Verheijen, E. [1998], Sound Reproduction by Wave Field Synthesis, PhD thesis, TU Delft.

Walker, A., Brewster, S. & McGookin, S. A. [2001], Diary in the Sky: A Spatial Audio Display for a Mobile Calendar, *in* A. Blandford, J. Vanderdonckt & P. Gray (eds.), *People and Computers XV: Interaction without Frontiers (Joint Proceedings of HCI2001 and IHM2001)*, Springer-Verlag, pp.531–40.

Author Index

Keyword Index