

INTERACTIVE SONIFICATION OF SPREADSHEETS

*Dr Tony Stockman,
Greg Hind, Christopher Frauenberger*

Department of Computer Science,
Queen Mary College
University of London
Mile End Road
London E1 4NS.
{tonys, frauenberger}@dcs.qmul.ac.uk

ABSTRACT

This paper describes the second phase in a project exploring the application of sonification to improve the accessibility of spreadsheets. The principal target population of the study has been visually impaired users, though the approach has potential in a range of other application areas, such as the eyes-free monitoring of large data spaces, or to preserve screen space in mobile use.

The paper begins by examining previous sonification research relevant to the auditory display of spreadsheets, before briefly reviewing the results and requirements arising from the first phase of this project. Phase 2 of the project began with a more detailed requirement investigation, the results of which are outlined. In particular, the requirement for increased interactivity is examined. The way in which these requirements are realised in a further prototype are described, and the results of evaluations of this second prototype discussed.

1. SONIFICATION AND SPREADSHEETS

Current speech-based screen-reader technology provides a fair level of access to spreadsheets through a spoken word interface. This enables users to navigate and examine individual or small groups of values of spreadsheet data. A major weakness of this approach is the difficulty of forming any kind of overview of the data, or of making comparisons across any more than a small number of cells. The reasons such comparisons are difficult appear to be as follows:

- i. The time required to navigate between cells to be compared is so long that users quickly find their memory for the foregoing sequence of cell values deteriorates.
- ii. The above problem is further compounded by the fact that during the process of speech-based screen-reader navigation, the values of cells are announced, whether or not they are relevant to the task. This problem can be offset providing the user has a good knowledge of keyboard-based spreadsheet navigation techniques, along side a knowledge of the coordinates of the target cells (but this is hardly realistic in many situations).

Therefore, whilst speech based screen readers may be good for providing exact values and words, they are not so good for presenting an overview of the data due to the limits of short term memory [10]. There is therefore the need for an additional tool for visually impaired users to use to analyse Excel

spreadsheet data in the way that a sighted user might use chart wizard to create a basic line graph of a set of data.

One approach to providing overviews or comparisons of data across more than very small ranges of cells could be data sonification. Mansur devised a method for line graph sonification called Sound Graphs where the y-axis of the graph is mapped to pitch and the x-axis to time. Movement along the x-axis in time causes notes of different pitches to be played where the frequency of each note is determined by the y value of the graph at that time. Mansur found that after a small amount of training, test subjects were able to identify the overall qualities of the data, such as linearity, monotonicity, and symmetry, on 79 to 95% of the trials [9].

Mansur's experiments only map one type of data against another. What of data with two or more dimensions? Brewster and Browne [10, 11, 12] conducted a number of experiments sonifying graphs containing two data series and showed that sonification allowed users to visualise graphs containing two data series while listening to them and assigning different instruments to each range. Brewster and others [13] conducted experiments exploring 2D tables with speech and non-speech sound and discovered that users found the use of pitch to be valuable in determining the shape of the data within the table.

Further extending the number of different data ranges, Bly's [14] direct mapping approach maps four different variables to pitch, volume, duration and the fundamental wave shape of a note. The resulting notes enabled users to place each sample in one of three related sets. This kind of approach has its limitations though as inevitably the user's attention is drawn more to certain variables than others, making the design of a balanced display all but impossible [5]. It is worth noting from this however, that Pitch is particularly strong as a means of displaying changing variable values. 'There is something fundamental about frequency in the response of the brain to sound, right up to the highest level, the auditory cortex' [4].

With specific reference to spreadsheets, Kramer [5] states that: 'While text-to-speech converters have been developed to serve the blind, one can envision auditory indicators (speech and otherwise) in public transportation, more sophisticated instrumentation, and data-searching tools for blind spreadsheet users'. Many of these ideas have been materialised since Kramer wrote this 11 years ago, but there has yet to be an application created that will allow visually impaired users to sonify spreadsheet data interactively.

