

Interactive Sonification in two domains: helicopter flight analysis and physiotherapy movement analysis

Sandra Pauletto and Andy Hunt

Abstract—This paper describes the role of interactive sonification in the analysis of data from two task domains. Helicopter flight recordings produce large multi-dimensional data sets which are difficult to analyse visually, so this project provides a method of navigating the data as it is converted to sound. Likewise, in the physiotherapy domain, data is gathered from EMG sensors which record the firing of various muscles. An Interactive Sonification Toolkit has been developed to aid the analysis of general data sets. This paper describes the toolkit, and the early results of interactive sonic analysis of both example domains.

I. INTRODUCTION AND BACKGROUND

A. Multi-disciplinary approach

The project ‘Improved data mining through an interactive sonic approach’ represents a multi-disciplinary approach to improving the exploration of large sets of data by converting them into sound and putting them under the real-time control of a human analyst. The research team consists of academics at the Universities of York (User Interfacing and Digital Sound) and Teesside (Physiotherapy), and engineers at Westland Helicopters. The project was launched in April 2003 and is funded by the EPSRC (Engineering and Physical Sciences Research Council).

What at first may seem an unusual combination of disciplines yields unexpected deep-rooted similarities in the ways in which data are gathered, stored and analysed by humans. This project brings together the knowledge and expertise of this range of disciplines to solve a common data analysis problem.

B. The problem of visual data overload

Advances in computing technology and digital instrumentation provide unprecedented opportunities for gathering data from a rich variety of multiparametric sources. However, the major drawback is the difficulty in portraying the sheer amount of data that is available for analysis. This must be converted into a form that can be readily assimilated by the human analyst. The primary medium by which this information is traditionally portrayed is the Visual Display Unit (VDU), either in the form of numerical values, or displayed as a graphical image. The VDU is often assumed to be the only method for computer output, but there are some inherent problems with this approach. Visual displays (particularly those involving text and numbers) demand the full visual attention of the operator, and much of his or her

verbal and logical thought processing. This is at least distracting from the task in hand, and at worst dangerous.

C. Aural interaction

In everyday life humans do not gather all their information about the world via the visual sense. Some tasks are inherently more suited to the use of the aural sense. The view taken by this project is to work with the natural tendency of clinicians and engineers to expect acoustic feedback.

The medical community has used sound feedback for a number of applications. The stethoscope is still a fundamental medical tool, widely used for a range of tasks. Clinicians use such interactive sonic feedback as an everyday diagnostic technique, allowing them to detect complex time-based events (such as heartbeat and air-flow), and hidden structural defects; such is the power of sound as an analytical aid.

Sonification has the potential of being able to portray many more simultaneous data parameters than visual displays [1], whilst freeing the eyes and hands for visual-spatial tasks (such as surgery, or communicating with colleagues and patients). This requires the sound to be rendered in real-time, and to react instantly to changes in the data. Real-time interaction is not limited to tasks where the data are being gathered in real-time. Even when the data has already been recorded, an element of user-interaction is needed to sift the data, tune the parameters, and focus onto the areas of interest. It is our belief that dynamic user-interaction with multi-parametric data streams can yield an effective method for controlling and interpreting those streams [2].

II. INTERACTIVE SONIFICATION IN TWO DOMAINS

One of the goals of this research is to provide a novel toolkit for experimentation with the interactive sonification of large data sets gathered from complex systems. It will provide designers with new methods to optimise the mapping between the input data, the human interface, and the output sound.

A. Commonalities between the two task domains

The work carried out in this project is generally applicable to situations where many simultaneous streams of data are gathered which then need to be analysed to discover the underlying state of the system.

In order to demonstrate the potential of the toolkit, we have chosen two different application areas that have high-

dimensional data sets which require exploration. These areas have been selected due to the remarkable similarity between the existing ways that their data is currently gathered and processed (see Figure 1).




Figure 1: System diagram showing current visual analysis of gathered data

Multiple sensors gather streams of data from different parts of the system. Together these data streams represent the state of the system. Currently the streams are displayed as numerical lists, or as individual graphs. Human analysts at present struggle to diagnose the relationships between the separate variables, and they have restricted methods at their disposal for moving through the data.

The Interactive Sonification Toolkit allows the user to access the recorded data streams and convert them into sound (see Figure 2). The human operator can then control the navigation through the data (for example scanning backwards and forwards) whilst listening to the resultant sonification.

