

THE DESIGN AND EVALUATION OF AUDITORY ACCESS TO SPREADSHEETS

Tony Stockman

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

ABSTRACT

This project is concerned with the development and evaluation of an integrated approach to sound-based access to spreadsheets. The primary target user population is visually impaired users, employing speech-based screen readers as their main means of accessing the data. Other users may benefit from the approach in situations where sound may provide an effective alternative or complementary means of exploring complex data relationships, when the visual interface being employed is too small to afford effective spreadsheet display or in monitoring situations where the visual attention of users must be employed elsewhere.

The paper begins with an introduction to the use of spreadsheets by visually impaired users, including the current state of auditory interfaces as provided by commercial screen reader (SR) systems. The development of a prototype system to enhance speech-based access through the use of data sonification is described. The results of a number of detailed evaluations are discussed, drawn from a number of "real work" situations. The paper concludes with an examination of the issues arising from the evaluations, and a program of future research.

1. INTRODUCTION

Screen readers (1) are the primary mechanism used by visually impaired computer users to access computer systems. Although screen readers exist for other systems such as the Macintosh and Unix, by far the most widely used screen readers, Jaws For Windows (JFW) (2) and Window-Eyes (WE) (3) run under the Microsoft Windows operating system. Both JFW and WE support output in both Braille and speech, but for a number of reasons, most users of screen readers either use speech output mode alone, or use Braille only in combination with speech. The reasons for the general preference for speech output can be briefly summarised as follows:

1. The additional hardware required for Braille output is relatively expensive, and is an additional component to be carried in mobile use
2. A significant proportion of visually impaired computer users are not fluent in reading Braille

In contrast, synthetic speech represents a relatively inexpensive option, which can often be provided via a built-in sound card, and which can deliver output at easily variable

speeds, which are only limited by the users' perceptive ability (4).

Spreadsheets are an extremely common tool for the storage and analysis of data, used throughout industry, commerce, government and education. The multi-dimensional nature of spreadsheets poses a number of problems for speech-based screen reader access, due to the temporal nature of speech and the difficulties of obtaining any kind of overview of the data being examined. Recent releases of SR systems such as JFW and WE have provided some increased support for spreadsheet-related tasks, but independent examination of spreadsheet data remains cumbersome and cognitively demanding for most speech-based SR users.

These difficulties can only partially be offset through the use of Braille (for the relatively small number of users for whom this is an option), since commercial soft Braille displays support presentation primarily in one dimension only. A few commercial Braille printers provide support for producing hard copy graphs of data displays, but this is generally an expensive and inflexible solution.

Research into the use of non-speech sound to convey information, known as data sonification, has grown rapidly since the early 1990s (5, 6). Data sonification holds considerable promise as a means by which visually impaired computer users can obtain improved access to data. The ability to overview large, possibly multi-dimensional data spaces rapidly, plus the possibility of using sonification to pinpoint segments of data for detailed examination, may lead to considerably reduced times for data exploration tasks.

This project is motivated out of the ongoing difficulties SR users encounter in obtaining effective access to spreadsheets, in spite of their being such a widely used tool in education and employment. Some of the research questions which the project seeks to address are as follows:

- 1) What means are currently employed by visually impaired users to address the shortcomings of current speech-based access to spreadsheet data?
- 2) Can these mechanisms be systematised in order to provide a partly or wholly structured auditory interface to spreadsheets?
- 3) Can data sonification provide an effective complement to speech-based access in real work contexts, and if so what are the characteristics of tasks where sonification is most effective? What are the nature of any improvements in task execution (e.g. in

terms of reduced task completion times and/or improved understanding of data)?

- 4) What are the issues in combining speech and non-speech access to data, and how can these two modes of access be unified to maximise the advantages of each mode and minimise their drawbacks.

In the first phase of the project, described here, we focus on exploring whether sonification can enhance spreadsheet access in the context of speech-based SR access.

2. SCREEN READERS AND NON SPEECH SOUND

The results of data sonification research have, so far, seen little take up in commercial SR technology. The release of JFW version 5, in September of 2003, represented the first significant use of non-speech sound in a commercial screen reader.

