Auditory Graphs of Real-Time Data

Edward Childs

Accentus LLC,
115 B Etna Road,
Building 3
Lebanon, NH 03766,
USA.
edward.childs@accentus.com

ABSTRACT

The advantages of auditory display for monitoring real-time data are discussed. Parallels between the structures of real-time data and music are emphasized as potentially fruitful areas of research. The Accentus LLC design philosophy is described, followed by several examples of auditory graphs. Areas of future research are recommended.

1. INTRODUCTION

Monitoring real-time data is a very promising practical application for auditory display. Most features of a fixed graph on a piece of paper, especially its structure, may be visually absorbed almost instantaneously. However, attending to a graph on a computer screen which is changing with time presents more challenges. When the task involves the monitoring of several time-dependent variables, there are problems both with how to present the data graphically, and how to attend to that presentation consistently.

1.1. Auditory Display Advantages for Real-Time Data

By its very definition, real-time data unfolds over time and exhibits certain kinds of "movement behaviors." Professionals who need to monitor this sort of data must coordinate these behaviors with time-pressured analysis and decision-making tasks. More often than not, these tasks require specific physical movements (such as the execution of a command at the computer keyboard, or placing a phone call).

There is substantial evidence, as documented by psychologists working in the field of music cognition and perception, that humans are innately more sensitive and responsive to time domain behavior when it is represented auditorially. Temporal discrimination is superior, and when both visual and auditory stimuli are present, the latter has more influence on the observer's judgment. Motor activity (such as tapping a beat) is controlled by auditory, rather than visual input [1].

Several phenomena which point to highly sophisticated interactions between music, performers and listeners have been noted. A very partial list:

1. The perception of motion (harmonic, melodic, rhythmic) on different cycles and time-scales [2].
2. Sensitivity to expressive performance timing (e.g. tempo fluctuations) [3].
3. Ability to recognize rhythm, meter and tempo despite performance variations [4].
4. It is easier to judge the duration of melodies when they are musically "coherent" [5].
5. "Brief excerpts of classical music generate behavioral and electrophysiological priming effects that are indistinguishable from those evoked by sentences." [6] [7].

Designing and testing sonifications of real-time data which draw on the "palette" of musical attributes in the above list is a compelling area for research and practical application. When the data monitoring task also requires the execution of commands using multiple, rapid keystrokes (such as the case with financial traders using electronic platforms) the opportunity is even more pronounced. Sonifications which are effectively designed could eliminate the need to look at the screen for the data which will lead to the keystrokes, i.e. the user would execute keystrokes in direct response to the musical sounds, bypassing the visual mode altogether. The process could be described as "playing a musical instrument in reverse."

1.2. Financial Data

For three years, Accentus LLC has been working with financial traders to design and build auditory displays that provide value [8]. Some traders focus their attention on a limited number of data streams but trade intensively (several hundred transactions per day), in order to exploit arbitrage opportunities which may last less than a second. Others focus on a larger number of data streams but trade less frequently, in response to specific "target" market events on which their strategy is based.

Regardless of their style, traders use multiple visual displays (six is common) and predominantly view their data as numbers, often in extremely small fonts. Graphs are sometimes displayed; however their use is limited due to the significant amount of screen "real estate" they occupy. Alternative graphical displays, (e.g. Heatmaps) are occasionally used in an attempt to display information in a more compact, graphical format [9].

The predominant mode of data monitoring is to scan sets of numbers continually on multiple screens for the occurrence of important events on which they will trade. While traders develop formidable visual skills in order to stay on top of multiple rapidly changing data streams, nearly all acknowledge that they have missed critical opportunities to realize gains or avoid losses because they "weren't looking."

The notion that auditory cues could play a significant role in financial trading is not new. The problem is that the modern electronic trading environment lacks many of the auditory cues which were available on a traditional trading floor. The sound of particular voices or the general level of hubbub on the floor

would provide substantive market information. Traders would trade on the basis of the sound of the ticker tape without looking at the actual print.

2. DESIGN PHILOSOPHY

The Accentus sonification design methodology has been to:

1. Leverage the experimental results of music cognition research in sonification design.
2. Simplify the design to the greatest extent possible, so that it can be learned as quickly as possible.
3. Use the highest audio production standards.
4. Test the design in practical, working environments.
5. Quantify the efficacy of the sonifications using rigorous psychological experiments and/or measurements of productivity and effectiveness.