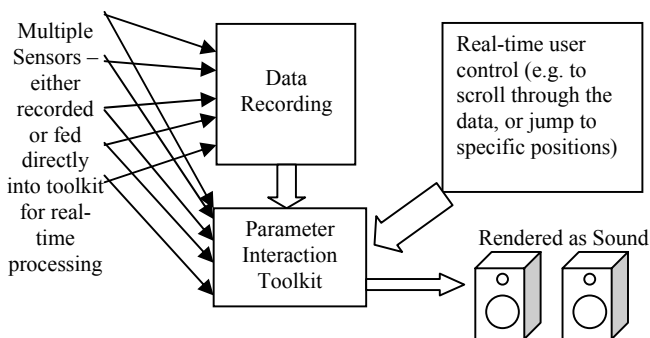


Figure 2: System block diagram

We now explain how each of the chosen target application areas will utilise the toolkit.

B. Physiotherapy movement analysis

Within the physiotherapy profession there are a few well-known 'gurus' who can locate and cure problems with ease, but are unable to explain to others how they do this. The majority of practitioners need extra information to aid their diagnosis of non-routine cases. Therapists are also being increasingly called to produce quantitative evidence of the progress of a particular condition (for example to assess the effectiveness of a particular form of treatment).

In order to gather such numerical data, sensors can be attached to the surface of the body to monitor muscle potentials in various locations over time as a task is carried out (e.g. standing up). Vast amounts of data can be easily collected in this way, but the problem comes in trying to analyse and make sense of the many large streams of numbers. Therapists are convinced that the data streams hold the key to understanding

the particular ailment, but they lack a tool for sifting, analysing and manipulating the data.

This project plans to provide a diagnostic tool to aid therapists in assessing a complex set of movements. Such a tool would be used for helping to diagnose complex conditions, and monitoring the progress of individual patients.

Therapists will use a set of sensors and a data acquisition system to gather data for a particular movement. The analyst will use the Toolkit to parse that data (without the patient having to be in attendance) by listening to a sound-plot of individual variable streams, or the interaction of two or more streams (e.g. the difference in tension between two muscles). Once the problem area has been located, the more conventional visual feedback will be used to study the graph plots and note down data values. In this way we are providing a *supplement* to existing analysis techniques.

However, we plan a further development of the system for use in real-time by patients to enable them to monitor their own muscle movements (in sound) as they attempt to carry out specific movement tasks. This will be particularly useful in the rehabilitation of stroke patients – who have to re-learn movements completely, and have little feedback to tell them what is going wrong inside their body. In this case, it might be possible to play the patient the sound of an 'idealised standing-up movement', for example. The patient can then try the task for themselves (many times if necessary), with sonic feedback given in real-time at all stages of the movement. The therapists would be able to give instructions such as "Try again, but make the sound less rough", which is innately easier for most patients than "Try not to put too much tension in the largest muscle of your leg, compared with the smallest".

C. Helicopter Flight data analysis

The handling of flight data and its subsequent analysis is an everyday activity for the designers of aircraft. An on-line data acquisition system is used to gather flight data from the pilot controls and a large number of sensors around the aircraft. Large multi-parametric data sets are generated as a result, and these are currently examined off-line using visual inspection of graphs. The graphs are printed out and laid across an open floor to allow the whole data set to be seen at a reasonable resolution. Engineers walk around this paper array looking for anomalous values and discontinuities in the signal.

We propose to improve the methods for handling such data, with the immediate benefit of saved technician time, and thus a faster turnaround in flight data analysis. Flight data is converted into sound, and sifted through interactively by the human analyst. This helps to pinpoint problem areas due to, for example, local non-linear instabilities, undesirable dynamics and flight control system failures, and will give the analyst a better feel for the relationships between the data streams.

III. A PERCEPTUAL APPROACH TO SONIFICATION

We aim to explore data sets from our two target domains using sonification techniques that exploit the inherent ability of the human auditory system to analyse sound and identify

and recognize features in it. These properties of the auditory system are well described by Bregman in the book *Auditory Scene Analysis* [3]. Bregman writes that the brain appears to have two main ways of analyzing, identifying and grouping features in sound. The first way uses primitive processes of auditory grouping, which are based on the way the auditory system is naturally built and are independent of learning and experience. The second method, schema-based grouping, relies on individuals' knowledge of familiar sounds and their experience of listening.