The recent release of JFW version 5 (7) contains limited use of non-speech sound to customise the feedback provided to users about the current state of the interface. Examples of the ways in which non-speech sounds may be used include:

- to identify when focus has moved to a particular type of interface widget
- to signal the different states of a checkbox
- to indicate upper/lower case
- to indicate the degree of indentation
- to identify the values of HTML attributes.

The inclusion of this limited support for the use of non-speech sound is intended to reduce the time and improve the accuracy of user tasks, for example when navigating between objects in the Windows GUI, or for carrying out screen-based proof reading of lengthy documents.

Data sonification has considerable potential to enable visually impaired people to have improved access to data (8, 9). Data sonification is potentially useful to visually impaired people in at least the 2 following data analysis tasks:

- 1) In providing a rapid overview of a data set. It seems likely that data can be heard rather faster using sonification than it can be browsed using speech and/or Braille.
- 2) Identification and navigation to points or areas of interest. It is usually the case that speech or Braille will be required at some point during data analysis to obtain specific values. However, a well designed sonification system could reduce the amount of data to be navigated in speech/Braille, both by helping identify the area of interest in the first instance, and if required, by providing a second, more localised sonification of the immediate data range of interest.

The recent inclusion of non-speech sound in JFW does not directly address these overview and data navigation tasks, though a mechanism to provide this support has been available for some time via JFW's scripting language. The following section describes how support for the above tasks has been prototyped within the Microsoft Excel spreadsheet environment. Issues arising from the development and evaluation of the prototype are discussed in the remaining sections of the paper.

3. PROTOTYPE DEVELOPMENT

Initial discussions with visually impaired users determined that by far the most common source of complex data they encounter is Excel spreadsheets. The most obvious issue in navigating and analysing spreadsheet data is simply the length of time required when using a speech and/or Braille interface. The objective of the first phase of the project was therefore determined: to explore the value of data sonification in speeding up the overview and data navigation tasks described in section 2 above, evaluating the approach within the context of a speech-based screen reader, in realistic, work-based situations.

3.1. Software Environment

The system was built and tested on a Toshiba laptop running Windows XP. Interactivity was encoded using a combination of Visual Basic for Applications (10) and the JFW scripting language (2), whilst Jaws version 4.51 was used as the screen reader. Csound version 4.23 (11) was used as the sonification engine, after Flowers (12). Csound uses two files in order to synthesise sound: an orchestra file, which defines the instruments to be used in the sonification, and a score file, defining the notes to be played by each instrument (note that in recent versions of Csound, these two files have been merged into one, integrated file format).

Within the prototype, A pre-written orchestra file is used to specify the instruments to be "played" during sonifications. An enhanced version of the system could include the functionality to change the orchestra file dynamically during spreadsheet interaction. Each line of the score file corresponds to a note to be sonified. In the prototype developed in this phase of the project, the start time and frequencies of notes in the score file are determined by scaled values of data items to be sonified.

3.2. Operation of the Prototype

An initialisation macro performs a pass through the data in the worksheet, identifying global maximum and minimum numeric values, which are used to calculate offset and scaling parameters, used to map the raw data into proportional frequency values within a comfortably audible range. When a user requires to sonify a range of cells, the required macro is invoked using the corresponding "hot-key" combination. The macro steps through the set of cells to be sonified, for each cell writing a line to the score file, the frequency value of which is calculated by offsetting and scaling the raw data value of the cell. The prototype provides the following functions:

- 1) Sonify a selected range of cells either row-wise or column-wise, respectively with end of row or column indicators, with optional move to global max or min within the selected range.
- 2) Sonify current column, current row, rest of col, rest of row, previous part of row or col.

4. EVALUATIONS

The evaluation of the system, which is still in progress, has so far included 8 visually impaired users, using 3 different spreadsheet applications. Major difficulties arise in gaining access to significant numbers of such a specialised user population. It was felt in these circumstances that it was

essential to build up a detailed, case by case understanding of user requirements and evaluation of the prototype, to make optimum use of the number of participants available.

