SONIFICATION SANDBOX: A GRAPHICAL TOOLKIT FOR AUDITORY GRAPHS

Bruce N. Walker and Joshua T. Cothran

Sonification Lab, School of Psychology
Georgia Institute of Technology,
654 Cherry Street, N.W.
Atlanta, GA 30332
bruce.walker@psych.gatech.edu, infinite@cc.gatech.edu

ABSTRACT

Motivated by the need for a multi-platform, multi-purpose toolkit for sonifying data, the Sonification Sandbox allows users to map data to multiple auditory parameters and add context using a graphical interface. The Sonification Sandbox is a cross-platform application authored in Java, using the Java Sound API to generate MIDI output.

The software allows users to independently map several data sets to timbre, pitch, volume, and pan, with full control over the default, minimum, maximum, and polarity for each attribute. It is also possible to add context to the sonification using a percussive "click track" or a constant, repeating, or notifying tone at the minimum, maximum, or mean of a given set of data.

Applications for the Sonification Sandbox include auditory display for the blind, science and mathematics education, data exploration, and experimenting with the effectiveness of various sonification techniques and parameters.

1. INTRODUCTION

Sonification is the use of non-speech audio to convey information such as that used in the interpretation of scientific results. That is, scientific data, of any sort, are used to change the parameters of a synthesized tone in order to create an "auditory graph". Walker has pointed out the benefits of auditory graphs in science and education, for both visually impaired and sighted users [1]. However, despite the potential scientific and educational benefits, there are very few widely available software tools for creating sonifications and auditory graphs. The few sonification toolkits that have been developed are either proprietary, dependent on a specific hardware or software platform (e.g., SGI/Irix), are not powerful enough, or are built for the expert sonification designer, and not the schoolteacher, student, or average researcher. There is a clear need for an effective, robust, highly usable, cross-platform, and widely available sonification and auditory graphing toolkit. Such a system would be useful both for creating sonifications and for further studying perception of auditory data recognition.

2. RELATED WORK

The Sonification Sandbox is largely based upon the work of Lodha and his students at the University of California, Santa Cruz, which has evolved from Listen, a relatively simple implementation for the SGI Platform, to MUSE, a more sophisticated iteration incorporating musical parameters, and most recently MUSART, which provides a more robust implementation with an improved graphical interface [2] [3] [4]. Joseph's MUSART produces many of the same effects as the Sonification Sandbox, with the ability to manipulate register, pitch, duration, silence, loudness, thickness, timbre, balance, and beats. MUSART differs from the Sonification Sandbox in its focus on the addition of musical elements to the sonification, while the Sonification Sandbox focuses more on the addition of context. In addition, MUSART (authored in C++, Csound, Fltk, and OpenGL) remains a platform-dependent application, and would require porting the code to several new platforms before widespread deployment would be possible [4].

Other similar work includes Upson's Sound Grid which, like the Sonification Sandbox, uses Java's Java Sound package and Swing to create a multi-platform graphical application for experimenting with sonification [5]. Sound Grid enables users to plot points of varying magnitude (mapped to note duration) on a Cartesian grid, then play back the sonification. Timbre and tempo can be manipulated in real-time while the sonification is playing, and the visual graph highlights which point is being sonified for each note during the sonification playback. Although the Sound Grid is a useful exploratory tool, it is primarily designed for educational applications, and has limited utility for the kind of precise technical parameterization necessary for research applications. As with MUSART, Sound Grid does little to provide auditory context, which Smith and Walker have shown to be valuable in auditory graphs [6].

Analysis of other sonification systems has yielded many of the same conclusions as Wilson—most have produced encouraging results in support of the potential usefulness of sonification, but as applications, they fail to address many issues critical to its widespread adoption as a research technique [2]. These shortcomings include the need for exotic hardware, support for very few platforms, domain-specific design, and limited numbers of data sets and/or parameters.