This paper will focus on the first three items (design). The last two (testing) have been documented in [8] and will be the subject of future papers.

2.1. Guidance from Music Cognition Research

There is increasing experimental evidence that music, as a language, possesses phonological, syntactic and semantic components which may be used to communicate information (beyond mood or feelings).

Another motivation for using musically-based sonifications (as opposed to other kinds of sounds, e.g. crickets) is the general notion that the brain processes speech, environmental and musical sound in a somewhat parallel fashion. In a city office setting where electronic trading typically occurs, the speech and environmental "channels" are already in use, however the music "channel" is idle.

It has been shown scientifically that the general population absorbs and responds to certain syntactic aspects of musical language by mere passive exposure [10]; those aspects are therefore available to be "loaded" with data in the design of sonifications. For example, the P600 event-related brain potential has been shown to indicate both syntactic irregularities in language and incongruous (or distant) chords in an otherwise predictable harmonic progression [11]. It follows that if a sonification of, say, the risk profile of a portfolio needs to convey a "distance" from the desired profile, that a "distant" harmonic representation could be effective.

The general principle could be extended to guide the choice of timbres (i.e. instruments), scales, etc. that are appropriate to the culture in which the sonification will be used.

2.2. Simplicity

There are many reasons to opt for the simplest, sparsest design for practical use. First, in order that the sonifications may be learned as quickly as possible and require little or no explanation or training. Second, to make efficient use of available auditory bandwidth. In a practical working environment, it is essential that we not introduce complex or repetitive sonifications that may interfere with other activities (such as speech communications) or be distracting or annoying. Furthermore, if the same sonification is being applied to multiple data streams, a sparse design will make it easier for the user to distinguish one stream from the other. Third, simplicity and clarity are aesthetic goals in many traditions of musical composition. For example, in response to a radio interview question [12] "If you had one piece of advice to young composers, what would it be?" the Finnish composer Jean Sibelius responded "Do not write too many notes, for each note must have a life of its own." Finally, to the extent that there are parallels between the design of auditory and visual displays, simplicity and sparseness is an established design principle for graphical presentation [13].

Following on in this tradition, Accentus sonifications follow the general principle that every sonic attribute (including the mere fact that a sound occurs at all) should arise from a specific data event. The dimensionality of the data being sonified should match the number of musical attributes that are manipulated.

2.3. Audio Production

Practical sonification design workers should consider themselves in the same league as Foley artists and music producers in the movie industry, where considerable effort and resources are expended to get just the right sound effect (e.g. the light saber sound in Star Wars), or the appropriate music to enhance the scene. Audio systems in theaters produce high quality sound, as do commonly available home entertainment systems used for watching DVD's.

Modern desktop and laptop computers (even hand-held devices) are routinely configured as multi-media systems. Even a modest investment in a 2.1 computer speaker system will provide a reasonable quality of sound, and enable the listener to hear timbral differences. If orchestral instruments are to be used in sonifications, high quality orchestral samples are notably superior to MIDI instruments.

As an anecdotal example, one Accentus client had been using an early prototype in which a staccato bassoon sample (three bassoons, in unison, with substantial room ambience, and a lot of character, extracted from an excerpt of *The Sorcerer's Apprentice*) was used to represent the Dow Jones Industrial Average. He was disorientated when he switched to a web-based applet version using a MIDI bassoon [14]. (A MIDI implementation had been used to minimize the download footprint of the applet).

3. EXAMPLES

3.1. Single Variable Time Series

The simplest type of graph in common use by financial traders is a trace of the "price" of some security as a function of time, over the course of the trading day. In Fig. 1, the price at the extreme left would be the "open." The trader typically uses the Accentus auditory display to be informed of "significant movement" relative to the "open." In this example, the "significant movement" is ten points, which is graphically represented in Fig. 1 by the vertical grid lines. The sonification, then, will occur whenever the trace crosses a grid line (a crossing event is shown by the red circles). The Accentus Three Note Movement Tracker uses the musical interval to express the current value (third note) relative to the previous grid line crossed (second note) and also relative to the open (first note).

