DESIGNING SYSTEMS FOR THE CREATION AND EVALUATION OF DYNAMIC PERIPHERAL SOUNDSCAPES: A USABILITY STUDY

Bradley S. Mauney and Bruce N. Walker
Sonification Lab, School of Psychology
Georgia Institute of Psychology,
Atlanta, GA 30332-0170
bruce.walker@psych.gatech.edu

Sonifications have proven useful as stand-alone displays and in bimodal audio-visual displays. Audio is powerful in a peripheral display due to its unobtrusive and omni-directional nature. This study examined the creation and usability of peripheral auditory displays or 'soundscapes' composed of ecological sounds. A system was created for dynamically generating soundscapes from a data source according to a threshold-based model. In this model, a percentage change in the data is mapped to an ecological sound that is played whenever that threshold has been met or exceeded. The system allows for the creation of soundscapes of arbitrary complexity, providing mechanisms for fading, relative gain, and random timing. With the system, a reference soundscape was developed for sonifying the stock market. This display was then presented to stock trader test participants as part of a usability evaluation.

INTRODUCTION

Modern science has seen the proliferation of vast quantities of data to be analyzed, such as that from the Human Genome Project. The integrative nature of audio and its support for pattern recognition can make auditory displays effective in the interpretation of these large, multi-dimensional data sets (Kramer et. al., 1999).

Data have typically been displayed by visual means; however, research has shown that in visually demanding situations auditory graphs can be more effective than purely visual or mixed visual and audio displays (Barras & Kramer, 1999). Sonification focuses on using non-speech audio to convey data relations into perceived relations in sound. Highly successful, albeit simplistic, sonifications include the Geiger counter and the pulse-oximeter. In general, sonifications are successful when either the user cannot see a visual display due to physical limitations or cannot afford to look because they are performing a competing visual task. Research has also shown that large, multi-dimensional datasets can be sonified without overloading the user, as would be the case with a visual display (Kramer et. al., 1999), and that the addition of auditory tick-marks, axes, and labels further enhance the perceivability of these displays (Smith & Walker, 2002). This study focuses on the addition of this contextual information in the form of gradients overlaid within a peripheral auditory display or soundscape.

Peripheral Auditory Displays

An advantage of peripheral auditory displays can be found in the perceptual modalities of multiple resource theory (Wickens, 1992). According to this framework, auditory-visual cross-modal displays have shown an advantage over intramodal displays (visual-visual or auditory-auditory) when there is competition in time-sharing between two or more intramodal channels. An improvement can be seen when some visual channels are displayed in the auditory modality, especially when there is considerable visual scanning required. Thus, not only can both auditory and visual channels be perceived in parallel, but in many cases a bimodal display is better than a purely visual or purely auditory display. Additionally, auditory perception is especially sensitive to temporal events and frequency changes, which has proven to be an advantage over visual displays when the data are rapidly changing (Kramer et. al., 1999). This argument supports the importance of auditory displays in peripheral awareness.

A peripheral auditory display must be designed to be unobtrusive, peripheral, and continuously accessible over the ambient noise of an indoor (e.g., office) environment. To be peripheral, such a display must not fall into the 'alarm' range of hearing identified as the general range of human voice or 200 – 2000 Hz (Mynatt, 1998). These considerations echo the literature on the creation of auditory warnings (e.g., Patterson, 1982). In contrast to the auditory component of most computer displays, by using natural, ecological sounds such as birdcalls, thunder, or running water the alarm range can be avoided and the display can be pushed further into the periphery. Additionally, higher (CD) quality sounds support the aesthetic, realistic, and distinguishable quality of the soundscape.

In building peripheral displays, the technology must support high-fidelity audio in a flexible, extensible, and cross-platform solution. JASS is a Java-based audio framework that encapsulates the complexities of audio rendering, while remaining robust and extensible enough to support arbitrarily complex manipulation and filtering (van den Doel & Pai, 2001). JASS supports rapid prototyping of auditory displays by providing a library of predefined unit generators that can be used as is or extended with user-defined functionality. JASS and an application built upon it can be run on any machine that is Java enabled and has basic audio hardware.

SYSTEM DESIGN

Sound Design

In this study, a system was designed on top of JASS for the real-time rendering of ecological sounds into a dynamic, peripheral soundscape. The objective of this approach was to enable livable soundscapes that are easily distinguished from the background, but can be allowed to fade from attention and not be tiring or intrusive.