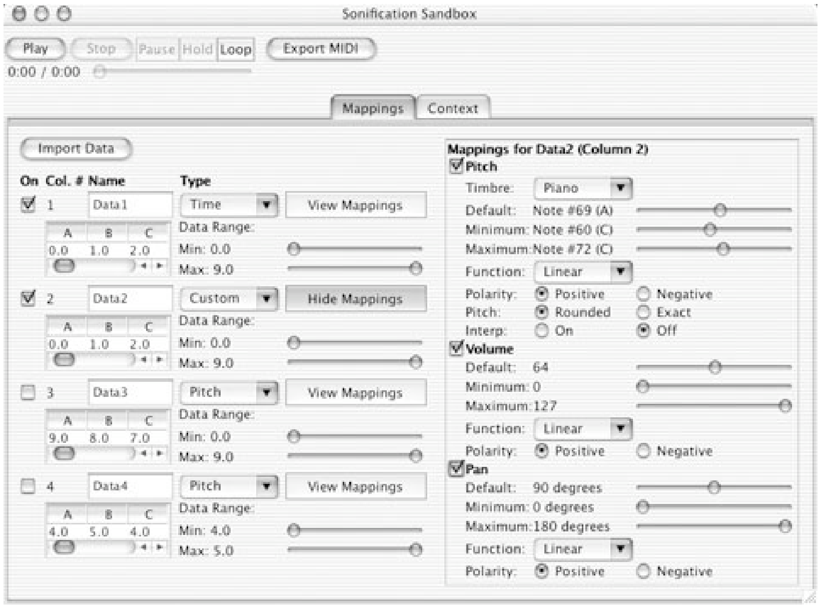

0:00 / 0:00

On	Col. #	Name	Type
☑	1	Data1	Time
☑	2	Data2	Custom
□	3	Data3	Pitch
□	4	Data4	Pitch

Data1: A 0.0, B 1.0, C 2.0; Data Range: Min: 0.0, Max: 9.0

Data2: A 0.0, B 1.0, C 2.0; Data Range: Min: 0.0, Max: 9.0

Data3: A 9.0, B 8.0, C 7.0; Data Range: Min: 0.0, Max: 9.0

Data4: A 4.0, B 5.0, C 4.0; Data Range: Min: 4.0, Max: 5.0

Mappings for Data2 (Column 2)

☑ Pitch
Timbre: Piano
Default: Note #69 (A)
Minimum: Note #60 (C)
Maximum: Note #72 (C)
Function: Linear
Polarity: ● Positive ○ Negative
Pitch: ● Rounded ○ Exact
Interp: ○ On ● Off

☑ Volume
Default: 64
Minimum: 0
Maximum: 127
Function: Linear
Polarity: ● Positive ○ Negative

☑ Pan
Default: 90 degrees
Minimum: 0 degrees
Maximum: 180 degrees
Function: Linear
Polarity: ● Positive ○ Negative

Figure 1: The Mappings Panel for the Sonification Sandbox

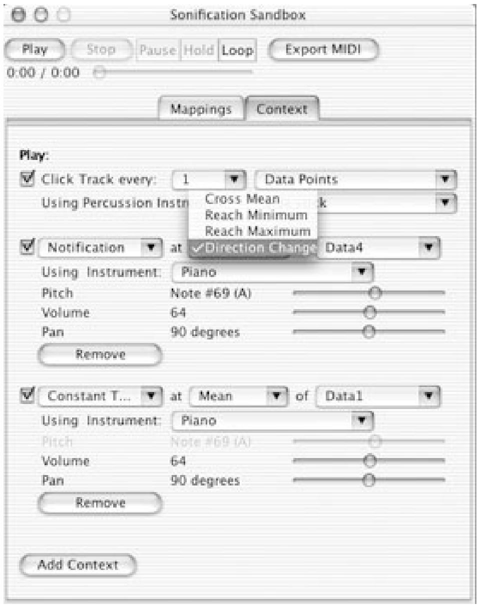

0:00 / 0:00

Play:
☑ Click Track every: 1 Data Points
Using Percussion Instr… …k

☑ Notification at Direction Change Data4
Using Instrument: Piano
Pitch Note #69 (A)
Volume 64
Pan 90 degrees