The toolkit developed so far creates sonifications which primarily appeal to the primitive processes of auditory grouping, since they are more generally applicable to a wide range of users. However, during the process of evaluation of the created sonifications, observations will be made on how easily a sonification is learnt and how the experience of hearing a particular sonification improves our ability to recognise features in it (see section VI).

IV. INTERACTIVE SONIFICATION TOOLKIT

This section reports on the first prototype of the Interactive Sonification Toolkit. This version is far from complete; however it has already been used to make observations on the gathered data in both task domains.

A. The Prototyping Environment

The programming environment used to develop the Toolkit prototype is the visual programming environment Pure Data (PD) [4] developed by Miller Puckette [5]. The reasons for choosing this environment are:

- it is cross-platform, working on Macintosh and PCs under Windows, OS and Linux operating systems;
- it allows real-time programming (no need for a compiling phase) which facilitates fast prototyping;
- it contains many built-in DSP (Digital Signal Processing) functions;
- it is free and open source;
- it is extendible: new specific functions and libraries can be programmed in C and added to the environment.

The Toolkit makes use of existing PD libraries. In particular, it uses the basic PD objects that come with the 0.36.0 version of PD, the Zexy [6] library 1.3 by Iohannes Zmoelnig, and the GriPD [7] library version 0.1.0 by Joseph A. Sarlo which is used to build the Graphical User Interface (GUI).

B. Data format and storage

For this toolkit every data set is represented by a matrix with various columns and rows. Each column of numbers in the matrix contains a particular time-stamped sequence of recorded data. In our two example task domains each column represents data gathered from one specific sensor. For instance, in the helicopter data there will be a column representing the speed of the rotor, another representing the current altitude etc. Sensors are normally very different from each other and measure various parameters. Their ranges and measurement units are often different. Within the Toolkit each column of data is referred to as a 'channel'.

The number of rows in the matrix corresponds to the number of samples in each channel. For example, if a sensor records data for 10 minutes at the rate of 2 samples per second, then there will be $(10 \cdot 60 \cdot 2 =)$ 1200 rows (samples per channel). The data set or matrix is written in a text file that can be edited using a simple text editor (see Figure 3).

```
matrix 5 2
1 3
2 4
1 45
1 43
4 46
```

Figure 3: Format of matrix for a data set containing 2 channels of 5 samples.

The first line of the file must contain the word *matrix* followed by the numbers of rows and columns. A return separates the lines. Separators between values can be either spaces or tabs. The file has to be saved with the file extension '.mtx'.

C. A description of the Toolkit and its functions

The user interface for the Interactive Sonification Toolkit consists of two main screens (PD "patches"): A *Data Scaling* page, and an *Interactive Sonification* page.

Data Scaling page

The function of this part of the program is to upload one or more data sets, and scale them appropriately prior to sonification. The screen shows details of the data set(s) loaded (see Figure 4).

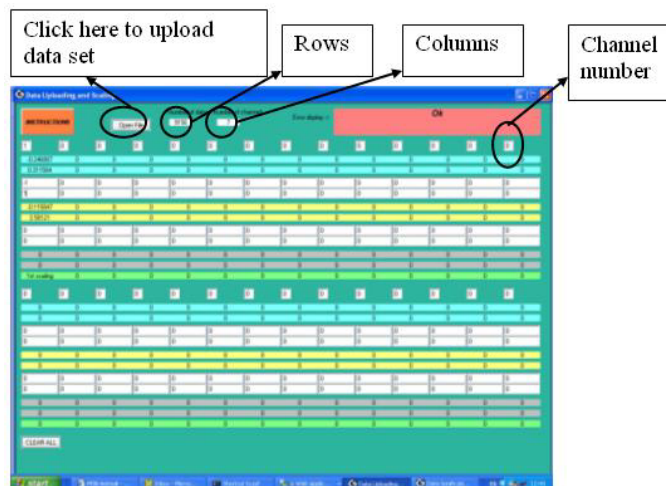


Figure 4: Data Scaling page showing 28 channels of data.

Within this part of the program the user can scale each channel separately (and opt to move to another screen to display the data of each channel as a graph). The Toolkit allows two different methods of scaling:

- defining new *minimum* and *maximum* values;
- defining new *transposition* and *stretching* factors.