Figure 1. *Three Note Movement Tracker*

The musical intervals are based on the half-step (i.e. chromatic). The open is considered to be a "reference pitch" and is always the same (for the given instrument). In the Fig. 1 example, the index is the Dow Jones Industrial Average, the significant movement is 10 points and the opening value is 10,080. Thus, the sonification event (red circles) occurs at 10,160 (8 significant movements above the open), which would correspond to the interval of a minor sixth. The previous sonification event occurred at one grid line higher than the current (10,170) or 9 significant movements (a major sixth) above the open. If the reference pitch was chosen to be G3 (G below middle C), the pitches generated by the event would be G3, E4, Eflat4. The information provided by the three notes is: "Dow Jones Average (bassoon), up 80 points on the day, down 10 points from the previous sonification." The auditory graph thus emphasizes the most recent events, rather than the overall shape of the day's movement, which is best absorbed visually, and which a trader will generally be familiar with anyway over the course of a trading day.

The Three Note Movement Tracker may be extended, through the use of a fourth note, to indicate the approach to a target. In the above example, if the trader wished to execute an order when the Dow reached 10,200 (which is 12 significant movements above the open), a fourth note, one octave above the reference pitch, would sound, when the current price is within 6 significant movements of the target (i.e. starting at 10,140). The fourth note becomes more insistent at the target is approached (of longer duration, and louder). At the point of crossing the target, the fourth note sustains.

The Three Note and Three Note with Target schemes work well when the significant movement is set high enough so that the calculated musical intervals do not exceed an octave. When this is not the case, a simplified Two Note scheme may be used which reports only an increase or decrease from the previous sonification event, but no longer relates the current value to a reference such as the open price.
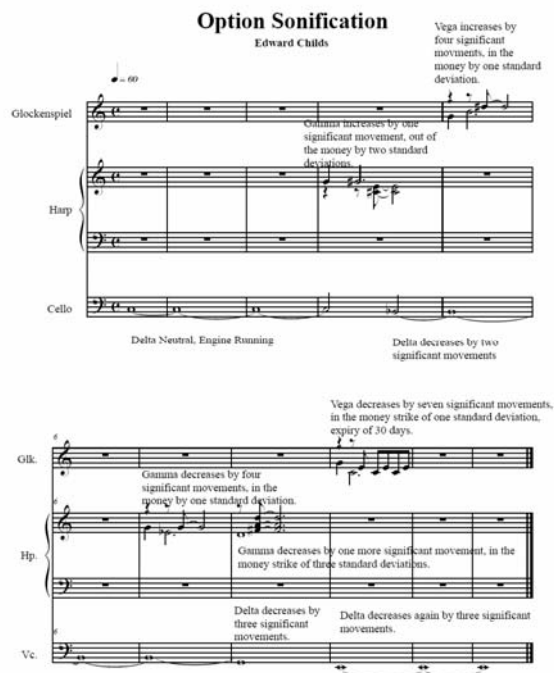
3.2. Multiple Variable Time Series

Figure 2. *Option Sonification*

Options, by which a trader purchases the right to buy or sell an "underlying" security at a certain price at some point in the future, are frequently used in "hedging" strategies, which are designed to protect the investor irrespective of the direction of market movements. In monitoring the overall risk of a portfolio, the "Greeks", delta, gamma and vega are frequently calculated. If delta is zero, then a portfolio is "perfectly" hedged. Gamma is the first derivative of delta, and vega is a measure of market volatility. Options may be priced at the same price as the underlying security (at the money), higher (in the money) or lower (out of the money). The option price relative to the underlying is known as the "strike" and is usually measured in standard deviations. The date at which the option may be exercised (in units of days) is known as the "expiry." Typical values are 30 days (short) or 300 days (long).

In designing a musical sonification to convey all of this information in real time, the following considerations were followed:

1. Delta, being the most important of the Greeks, was sonified in a continuous fashion. A low, soothing String Ensemble sound was chosen as a timbre which could be tolerated by a user over long periods of time. It was cast in the low register to suggest a bass line or an underlying tonality.

2. As with the single variable time series, discrete musical intervals were chosen to represent an increase or decrease in the value of the Greeks (with a harp representing gamma, and a glockenspiel representing vega).