There have already been applications created that allow for the sonification of Excel spreadsheet data. Sonification Sandbox [1] is a flexible tool that supports experimentation using a wide range of sound parameters. From the point of interactivity however, it suffers from the fact that data input is

achieved through a comma delimited, numeric only file, which is created by extracting row and column subsets from a spreadsheet. Whilst the application allows considerable flexibility in exploring different parameter mappings, it does not afford interactive use through screen-readers such as Jaws or Window-eyes.

Kramer [5] defines a general auditory display system in the following way:




Figure 1. Based on a diagram of an auditory display system in Kramer 1994, but modified for spreadsheet sonification.

In the case of a system to sonify Excel spreadsheet data, the spreadsheet is the Information Generator, and the Information Receiver is the user listening to the output. The communicative medium is the sonification system itself, which will consist of code to extract the spreadsheet data (Data Receiving), code to map these values to a variable such as frequency (Intermediary Structures), and a means of generating the sound e.g. Musical Instrument Digital Interface (MIDI) or synthesiser software such as CSOUND or Puredata.

2. REQUIREMENTS FOR PHASE 2 OF THE PROJECT

The first phase of the project produced a basic prototype for spreadsheet sonification, the details and evaluations of which were reported at last year's ICAD conference (2). The second phase of the project began with a larger requirements gathering stage, in which the prototype from phase 1 was used to demonstrate the concept and as a catalyst to explore further ideas for development. In addition to the 8 users involved in stage 1, a further 10 visually impaired users were involved during this process, all of whom had had some exposure to spreadsheet systems, but only 2 of whom regularly used them without additional human support. All 18 users were sufficiently experienced within their problem domains to have an adequate to good understanding of the type of data analysis required, but because the majority of them preferred to carry out their analysis with the help of secretarial support, most users had interacted with spreadsheets to perform only relatively straightforward tasks.

The requirements gathering process for the 2nd stage of the project comprised a two-part interview with each participant: the first to establish the data analysis requirements for that user, and the second to discuss specific issues and preferences they had regarding the use of spreadsheets.

A consolidated list of requirements was established from this process, which is summarised below:

- i. User Requirement 1 (UR1) – To simplify the process of range selection, to seek to reduce the number of errors in the process of highlighting cells.

- ii. UR2 – Sonification of compound ranges, i.e. row or column expressions such as a - b + c. These compound values should be immediately recognisable in the audio from base values, by using a different instrument.
- iii. UR3 – Multiple range sonification: to be able to sonify one cell range followed by another, or sonify corresponding values from ranges adjacent to one another.
- iv. UR4 – Spatial data sonification where space is used either to map part or all of the spreadsheet structure, or where the spatial location of a sonified cell is based on its value.
- v. UR5 – Cell-wise sonification: provide immediate (manual) mode sonification by sonifying each cell as the user cursors to it.
- vi. UR6 – Support for two logical cursors, one used for controlling the speech output of a screen reader, the other for sonification.
- vii. UR7 – Sonification of some character-based information, typically where a small set of values are used in a particular column, e.g. "good", "satisfactory", "bad". Users should be able to specify how this is done interactively, i.e. allocate pre-programmed sounds that may be referenced using mnemonics.
- viii. UR8 – Interactive change of sonification parameters such as length of tones, gap between tones, frequency range etc
- ix. UR9 – Interactive specification of bookmarks for use in specifying cell ranges and to assist orientation and navigation back to cells of interest.
- x. UR10 Alarms – to be able to specify upper and lower thresholds so that values occurring outside these thresholds would be easily identified
- xi. UR11 Boundary specification – to be able to specify sub-areas within the spreadsheet of particular interest. This is to enable the frequency range to be re-calibrated to the maximum and minimum data values appearing in the sub-area, in order to make maximum use of the available frequency range to improve resolution between data values.

3. DEVELOPMENT

The second phase prototype was constructed within the Excel-VBA environment, using CSOUND as the primary means of sound production.