If the user employs the first scaling type, the program automatically calculates and displays the corresponding

transposition and stretching factors, and vice-versa. Having the two types of scaling has proven to be very useful in the following situations:

1. When employing the second (stretching) scaling, the automatically generated minimum and maximum values can be used to determine whether the sound will be within the perceivable range.
2. When the user needs to display two channels using the same reference system (to allow direct comparison), then the second scaling has to be used because both channels need to be transposed and stretched by the same amount. However, the user wants to ensure that the resultant values are still perceivable. So the first channel can be altered (with type 1 scaling) to be within a perceivable range, then the associated transposition and stretching values can be imposed on the second channel.

Interactive Sonification page

This constitutes the second part of the Toolkit and has its own separate GUI window (see figure 5).

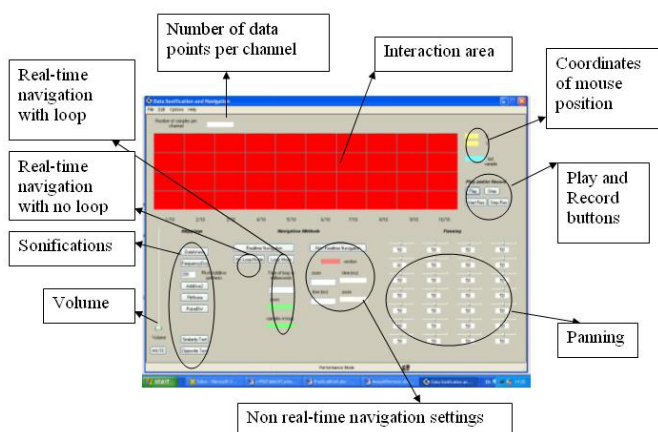


Figure 5: Interactive Sonification page showing main functional areas

When a data set is uploaded in the Data Scaling page, it appears in this page ready for sonification and navigation.

D. The sonification methods

Eleven different types of sonification have been explored to date. The first eight are currently implemented in the Toolkit, while the last three exist as separate prototypes.

- 1) *DataWave*: the rescaled data points are read directly as sound samples. This mapping is often referred to as *audification*.
- 2) *FrequencyOsc*: the scaled data points of a channel are fed into the frequency input of an oscillator, producing a sine wave with a pitch contour that varies with the data.
- 3) *Additive Synthesis with constant pitch*: the data points of each channel drive the amplitude envelope of a harmonic partial of a sound synthesised by additive synthesis, with a constant fundamental frequency (chosen by the user).

- 4) *Additive Synthesis 2*: the data points of one specified channel control the fundamental frequency of an additive synthesis sound; the other channels drive the amplitude envelopes of the other partials and thus control the harmonic content.
- 5) *Filter Noise*: Noise is played through a band-pass filter. One data channel controls the centre frequency of the filter, while a second channel drives the filter's 'Q' factor (CentreFrequency/BandWidth).
- 6) *Pulse*: one channel controls the frequency of a pulse waveform; while a second channel drives its bandwidth.
- 7) *Similarity Test*: the data of two channels are re-scaled between -1 and 1. One channel is then subtracted from the other (data point by data point) and the result (scaled to audible frequency range) fed into the frequency input of an oscillator. If the two channels' data are identical, then the frequency is zero and therefore no sound is heard.
- 8) *Opposition Test*: the data of two channels are rescaled between -1 and 1. The two channels are then added (point by point) and the result (scaled again to audible frequency range) fed into the frequency input of an oscillator. If the two channels' data are completely opposite to each other, then the frequency is zero and therefore no sound is heard.
- 9) *Data as notes durations*: the data points of one channel control the durations of successive notes, played on a MIDI synthesiser.
- 10) *Data as note pitches*: different instruments are assigned to different channels and the data points of each channel drive the pitches of a MIDI synthesiser.
- 11) *Pitch and duration*: One channel is mapped onto note durations, the other onto note pitch.

Left and right panning can be applied to sonification methods 1, 2 and 10, because these involve several independent sound streams. This can aid in comparing more than one channel, by placing them in different virtual positions in the stereo field.

E. Interactive navigation

In the Interactive Sonification Toolkit the sound can be navigated in three different modes: two real-time and one non-real-time. The interaction area (see the rectangle within the screenshot in Figure 5) represents the sound in its totality (with time portrayed from left to right). It can be explored by dragging a mouse over it. The co-ordinates of the mouse are tracked by the program and used to generate sound. A future version of the Toolkit will provide the option of displaying, as a graph, the data array in the interaction area, so that the user can have both an audio and visual display.