☑ Constant T… at Mean of Data1
Using Instrument: Piano
Pitch Note #69 (A)
Volume 64
Pan 90 degrees

Figure 2: The Context Panel for the Sonification Sandbox

3. DESIGN

The Sonification Sandbox is, as the name suggests, intended to provide the user with a simple, user-friendly, moldable, responsive environment for creating sonifications. As with the sandboxes on a playground, one size fits all—the Java Virtual Machine allows users on all popular platforms to use the software with little or no setup, and open-source code enables developers to shape the application to meet their individual needs and preferences. Once a sonification is satisfactory, it's easy to "show and tell"—the application supports MIDI export for easy sharing of auditory graphs across multiple platforms. On the other hand, designing for a range of users from novice to expert required a significant reduction in complexity compared to some sonification toolkits; as such, the application is limited by the capabilities of MIDI, and has a comparatively smaller range of potential combinations of settings for sonifications.

3.1. Current Status

At present, a stable release of the Sonification Sandbox is available for download online (with the source code) at http://sonify.psych.gatech.edu/research/. The program is in a runnable JAR file, and is executable on MIDI-enabled Windows, Macintosh (OS X), Linux, and Solaris SPARC or Intel platforms with Java 2, version 1.3.1 or higher.

3.1.1. Data Input

Data is provided to the Sandbox by importing a comma separated value (CSV) file from the user's computer. The file can be edited by hand, or more likely, in a spreadsheet application such as Microsoft Excel. A valid input file must consist of all-numerical input in the form of an MxN matrix; users may specify whether the data are to be sorted by rows or columns. Users can update data while the Sonification Sandbox is open by saving the data in an external editor and re-importing it into the application.

3.1.2. Mappings

Each data set can be mapped to pitch, timbre, volume, and pan independently (see Figure 1). If data are not mapped, the user can select a default value using a slider, which ensures that the data falls within the correct range. If the data are selected for mapping, the minimum and maximum slider values are used, and a transfer function performs a linear transform from the data's initial range to the range specified in the Mappings Panel. The toolkit enables the user to control for outliers by allowing the maximum and minimum of the data range to be set; data outside of the range are automatically compressed to the revised maxima or minima during the transform process. In addition, it is possible to reverse the polarity of a mapping to match the user's conceptual model of magnitude for a given combination of data type and tonal attribute.

Only one data dimension may be mapped to time at once for a given sonification. By default, the first valid data dimension is mapped to time. The time dimension is the only setting for which a specific trend restriction is imposed; the data set must increase monotonically. If the data points are not uniformly spaced, the length that each data point is played corresponds to the relative spacing of the data points in the time dimension.

To prevent user error, all widgets (check boxes, sliders, etc.) have checks and balances to ensure that the settings for the auditory graph are intuitively and algorithmically within the appropriate ranges.

3.1.3. Context

The Sonification Sandbox provides for the addition of context to auditory graphs in the form of a click track, constant or repeating tones, or notifications, which are analogous to the tick-marks, axes, and labels found in visual graphs (see Figure 2). The click track, which uses a percussive instrument, can be set to play at intervals of data points or seconds. Constant or repeating tones can be played at the mean, minimum, or maximum of a data set. For constant or repeating tones, pitch cannot be set, since it must be relative to the pitch of the tones being sonified in order to be meaningful, however, volume and pan can be adjusted to improve the sound design of the auditory graph. Finally, a notification can be played when a data set crosses the mean or reaches the minimum, maximum, or an inflection point (slope change). For notifications, pitch, volume, and pan can be set. The Sonification Sandbox permits the user to add as many sets of context as they like and turn contexts on and off, and will play as many contexts as the MIDI specification allows, once the primary sonification has been rendered (General MIDI places a limit of 16 channels of output at once).