The level 1 Data Flow Diagram (DFD) below describes the system, showing that once the user has activated the main menu by pressing CTRL + K the system automatically gathers data from the active spreadsheet (which is the information generator). Then any changes that the user makes to these settings are stored within the system. At this point there is a wide range of different types of auto-sonification that the user could carry out, but they all share common characteristics that are shown as one auto-sonification path in the DFD. This shows that they all require a range or ranges from the spreadsheet to be selected, and then this information is turned into score (.sco) and orchestra files (.orc) and then shelled out

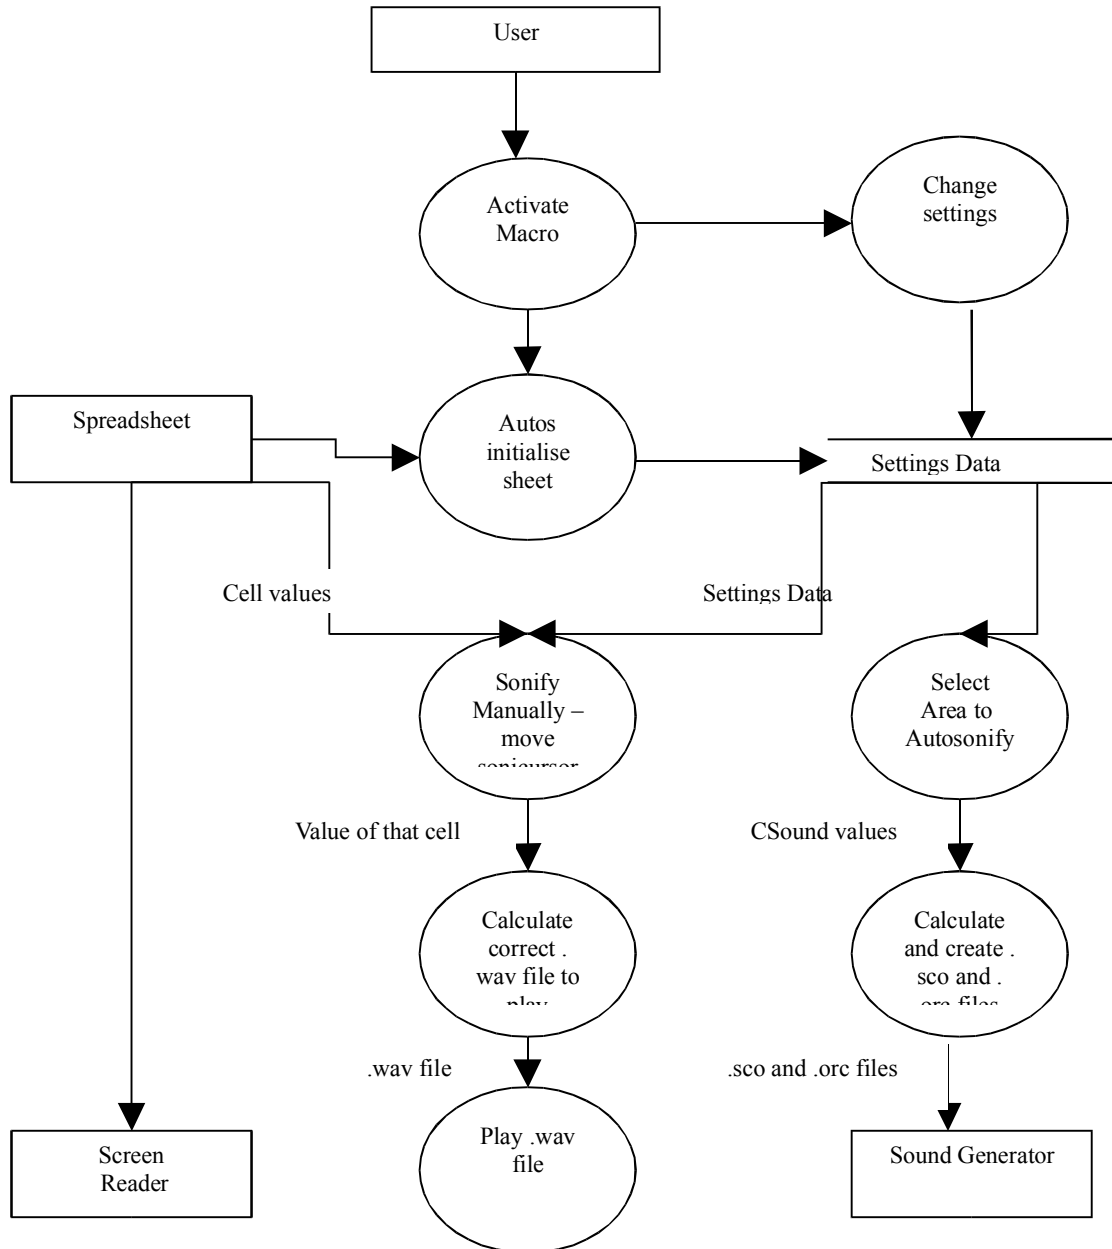


Figure 2. Data Flow Diagram of the second phase prototype

to CSOUND to actually generate the sound. If the user opts to go down the manual sonification route, then the system must calculate the correct pitch from 100 pre recorded .wav files and then render the appropriate file.

3.2. Auto Sonification

By selecting Auto Sonification a basic auto sonification can be performed from a small sub menu, which allows the user to specify the range and then choose to sonify, either by row or column. Because phase 1 evaluations indicated that users wanted something to facilitate the easy selection of arbitrary cell ranges that are not complete rows or columns [2], the user is now able to select the range from within the interface using just 5 or 6 keystrokes rather than by having to hold the shift key and highlight using the cursor keys. Providing the ability for users to specify the range in a cell:cell format was an

attempt to reduce difficulties of overshoot and undoing highlighting users reported in the usual, screen-reader based method of highlighting cells holding down the shift key and cursoring.

3.3. Manual Sonification

As a core requirement of phase 1 evaluations, the new prototype gives the option to manually sonify cells as the spreadsheet is navigated. On selecting this option from the main menu, a label switches from off to on and an audio confirmation speaks the status change. The user can now sonify cell by cell using 4 keys on the keyboard that mirror the layout of the 4 arrow cursor keys. The full range of keys is as follows:

E = up
D = down

S	=	left
F	=	right