The navigation modes are as follows:

- 1) *Real-Time Navigation*: the user clicks the mouse and drags it within the interaction area. At each instant the co-ordinates of the mouse drive a pointer to the sound/data array. The sound is played in immediate response to the mouse movement, and therefore the

user can freely move backwards and forwards through the sonification and hear the result. (The y axis of the interaction area does not have a function in this mode). Note that sound is only heard when the mouse is moving. This is rather like dragging a piece of recorded magnetic tape past a tape head. Sound is heard only when the tape is moving, and the rapidity of sonic events is determined by the speed.

- 2) *Real-Time Navigation with Loop*: this time the mouse x position controls the start-point of a sound loop. The user inputs the length of this loop in milliseconds before navigation. Even when the mouse is not moving, a sound can be heard, as the loop plays continuously. In this mode, the y axis has a *zoom* function: as the mouse pointer is moved down, the data is played through more slowly.
- 3) *Non-Real-Time Navigation*: in this case the user selects a zone of the sound he/she wants to hear. The user clicks, drags, and releases the mouse. The x position of the initial click sets the start position, and the release sets the end. The user enters the *playback time* in milliseconds and the program calculates the necessary *zoom* factor (alternatively the *zoom* factor can be entered directly). To hear the sound the user presses the *Play* button.

The sonic result of any navigation can be recorded as a wave files and saved to disk for later analysis.

V. CONSIDERATIONS AND OBSERVATIONS SO FAR

A major aim of this project is to evaluate the efficiency of various interactive sonification methods.

Two fundamental criteria will be used in such evaluation:

- 1) which mappings clearly portray the structures encoded in the data streams?
- 2) which mappings are best at portraying correlations between different data streams?

The following observations are based on the above criteria and they were developed by using the prototype Toolkit with the Helicopter and Physiotherapy data. They are not scientifically proven observations, but they are useful as preparatory work prior to experiment design.

A. DataWave sonification (audification)

If the data points of a channel become the samples of a waveform, then the data can be listened to very quickly (e.g. at CD sampling rate, 44100 data points per second are converted to sound). Using audification the user can expect to hear:

- *a particular frequency* if there is periodicity in the data;
- *noise* if the numbers are randomly distributed;
- *increased loudness* for increased activity in the data;
- *nothing* if the data string is constant or moves smoothly and slowly;
- *a click* if there is a discontinuity in the data stream.

It is therefore possible to rapidly deduce some very basic, but fundamental, characteristics of the data stream just by listening to it as a waveform.

When listening with this method to the physiotherapy data the muscles' firings are like bursts of noise with an amplitude and timbral envelope. It is clearly perceivable whether they are shorter or longer, louder or softer than expected. Our therapist collaborators are keen to investigate how the characteristics of the amplitude envelope and the overall timbre relate to specific muscle conditions. Recorded within the data set are also the on/off state of switches which indicate (for example) when a person's weight presses down on a chair (for sitting and standing exercises). These signals simply become clicks under this sonification method.

The helicopter data set is much more complex in terms of the number of channels, and the huge numbers of samples. A typical half hour flight records the output of 400 sensors every 10 milliseconds – yielding a data file of 72 Megabytes. We are studying a sub-set of the gathered data that has 28 channels describing the behaviour of very different flight and user-control parameters (velocities, altitudes, movements of levers etc.). Using the audification method demonstrates quickly which parameters have a similar structure. Clearly noticeable are noise, discontinuities, the presence of a repetitive element, periodicity and amplitudes envelopes. The sounds produced have a complex structure, often a mixture of continuous smooth changes and noisy parts.

The *real-time navigation with loop* method was implemented with the intention of finding a solution to the fact that audification of a smoothly changing curve does not produce a perceivable sound. If a section of that smoothly changing curve is looped however, then a sound is heard because the loop becomes one period of the looping section of the waveform. Clearly there will still be no sound if the data set is completely constant. Jumps and clicks can be located very easily with this method, as the discontinuity can be heard as a suddenly rich signal as the mouse is moved across the data.