3. A major/minor chordal "decoration" was provided for gamma and vega only, to indicate in the

money (major) or out of the money (minor) strike prices.

4. An echo was used to indicate distance into the future, the longer the echo, the further the distance. Echoes generated by shouting at a distant hill will take longer to return the further away the hill is.

In Fig. 2, various musical events are illustrated in response to the textual explanations of the changes in the data.

4. RESEARCH DIRECTIONS

Auditory display was recently named in the Harvard Business Review's list of break-through innovations for 2005 [15]. The theme of "too much data, too much information" is prevalent in any field requiring the monitoring of real-time data, or searching large databases (e.g. data mining). The design of auditory displays for situations in which time is a factor seems the most compelling research area. In the author's view, research into the parallels between structures in time domain data and music may be the most promising area of exploration. For one thing, music possesses structure on many different time scales, all the way from note by note to an entire movement, work (e.g. opera) or cycle of works (e.g. Wagner's Ring).

Since there is so much research into how humans perceive these structures from a purely musical standpoint, it would seem that we already have a "head start" in this line of inquiry. The task, then is to design sonifications which evoke similar perceptions (which must of course be faithful to the data relationships).

The Accentus LLC research and development program has thus far established the effectiveness of the simple two note scheme [8]. Further research into the Three Note and Three Note with Target schemes must take into account the effect of tonality on the perception of the intervals. For example, since the reference note always sounds in all sonification events, there is a strong implied tonality. The scale of upward and downward movement (the second and third notes), however, has been chosen in existing designs to be chromatic (in order to afford more steps within the octave). It would be interesting, for example to investigate whether the addition of harmonic content to reinforce the implied tonality would assist in the perception of interval.

5. REFERENCES

[1] B. Repp and A. Penel, "Rhythmic movement is attracted more strongly to auditory than to visual rhythms," *Psychological Research*, vol. 68, pp. 252-270, 2004.

[2] P. Shove and B. Repp, "Musical motion and performance: theoretical and empirical perspectives, in *The Practice of Performance,* ed. John Rink, Cambridge University Press, New York, USA, 1995, pp .55-83.

[3] B. Repp and G. Knoblich, "Perceiving action identity: How pianists recognize their own performances," *Psychological Science,* vol. 15, pp. 604-609, 2004.

[4] E.W. Large and M. R. Jones, "The Dynamics of Attending: How People Track Time-Varying Events," *Psychological Review,* vol. 106, no. 1, pp. 119-159, 1999.

[5] S.W. Brown and M.G. Boltz, "Attentional Processes in Time Perception: Effects of Mental Workload and Event Structure," *Journal of Experimental Psychology: Human Perception and Performance,* vol. 28, no. 3, pp. 600-615, 2002.

[6] P. Janata, "When music tells a story," *Nature Neuroscience,* vol. 7, no. 3, p. 203, March 2004.

[7] S. Koelsch et al., "Music, language and meaning: brain signatures of semantic processing," *Nature Neuroscience,* vol. 7, no. 3, p. 302, March 2004.

[8] P. Janata and E. Childs, "MarketBuzz: A Sonification of Real-Time Financial Data," *Proceedings of the Tenth International Conference on Auditory Display,* Sydney, Australia, July 6 – 9, 2004, pp.

[9] http://www.heatmaps.com

[10] J.J. Bharucha, B. Tillman and P. Janata, "Cultural Adaptation of the Brain to Music and Speech: an fMRI Study of Listening to Indian and Western Music, Hindi and English, Department of Psychological and Brain Sciences, Dartmouth College, *Unpublished.*

[11] A.D. Patel, E. Gibson, J. Ratner, M. Besson and P.J. Whitcomb, "Processing Syntactic Relations in Language and Music: An Event-Related Potential Study," *Journal of Cognitive Neuroscience,* vol. 10, no. 6, pp. 717-733, 1998.

[12] http://www.sibelius.fi/

[13] E. Tufte, *The Visual Display of Quantitative Information,* Graphics Press, Cheshire, CT, USA, 1983

[14] Accentus Trial Edition, http://www.accentus.com/

[15] E. Bonabeau, "You heard it here first," *Breakthrough Ideas for 2005, The HBR List, Harvard Business Review,* Feb. 2005.