Q	=	move the sonicursor to the active cursor
P	=	move the active cursor to the sonicursor
U	=	unlock the two cursors
L	=	lock the two cursors together

The user can also control the cursor with the existing cursor keys using their right hand, and as they cursor the screen reader will speak the value of the cell and its reference. In this way the right hand is controlling the cursor as normal, whilst when the left hand moves the cursor it sonifies the cell. In this way it is possible to navigate the spreadsheet with ease taking advantage of both auditory and non-auditory sound. By unlocking the two cursors it is possible to have a separate sonification cursor for sonification in addition to the usual ActiveCell cursor for spoken values, and these can then be steered around the spreadsheet independently. They can also be rejoined at any point by either snapping the sonicursor to the active cursor or vice versa. They can then be locked together again if desired.

For example, if a user wishes to examine manually sets of cells associated with specific column or row headers, this can be done by moving the active cursor to the first header of interest to get its value spoken by the screen reader, and then sonifying the values associated with the header using the sonification cursor. The user would then navigate to the next header and repeat the process.

Alternatively, the analysis process may be driven from the sonification cursor. A user may sonify down a set of exam results, looking for particularly low or high marks, or marks that came above or below pre-specified alarm thresholds. Having found an exam result of interest, the user can use the active cursor to identify the student by moving to the name column so that the screen reader speaks the required name. Should the name column be some distance from the exam marks column, the user may choose to detach the active cursor to check the name, before snapping the active cursor back to the location of the sonification cursor to continue the sonification of exam marks.

As we are dealing with modes here it is important that it be easy for the user to know which mode they are in at any time. For this reason there is a status update button on the main menu. When pressed it speaks the status of the manual sonification mode (on/off) and whether the cursors are locked or unlocked.