B. Mapping data to frequency

Frequency is an important sound variable for mapping. The human ear is highly trained to follow the pitch of sounds. For people with perfect pitch, this dimension of sound could also be quantitatively exact. The perceivable range of pitch is quite large and a lot of detail can be portrayed. People easily follow and remember frequency contours, and even communicate them vocally to other users. Smooth changing curves (which as we have discussed above are not portrayed with audification) can be better displayed with this method. Jumps, noise and periodicities can be also heard as corresponding movements in frequency.

C. Mapping to durations

Mapping to durations can be used to highlight certain patterns in the data. For instance, in a data set that alternates between two values, duration mapping produces a clear repetitive rhythmic pattern, and any change in that pattern would be

clearly noticeable. There are, though, limitations associated with this method. If the data strings are taken from a system which evolves naturally in time (such as our two example task domains!), then duration mapping destroys the analogical relation between the data and time. Moreover, two channels mapped in this way cannot be mixed and compared simultaneously because they occupy different time scales.

D. Listening to two channels

Channels can be compared by playing them simultaneously, and relying on the grouping principles (mentioned in section III). If the channels are mapped to different timbres of sound, the brain can follow the various streams separately. Thus it is possible to tell if one channel is related to the other, or if they have the same general behaviour.

The use of stereo panning enables a greater audio separation of multiple channels. This method has been used by the physiotherapists to verify a hypothesis that a particular set of supposedly symmetrical muscles were not in fact firing simultaneously. The sensors from the left side of a knee were panned left, and from the right side panned right. The resulting sound was rather like a complex stereo ‘ping-pong’ effect, demonstrating that when one muscle showed high activity the other did not, and vice versa. The physiotherapist working on this project found this fact very interesting because she had never seen a good visual display of this behaviour, yet it was immediately obvious in the aural rendition.

We will be investigating another approach to comparing channels. We plan to map the two data streams to variables that build the spectrum of the sound. Various mappings will be designed (for instance the more similar the data on both channels, the more harmonic the synthesised sound).

E. Listening to many channels

When there are a large number of channels to be mapped onto many sound variables, it is difficult to do this so that the data streams can be perceived separately. It should be possible, though, to synthesise a sound whose timbre changes in relation to the data. In this case, a certain timbre will become associated with a certain data configuration.

We plan to experiment with finding a way to represent the entire helicopter data set as a single complex timbre at each moment in time. This will require much work on appropriate mapping and scaling, and combining parameters by identifying redundancy in the signals. One idea is to compare the helicopter state with some ‘idealised’ state by playing the real and idealised timbres in the left and right ears respectively. When the real state is changed (via a user interface) to be the same as the idealised state, the timbres will be identical, and the left and right sound fields will meld to a single mono timbre.

VI. FUTURE WORK

Development and research will continue on the EPSRC-funded project for the next two years. At this stage there are various ideas for further work. Some of these concern specific improvements of the toolkit, while others are more general research developments.

We plan to investigate how the nature of the *interface* and the *interaction style* affects the perception of the data in sound. A number of interface devices will be used in conjunction with the toolkit, to extend the range of interaction possible with the mouse. For instance, we plan to use a fader-box, where each slider controls an aspect of the sonification in real-time, and to supplement this with a rotary dial with tactile feedback. The synchronous use of tactile cues in the data may aid the user’s perception of their location within the data. Within the physiotherapy work, we may render the on/off switches (which act to the therapists as ‘markers’ for key moments in the data) as tactile pulses. The therapists will then be able to more intuitively relate the current sound to the key positions in the recorded movement.

There will be much exchange between Pauletto’s doctoral research work and her work in this project. The doctoral work aims to explore and evaluate data-to-sound mappings on the basis of the perceptual characteristics of the resulting sounds. The goal of this research is to improve our knowledge of data-to-sound mapping techniques in order to improve their application in both the creative (composition, interactive art) and the scientific (auditory display, multimodal interface) domains.

Experiments will be designed to compare various mappings using specially constructed data sets. These allow the researcher to know exactly which structures are present in the data, and thus to verify if they are perceivable. Experiments will also utilise real system data from our two target domains. The perception, by users, of any structure heard in the data will be verified using graphical and/or mathematical methods. Subjects will repeat the experiments on several different occasions so that observations can be made on their ability to learn and interpret each mapping. Various styles of data navigation will also be tested in order to explore the relationship between data mapping and data interaction.