User U1 has evaluated the system using application A1, which is a spreadsheet of children and families social work data, drawn from this user's real work context. Similarly user U2 has used spreadsheet application A2, a university admissions spreadsheet, again drawn from that user's work domain. Applications A3 and A4 both contain results of students' assessments, detailing the module performance of students in courseworks and examinations. Application A3 has been used by users U3-U5, and application A4 by users U6-U8. Student numbers and names have been removed from both of these applications.

4.1. Evaluation Method

7 of the 8 users reported that they used a support worker or secretarial support to assist them in examining spreadsheet data. For those users less familiar with Excel (see section 4.2) this was the only means they used to examine the data, whereas for 3 of the 4 users more experienced with Excel, it formed a complementary, though important means of data examination. Whenever possible this process was observed, in order to understand the methods used to co-explore the structure and content of spreadsheets.

An initial preparation session of about 30 minutes was used to do the following:

- 1) Ascertain each user's previous experience of spreadsheet applications in general, and Excel in particular.
- 2) Explain the built-in keystrokes used in the Excel-JFW environment for selecting rows, columns and arbitrary data ranges.
- 3) Explain and demonstrate the basic idea of sonification, and show users how to invoke the macros to perform the available sonification functions described in section 3.

Users were then given an amount of time determined by themselves to explore the spreadsheet used for their evaluation, employing a combination of sonification and spreadsheet information spoken by Jaws. During this time they were free to ask questions about operation of the system and make comments about its use. The times for these exploratory sessions varied from 15 to 25 minutes. The evaluation session concluded with an informal discussion of each users experience of using the system.

4.2. Evaluation Results

4.2.1. Co-Exploration

The objectives of these observations are two-fold:

- 1) To develop an understanding of the detailed requirements of each users spreadsheet access needs
- 2) To gather information that may be useful in the next phase of the project in constructing an interface which may partly or wholly automate the support workers role in the co-exploration process.

Most of the observed co-exploration sessions took place with the sighted assistant working from a hard copy of the spreadsheet, and the visually impaired user making notes but not interacting with a machine-based version of the spreadsheet. Some commonalities emerged concerning the way spreadsheets would be co-explored as follows:

- 1) Important header information and the names and order of columns are established
- 2) Row titles are identified
- 3) The availability of row/column totals and other statistics would be ascertained.

A row or column-based examination of specific data items would follow, often leading to annotation of the hard copy of the spreadsheets and/or notes being made by the visually impaired user to capture questions or facts about the data.

4.2.2. Background

All the users involved so far possess a competent knowledge of the JFW screen reader, using it regularly for work and leisure. All users said that Excel was the only spreadsheet application they had used.

Users U1, U2, U3 and U6 possessed a similar level of knowledge of Excel, having used it from time to time in a basic way, to review spreadsheet data and being comfortable navigating to specific cells. Users U4, U5, U7 and U8 reported they had very little knowledge of Excel, avoiding its use where possible. All participants, excluding one of the more competent Excel users, said they used additional clerical support for spreadsheets containing data important in their jobs, and expressed reservations about their ability to extract the information they required independently and efficiently.

4.2.3. Interaction with the Prototype

The main points arising from the free form spreadsheet exploration sessions can be summarised as follows:

- 1) All users seemed to develop an understanding of the idea of using sonification to overview data, and appreciate how it could complement spreadsheet exploration with a screen reader.
- 2) The up-up polarity mapping of data values to frequency seemed to be a natural one for the variables in the 3 applications employed here.
- 3) The ability to be able to select entire rows or columns and sonify them directly provided a natural means of over-viewing relevant subsets of data. There occurred however a number of situations where users would have liked sonification of some compound combination of rows or columns, such as the quarterly sum of 3 monthly columns, or the difference between target and achieved results.
- 4) The ability to be able to sonify arbitrary ranges of cells was useful. A particularly striking use of this occurred in application A1, where U1 employed it to verify that the values in 4 related columns were the same across a range of rows. In general however this feature was less frequently used than the examination of entire rows and columns. This may be partly

attributable to the relative ease with which entire rows and columns can be selected, keyboard-based selection of arbitrary ranges is rather slower than with a mouse, and might be better supported in future versions of the interface.