In this model, the user defines a set of thresholds relative to some baseline value and associates a sound clip (e.g., cricket chirping) with each. Each threshold represents a *percentage change* relative to the baseline; when the threshold is reached, the system activates the associated sound, to be played randomly within a certain interval or looped continuously. While the data value exceeds the threshold, the sound clip continues to play even as progressive thresholds are activated, resulting in a cumulative effect the further the value moves from the baseline.

By this design, the system is generic enough to be applied to any continuous data source where the user is concerned with percentage change. Each threshold and its corresponding frequency, relative gain, attack, and decay is completely user defined. By providing the random factor to clip timing, the soundscape becomes more natural, less mechanical, and not repetitive.

Peripheral Soundscape System

To implement such a design, we chose a multi-threaded, unit generator based approach, leveraging JASS libraries and written in Java 1.4. The system loads a user defined properties file specifying all thresholds and associated sound clips. Each threshold mapping is encapsulated and run in a child thread responsible for activation, frequency, and gain. A monitor module parses the data stream or reads from a flat file to simulate real-time data. At set intervals, the sound-mapping module updates the current value and alerts each child thread, which performs the appropriate action on its respective sound clip. The audio buffer from each child thread is output to a player module, which renders sound to the audio hardware.

Stock Market Soundscape Reference Implementation

As a reference implementation, we created a stock market soundscape for monitoring the price of a certain stock or market index. The dynamics of the stock market application provide for a good test of the soundscape concept, and would generalize to other situations like process control and monitoring, or even status and situation awareness on a ship or space station.

In the primarily visual world of stock trading, traders rely heavily on visually attending to 'tickers' or 'streamers' for the latest price movements. However, with so much information being displayed visually through these tools, there is a strong chance of visual overload or at least difficulty in finding a certain type of information.

Several efforts have been made to sonify stock market data. These range from simply audifying the data (Frysinger, 1990) to much more complex multimodal displays created by Nesbitt and Barrass (2002). In most cases, sound attributes such as pitch and loudness are driven by stock price and trading volume. It is important to note that the majority of these approaches have been aimed at the auditory display being either the sole information display, or a significant partner in a multimodal display (e.g., Nesbitt & Barrass, 2002). The auditory display often requires considerable attention in order to function as intended.

The intra-day trader, on the other hand, is already immersed in visual displays, so we sought to provide a *supplemental and peripheral* source of information, rather than an alternative display method. In the soundscape implemented here, the baseline data came from the *simple moving average* of the issue—an indicator favored by technical analysts—and thresholds were defined above and below this baseline. For each threshold a sound clip was added to the soundscape. For example, at the "+1.0%" threshold, crickets begin to chirp at the rate of two samples per minute when the moving average departed from its initial baseline by more than one percent. As the data value moved farther away from the baseline, the sounds built upon one another to create a more complex and layered soundscape. The details of a complete soundscape design can be found in Mauney and Walker (2004).

EVALUATION METHOD

The resulting soundscape system was evaluated both in terms of sound design and system design. To evaluate the sound design with real expert listeners (our focus here), the system was presented to four stock traders of a private equity fund that manages tens of millions of dollars of investor funds. These traders place several large trades per week averaging $20 million in volume per month. Even though they only place a few trades per week, they nearly constantly monitor the stock, currency, commodity, and futures markets to identify trading opportunities. Each trader typically sits at a desk and uses a computer based trading platform with extensive visual graphing tools but nothing particularly auditory related. The traders interviewed each had their own private office, but this is an exception to the industry norm–generally a trader works from a 'bullpen' with several other traders.

The evaluation took the form of a semi-structured interview in which the participants were given progressive training. The soundscape was first presented with the participant only having been told that the sounds represented the price movement of the S&P500. The participant was then asked to describe what he had heard. Next, the concept of threshold mapping was explained and a description was provided for each sound clip. The soundscape was then played a second time, and the participant was asked to describe what he had heard. The participant was then shown the underlying price quote data and given a brief description of the system's threshold activation and events for which they should listen. After a third listening, the participant was asked to describe the trends in the data from what he had heard. After the listening task, each participant was then interviewed as to his perception of the system and its usefulness in a stock-trading environment.

RESULTS

On the first listening the users were generally confused as to the nature of the soundscape. They could distinguish the sound clips and recognize them as events of some type, but none perceived the threshold mapping that was taking place. After the second listening, each participant reported he could roughly follow what was being displayed and could identify when the price was moving up and down. Upon the third listening, the participants reported they could better distinguish the trends while knowing specifically the events to listen for and could identify when a threshold had been reached and when a local maximum or minimum value had been passed.