Other ideas for future work include implementing new sonification methods using subtractive synthesis and physical modelling. Vowel sounds can be produced by using subtractive synthesis (e.g. Klatt synthesizer [8]). By changing the values of formants a vowel-like sound changes into another sound with a very different tone ‘colour’ (e.g. from the sound ‘a’ to the sound ‘u’). The fact that humans are innately skilled at perceiving speech-like sounds makes this sonification method particularly worth investigating.

It could be very interesting to map data onto the parameters of physical models of sounding objects, (instead of directly onto sound parameters). Physical models offer both the promise of natural human interaction and many potential inputs (for example, a model of a string can be built with as many masses as required, each of which could potentially receive an input from a data stream).

Finally, a new parameter will be implemented in the Toolkit. This parameter will indicate to the user what the minimum perceivable step in the data is after a sonification method, and

a particular scaling, has been chosen. This parameter will be a kind of 'resolution' indicator for the sonification.

VII. CONCLUSIONS

This paper has described the initial phases of the '*Improved data mining through an interactive sonic approach*' project, and presented the prototype Toolkit with its first informal results. Eleven sonification methods have been implemented. Early experiments have revealed structures in the data that were not at all obvious in the traditional visual-only analysis. Very large files can be listened to rapidly (e.g. one channel of a half-hour helicopter flight, with 100 sensor samples per second, can be reviewed by audification in 4 seconds).

We have begun to assess the relative advantages and disadvantages of the various sonification methods, and noted how certain data structures are highlighted by different sonifications.

The next phase of the project will aim at improving the quality of the interaction by investigating multi-dimensional interfaces and the control of data-to-sound mapping. Experiments will be run to determine how successful each sonification method is at allowing certain structures in the data to be perceived.

REFERENCES:

- [1] Scaletti, C., & Craig, A. B., Using sound to extract meaning from complex data. In Farrel, E.(ed.), *Extracting meaning from complex data: Processing, display, interaction, Proceedings of SPIE* Vol. 1459, 207-219, 1990.
- [2] Hunt A. & Kirk R, Multiparametric control of real-time systems, (chapter in Trends in Gestural Control of Music, eds. Marc Battier, Joseph Rován & M. Wanderley, FSU - Florida, (April 2000).
- [3] Bregman A. S., Auditory Scene Analysis, MIT Press, 1994
- [4] Pure Data: <http://www.pure-data.org>
- [5] Miller Puckette: <http://www.crcra.ucsd.edu/~msp/>
- [6] Zexy externals: <http://www.pure-data.org/downloads/>
- [7] Gripd: <http://crcra.ucsd.edu/~jsarlo/gripd/>
- [8] Klatt, D. H., 1980, Software for a cascade/parallel formant synthesizer, *Journal of Acoustical Society of America*, 67:3, pp. 971-995



Sandra Pauletto obtained a Classical Guitar Diploma in Performance at the Conservatory 'G. Tartini' of Trieste (Italy) in 1995, she obtained a Master of Physics Degree at The University of Manchester (UK) in 1998 and, in 2001, and she obtained a Master of Science Degree in Music Technology at the University of York (UK). In her final project, Sandra developed a VST Plug-in that synthesises vowel sounds. In 2002, Sandra was a Research

Associate at Loughborough University (UK) for the EPSRC funded project 'COSTART'. This project explored collaborations between artists and technologists. In April 2003, she was appointed Research Associate for the EPSRC funded project 'Data mining through interactive sonification' at the University of York. In this University she started her PhD. Sandra's research interests are data to sound mapping strategies, sound perception, interactivity and the relationship between art and technology.



Andy Hunt obtained a BSc in Electronics as the University of York in 1987, followed by a PhD in Music Technology in 2000. He is a lecturer in Music and Media Technology systems, and specialises in Human-computer interaction, interactive sonification systems, multimodal mapping techniques, and new computer-based musical instruments – especially for people with disabilities. He is chair of the Working Group on interactive multimedia systems as part of COST-287 ConGAS (Control of Gestural Audio

Systems). He is principal investigator of the EPSRC funded project 'Data mining through interactive sonification' at the University of York. He has authored over 35 papers in these areas, and five books/chapters.

Sandra Pauletto and Andy Hunt are members of the Media Engineering Research Group, Department of Electronics, University of York, Heslington, York YO10 5DD, UK. Their emails are sp148@ohm.york.ac.uk and adh@ohm.york.ac.uk respectively.