- 5) 3 Users commented that it would be nice to temporarily reorder columns and/or rows to facilitate highlighting for sonification. This mechanism could be provided in future versions either directly, or via the more keyboard friendly mechanism for selecting cell ranges eluded to in 4 above.
- 6) Users occasionally forgot to collapse the selection of the currently selected range before selecting the next range for sonification. This may be avoidable by providing some form of background sound indicating that a range is currently selected, to provide an equivalent to visual highlighting for sighted users.
- 7) In general, the more values involved, the more substantial was the advantage obtained from using sonification. For example, in the evaluations using application A3, sonifications of all 134 students results, or of the results of specific examination questions saved substantial time in screen reader navigation. A precise quantification of this gain will of course vary depending on the pace of interaction between any given user and screen reader. From observations during these experiments however, experienced users would typically review cell data values at a rate of approximately 3 per second, whereas sonifications were delivered at the rate of 5 values per second.
- 8) Users commented that the sonifications helped to provide an overview perspective of the data in a way that screen-reader-based review does not, i.e. it seems easier, at least for a short time, to remember the "shape" of the data as conveyed by the sonification, than it is to remember the "shape" of the data as conveyed by the list of numbers spoken by a screen reader. The ability to comprehend the sonified overview was illustrated by the following examples:
 - a. User U1 being able subjectively to compare the numbers of referrals in different months of the year in application A1,
 - b. User U2 being able easily to perceive that the University frequently exceeded its targets when recruiting to courses.
 - c. Users U3-U5 being able to compare the relative difficulty students had with different exam questions
 - d. Users U6-U8 being able to compare relative progression rates across different modules.

5. SUMMARY AND FUTURE WORK

The basic prototype and completed evaluations demonstrate that there is value in using sonification in real-world applications to convey overviews and support navigation of complex data. The advantage that this brings increases with the amount of data to be analysed, and the complexity of relationships in the data.

Improvements to the phase 1 prototype arising from the evaluations include: sonification of user-selected compound values of data, simplification of the keyboard-based selection

of arbitrary data ranges, and placement of more of the data to sound mappings under user control.

Major elements of the program of work in phase 2 of the project will include: examination of approaches to providing automated support for co-exploration of spreadsheets, further characterisation and quantification of the benefits data sonification can bring to spreadsheet analysis and a more detailed examination of the issues concerned in integrating speech-based screen access with sonification.

6. REFERENCES

- [1] Royal National Institute of the Blind technology homepage: <http://www.rnib.org.uk/technology>.
- [2] Freedom Scientific homepage: <http://www.freedom-scientific.com>.
- [3] GWMICRO homepage: <http://www.gwmicro.com>.
- [4] C. Asakawa, H. Takagi, S. Ino, and T. Ifukube, "Maximum listening speeds for the blind," in *Proc. of the 9th Int. Conf. (IC) on Auditory Display*, Boston, Ma, 2003, pp. 276-279.
- [5] G. Kramer, *Auditory Display, Sonification, Audification and Auditory Interfaces*, Addison Wesley 1994.
- [6] G. Kramer et al., *The Sonification Report*. <http://www.icad.org/websiteV2.0/References/nsf.html>.
- [7] Introduction to JFW release 5: http://www.freedom-scientific.com/fs_products/software_jaws5intro.asp.
- [8] L.M. Brown, S.A. Brewster, R. Ramloll, M. Burton, and B. Riedel, "Design guidelines for audio presentation of graphs and tables," *Proc. of the 9th Int. Conf. (IC) on Auditory Display*, Boston, Ma, 2003, pp. 284-287.
- [9] T. Stockman, "The Sonification of SQL constructs in multi-modal database interaction," *Proc. of the 9th Int. Conf. (IC) on Auditory Display*, Boston, Ma, 2003, pp.308.
- [10] Microsoft developer's network: <http://www.msdn.com>.
- [11] Csound support website: <http://www.csounds.com>.
- [12] J. H. Flowers, D. C. Buhman, and K. D. Turnage, "Data Sonification from the Desktop: Should Sound be part of Standard Data Analysis Software," *Proc. of the 3rd Int. Conf. (IC) on Auditory Display*, Palo Alto, Calif, 1996, <http://www.icad.org/websiteV2.0/Conferences/ICAD96/pr oc96/flowers.htm>.