There was some feedback from some of the listeners that that the system, although novel and interesting, might be too ‘busy’ to be used in an office or trading floor environment citing that these environments require considerable concentration and are naturally noisy. Of course, this study involved a limited exposure duration (approximately 30 minutes), so it does not factor in the habituation and familiarity with the soundscape that would occur over days and weeks. It was also noted that most traders are primarily dependent on visual graphs, are quite accustomed to using exclusively visual tools, and might be resistant to the adoption to such a seemingly foreign auditory tool. Certainly, users are often resistant to change at first, even if an enhanced technology can improve their performance. Overcoming technological inertia is a common challenge in implementing novel approaches, and auditory displays are no different in that sense. Regardless, taken together these points of feedback suggest the need for further research, including a longitudinal study involving training and an adaptation period.

DISCUSSION

In general, participants felt that the threshold-based soundscape was a fairly intuitive way to perceive price fluctuations, especially with regard to percentage deviation from moving averages. They also felt that the sound for each threshold seemed to be unique enough to distinguish it from the others when multiple samples are overlapped. They agreed that the overall distinguishable characteristic of nature sounds versus the sounds from an office environment aided in the perception and interpretation of the data and was ‘relaxing’ to listen to. This provides support for the use of sounds that are natural, but not necessarily part of the listener’s everyday acoustic ecology. They also agreed that it would be more useful than visual graphs when they are away from their desks and moving around their offices, which supports our user-centered decisions to make the display fit into the existing work flow. They felt that the peripheral nature of the system would keep it from ‘getting in the way’ once they became more acclimated to it being there. They also felt that significant training would be needed in order to use the system with the accuracy and equivalency of their standard visual tool. However, it should be noted that the goal of this project was not the identification of a specific data value (i.e., point estimation), as might be the case with a visual system, but rather the notification of threshold crossings. For this purpose the listeners felt it would be very effective and useful.

The user evaluation raised many new questions regarding performance, training, and visual versus auditory attention. As Smith and Walker (2002) have pointed out, training is an important element in performance with an auditory display. Further research on the system would involve quantifying with test participants many of the very qualitative results from the user interviews such as performance and equivalence versus or in combination with visual graphs.

With regard to the software application, it proved robust and flexible, and should serve as an effective platform to study all manner of peripheral auditory displays. Planned enhancements include an extended user interface, an internet module with live data feeds, and a library of pre-assembled soundscapes ("themes") such as rain forest, city street, or sea shore that would help other researchers quickly build up a display that is coherent and balances aesthetics and performance from the very start.

REFERENCES

Barrass, S., & Kramer, G. (1999). Using sonification. *Multimedia Systems, 7*, 23-31.

Frysinger, S. P. (1990). Applied research in auditory data representation. In E. Farrell (Ed.) *Extracting meaning from complex data: processing, display, interaction. Proceedings of the SPIE, Vol 1259*, pp. 130-139.

Kramer, G., Walker, B.N., Bonebright, T., Cook, P., Flowers, J., Miner, N., et al. (1999). The sonification report: Status of the field and research agenda. Report prepared for the National Science Foundation by members of the International Community for Auditory Display. Santa Fe, NM: International Community for Auditory Display.

Mauney, B.S., & Walker, B.N. (2004). Creating Functional and Livable Soundscapes for Peripheral Monitoring of Dynamic Data. *Proceedings of the International Conference on Auditory Display ICAD2004*, Sydney, Australia. July 6-10.

Mynatt, E.D., Back, M., Want, R., Baer, M., & Ellis, J.B. (1998). Designing Audio Aura. *Proceedings of the CHI Conference, CHI 1998,* (pp. 1-8).

Nesbitt, K. V., & Barrass, S. (2002). Evaluation of a multimodal sonification and visualisation of depth of market stock data. *Proceedings of the 2002 International Conference on Auditory Diplay, Kyoto, Japan, July 2-5 2002.*

Patterson, R. D. (1982). *Guidelines for auditory warning systems on Civil Aircraft*. London: Civil Aviation Authority.

Smith, D.R., & Walker, B.N. (2002). Tick-marks, axes, and labels: the effects of adding context to auditory graphs. *Proceedings of the 2002 International Conference on Auditory Diplay,* Kyoto, Japan, July 2-5 2002.

van den Doel, K., & Pai D. (2001). JASS: An Audio Synthesis System for Programmers. *Proceedings of the 2001 International Conference on Auditory Display*, Espoo, Finland, July 28-August 1, 2001.

Wickens, C. D. (1992). *Engineering psychology and human performance (2nd ed.)*. New York: Harper Collins.