3.4. Bounded Sonification

On entry to a spreadsheet the system automatically calibrates the available frequency range to the maximum and minimum numeric values in the active area of the spreadsheet. If, however, one wishes to examine values which fall within a subset of the entire range of numeric values, and these areas are localised to a definable cell range, then a boundary can be set within which the system recalibrates to optimise its use of the frequency range for values found in the defined range. The effect can be likened to applying an auditory lens to a specific set of values to be compared. If users navigate outside of the defined boundary, this is indicated by a harsher square wave sound.

3.5. Compound Sonification

This option allows for the selection of up to 3 different ranges and for the application of 2 different operators to these ranges, for example $a + b - c$. The types of operators available are plus, minus, multiply and divide. This enables the automatic sonification of virtual rows/columns based on calculated base ranges. For example, quarterly sales values can be displayed by sonifying the sum of 3 columns of monthly values.

3.6. Multi Sonification

Here the user can select 2 ranges and allocate a different sound style to each one. The resultant sonifications can then be played back one after the other (sequentially), at the same time (synchronously) and with pairs of values adjacent to one another (interleaved (interleaved style still to be implemented)).

3.7. Bookmarks Feature

Bookmarks provide a way to remember key cell ranges for later use. Pressing CTRL + B at any point whilst browsing the spreadsheet will bring up the bookmarks menu. This will allow the user to bookmark up to 4 cell ranges. Then when selecting a sonification from the Advanced Sonification screen the user can choose to select a cell range by picking cell addresses from the previously saved bookmarks.

3.8. Note Length and Frequency Range

Sliders are provided to enable users to adjust the lengths of notes, pause between notes and maximum/minimum limits of the frequency range used for sonification. These sliders can be operated via a mouse, through cursor keys or by typing specific values into edit boxes.

3.9. Alarms and Sonification of Text

Alarms offer the ability to set upper and lower threshold values that sound a minimum or maximum alarm sound if the value in a cell is over or under the level. In a similar way the user can opt to select up to three text values to be sonified by pushing the button marked sonify text and entering the details of the text string of interest and the sound to be associated with it.

4. EVALUATION

15 of the 18 users who had provided formative feedback during the 2 phases of the project were involved in individual evaluation tests using their own spreadsheet data. The format of these individual evaluation sessions was as follows:

1) A training session first introduced the users to the facilities in the new version of the prototype. This started by the evaluator firstly demonstrating the functionality, and as the session progressed the participant increasingly interacted directly with the system. The durations of these sessions range from 45 minutes to 2 hours. During the user interactions with the system the evaluator noted incidents of interest such as user errors, breakthroughs in user understanding or difficulties in conveying particular ideas, feedback from users. Where possible examples used spreadsheets from the users own work

domain, though this was limited in a few cases due to the fact the users' spreadsheets were substantially textual, whereas the majority of the features developed so far support sonification of numeric data.

2) Following the session a debriefing interview took place in which the evaluator and user discussed the features of the prototype, relevance to the users needs for data analysis as well as more general issues relating to spreadsheet interaction and information processing.

Additionally, the system was presented to the Annual General Meeting of the UK British Computer Association of the Blind (BCAB), where feedback was elicited from a further 25 visually impaired computer users both in plenary in response to the presentation and with 14 individual users in mini demonstration sessions subsequently. The principal findings from these evaluations are summarised below:

- i. Automatic sonification of ranges of cells was a straightforward feature for most users to understand. Most users did take a few sonifications to be able to follow properly the row or column-wise structure of the sonification, and be able to extract notable differences between rows and columns being sonified. The ability to interactively adjust the lengths of notes and pauses used in each sonification was valuable at this stage, as it enabled users to adjust the overall pace of the displays while they became accustomed to assimilating information from the sonification.
- ii. The cell:cell format for specifying ranges (as opposed to highlighting) typically cut by at least a third the time required to select ranges in situations where users knew the relevant cell coordinates, and this saving increased with range size. In situations where users needed to locate cells, the cell:cell format was still seen as preferable where users had to repeat the same tasks across different spreadsheets.
- iii. Manual sonification was found to be straightforward to understand, and proved useful for exploring data that is mostly numeric. Again the speed advantage of this approach improved the larger the set of values being explored, and where the task requirements were a good match to the general characteristics of non-speech sound (e.g. identifying outliers or patterns as opposed to finding exact values). In conjunction with manual mode, the alarms mechanism was commonly used to identify outliers, while the boundary mechanism was used less often, but occasionally employed to zoom in on a specific area of interest. Most users found the idea of the threshold alarms straightforward to understand, while the boundary mechanism, although well understood and appreciated by some users, proved hard to grasp for about a third of the users.
- iv. The sonification cursor and the provision of dual and independent modes for sonification and ActiveCell cursors took longer for users to understand. Some users simply didn't need to perform tasks sufficiently complex to warrant the use of this mechanism. For several users however this proved extremely valuable, typically where the task involved outliers or other data phenomena, which, having been identified using sonification, required the user to look up other details (using the ActiveCell cursor and jaws speech) without moving the sonification cursor. With about 15 minutes training and practice, users with this type of task were able to reduce task completion times by around 50%. A case can be made for simpler information where cursoring is based solely on the usual single set of cursor keys, where these are toggled between activating jaws and sonifying values. The implementation described here does provide the advantage of efficient relocation of one cursor to the other where the analysis task has required them to become some distance apart. This is particularly useful when it is required to relocate one cursor to the other, where the target cursor is not located at or near the beginning or end of a row or column. Some useful suggestions were made concerning the separate use of the two cursors in fixed positions relative to one another.
- v. The bookmarks mechanism was found mainly to be useful in conjunction with manual sonification, where users could simply mark a point without having to be concerned about cell coordinates and then use the bookmark later in a range specification. Bookmarks also help to address part of the problem identified earlier concerning remembering cell locations. The requirement to remember which bookmark has been used to mark which point of interest is less demanding than having to remember which pair of cell coordinates identifies a point of interest. The usability of bookmarks could be further improved by associating descriptions with each bookmark. A further interesting suggestion was to expand the area definable by a bookmark to be greater than an individual cell, i.e. bookmark areas.
- vi. Sonification of text was found to be more popular than anticipated. Even in the limited way it is supported in the current system, several users found applications for this, and said that in conjunction with the manual mode this improved their ability to locate items of interest rapidly. Some of these situations are actually an overclaim, as users sometimes hadn't thought of employing typical text search mechanisms (e.g. control-f) within the spreadsheet, when this was pointed out to them they typically said they associated this with word processing. In spite of this, because the text sonification simultaneously checks for a number of values, there are certainly situations where it provides gains.

5. CONCLUSIONS

The interactivity of the system described here is the central point of interest and the main contribution of this research. Sonification can provide a more immediate and in some ways intuitive representation of numeric data than speech-based screen readers. The work here demonstrates that automatically generated sonifications can help to provide overviews of data in an intuitive way that is less cognitively demanding than trying to remember spoken numeric values. This is particularly the case when interaction with the system is tailored to user characteristics (in this case use of the cell:cell format for range specification).

Manual sonification can speed up the process of examining groups of cells for particular features such as outliers or trends. The use of a separate cursor for sonification can help to integrate the use of non-speech with the use of a speech-based screen-reader, as it becomes possible to allocate the two components of the auditory display to the tasks for which they are best suited. Manual exploration in sound can be further supported through the alarm and boundary mechanisms described. The sonification of text played a more significant role in evaluations than was expected, and how much can be achieved through the extension of this mechanism appears to warrant further work.

Future work will include implementation of some of the identified requirements not yet supported, such as multiple, interleaved sonification of ranges and particularly spatial sonification.

Interactive control of the sonification [19] can be considerably improved by removing the reliance on CSOUND and generating all sonifications using pre-recorded sounds.

Interactivity could then be enhanced by providing users with tape recorder like controls over automatic sonifications, including repeated playback at different speeds, pause, reverse and forward wind, as well as being able to step forward or back at different rates. Modifications to these pre-recorded sounds can be applied to map data characteristics of interest using volume, frequency, effects and spatial orientation parameters available in the Microsoft DirectX library. These sonification design options would provide opportunities for exploring some of the perceptual and cognitive issues raised in studies by Walker et al in the design of auditory graphs [17, 18].