3.1.4. Playback and Output

Playback for the Sonification Sandbox is provided via a standard Play/Stop/Pause set of buttons, with an interactive progress bar for viewing and adjusting the playback position while the sonification is playing. A toggle Loop button allows the user to loop the sonification, and a Hold button causes the sonification to loop on the current second of the auditory graph. An Export MIDI button allows the user to save the current sonification as a MIDI file, a lightweight format that can be played on most multimedia-capable computer platforms using an application such as QuickTime.

3.2. Work in Progress

The Sonification Sandbox represents a strong step in the right direction for sonification toolkits, and it continues to evolve and improve as new features are added. Several efforts are currently underway to bring the software even closer to the goals that have been established for it. First, we will be adding settings for different kinds of data (e.g. money, weight, velocity, etc.) that will automatically generate mappings to reflect the results found by Walker in his work on magnitude estimation of conceptual data [7]. For instance, it has been found that, although most mappings have a positive polarity (e.g., increasing velocity to increasing pitch), some polarities are negative (e.g., increasing mass to decreasing pitch). Further, these mappings vary for sighted and non-sighted individuals (e.g., money).

Second, features will be added to improve the usability of the software for members of the non-sighted community. In its current implementation, a math or science professor could use the software to sonify graphs for non-sighted students, but the visual dominance of the application interface lends itself poorly to use by the blind. To mitigate this, alternative input options such as keyboard shortcuts and/or speech input will be added, and the interface will be augmented to provide more robust support for screen readers such as JAWS and Window-Eyes.

Third, a Java spreadsheet application will be integrated into the interface, with improved support for file types, support for equations, enhanced data selection, a sonification wizard, and other improvements. This will allow users to seamlessly transition between data editing and sonification. It will also be possible to save the sonification settings for a set of data in a project file.

Finally, the application will integrate a visual graph alongside the auditory graph output. Initially the graph will be static, but later iterations will include animation (e.g., highlighting the data point currently being played) and interaction (e.g., a data point will be played when the user clicks on it).

4. CONCLUSIONS

While there have been several excellent sonification toolkits to date, the increasingly widespread use of sonification will continue to rely on broadly usable, platform-independent, and user-friendly tools. Further, these tools need to continue to implement the results from experimental studies in the sonification and auditory display literature. The Sonification Sandbox project represents a step in this evolutionary process, intended to provide a useful way to create auditory graphs and also provide a platform for further research and development.

5. REFERENCES

[1] Walker, B.N. and D.M. Lane. “Psychophysical scaling of sonification mappings: A comparison of visually impaired and sighted listeners,” in *Proceedings of the International Conference on Auditory Display*. 2001. Espoo, Finland.

[2] Wilson, C.M. and Lodha, S.K., “Listen: A data sonification toolkit,” in *Proceedings of the International Conference on Auditory Display*. 1996. Palo Alto, USA.

[3] Lodha, S.K., Beahan, J., Heppe, T, Joseph, A.J., and Zne-Ulman, B., “MUSE: A musical data sonification toolkit,” in *Proceedings of the International Conference on Auditory Display*. 1997. Palo Alto, USA..

[4] Joseph, A.J. and Lodha, S.K. "MUSART: Musical Audio Transfer Function Real-Time Toolkit," in *Proceedings of the International Conference on Auditory Display*. 2002. Kyoto, Japan.

[5] Upson, Robert. "Educational sonification exercises: Pathways for mathematics and musical achievement," in *Proceedings of the International Conference on Auditory Display*. 2002. Kyoto, Japan.

[6] Smith, D. R. and Walker, B.N. “Tick-Marks, axes, and labels: The effect of adding context to auditory graphs,” in *Proceedings of the International Conference on Auditory Display*. 2002. Kyoto, Japan.

[7] Walker, B.N. “Magnitude estimation of conceptual data dimensions for use in sonification”, *Journal of Experiment Psychology: Applied*, 8, 2002.