Acknowledgements

I would like to extend my warm thanks to the members of the UK BCAB for their most helpful cooperation, and to all the participants who gave of their time during the stages of this project.

6. REFERENCES

- [1] B. N. Walker, and J. T. Coltran, "Sonification Sandbox: A Graphical Toolkit For Auditory Graphs," in *Proc. of the International Conference on Auditory Display*, Boston University, Boston, MA, USA, July 2003. pp.161-164.
- [2] T. Stockman, "The Design and Evaluation of Auditory Access to Spreadsheets," in *Proceedings of the International Conference on Auditory Display*, Sydney, Australia, 2004.
- [3] G. Kramer, B. Walker, T. Bonebright, P. Cook, J. Flowers, N. Miner, and J. Neuhoff, "The Sonification Report for the National Science Foundation," *International Community for Auditory Display (ICAD)*, 1998.
- [4] M. Matthews, "The Auditory Brain," in P.R. Cook, *Music, Cognition, and Computerized Sound an Introduction to Psychoacoustics*. MIT Press, pp. 11-20, 2001.
- [5] G. Kramer, "An Introduction to Auditory Display" in G. Kramer, *Auditory Display*, Addison-Wesley, pp. 3-77, 1994.
- [6] S. M. Williams, "Perceptual Principles in Sound Grouping" in G. Kramer, *Auditory Display*, Addison-Wesley, pp. 95-125, 1994.
- [7] T. M. Madhyastha, D. A. Reed, "A Framework for Sonification Design" in G. Kramer, *Auditory Display*, Addison-Wesley, pp. 267-287, 1994.
- [8] A. Dix, J. Finlay, G. D. Abowd, and R. Beale, *Human-Computer Interaction*, Prentice Hall, 2004.
- [9] D. L. Mansur, "Graphs in Sound: A Numerical Data Analysis Method for the Blind", *Unpublished Masters Thesis*, University of California, Davis, USA, 1984.
- [10] S. A. Brewster, L. M. Brown, M. Burton, R. Ramloll, and B. Riedel, "Design Guidelines For Audio Presentation of Graphs and Tables," in *Proc. of the International Conference on Auditory Display*, Boston University, Boston, MA, USA, July 2003.
- [11] S. A. Brewster, L. M. Brown, "Drawing By Ear: Interpreting Sonified Line Graphs," *Proc. of the International Conference on Auditory Display*, Boston University, Boston, MA, USA, July 2003.
- [12] S. Brewster, L. Brown, R. Ramloll, W. Yu, "Browsing Modes For Exploring Sonified Line Graphs," *Proc. of the 16th British HCI Conference London*, vol. 2, pp. 2-5, September 2002.
- [13] S. Brewster, R. Ramloll, B. Riedel, W. Yu, "Using Non-speech Sounds to Improve Access to 2D Tabular Numerical Information for Visually Impaired Users," *Proc. of the HCI 2001 Conference*, 2001.
- [14] S. Bly, "Sound and Computer Information Presentation" *PhD Thesis*, University of California, Davis, USA, 1982.
- [15] E. M. Wenzel, "Spatial Sound and Sonification," in G. Kramer, *Auditory Display*, Addison-Wesley, pp. 185-221, 1994.
- [16] E. M. Wenzel, M. Arruda, D. J. Kisller, and F. L. Wrightman, "Localisation using nonindividualised head-related transfer functions," *Journal of the Acoustical Society of America* 94, pp. 111-123, 1993.
- [17] B. N. Walker, G. Kramer, and D. M. Lane, "Psychophysical Scaling of Sonification Mappings," in *Proc. of the International Conference on Auditory Display*, 2000.
- [18] B. N. Walker, and D. R. Smith, "Tick-Marks, Axes, and Labels: The Effects of Adding Context to Auditory Graphs," in *Proc. of the International Conference on Auditory Display*, Japan, 2002.
- [19] International Workshop on Interactive Sonification, 2004. www.international-sonification.org.