# MARKETBUZZ: SONIFICATION OF REAL-TIME FINANCIAL DATA

*Petr Janata and Edward Childs*

[1]Department of Psychological and Brain Sciences and Center for Cognitive Neuroscience
Dartmouth College
6207 Moore Hall
Hanover, NH 03755
petr.janata@dartmouth.edu
[2]Accentus LLC
14 South Main St. Suite #2F
Hanover, NH 03755
edward.childs@dartmouth.edu

**ABSTRACT**

A system for the sonification of real-time financial data, currently in use by financial traders in five pilot projects, is described. Anecdotal feedback from the pilot projects suggests that the auditory display is more effective and consistent for monitoring the movement of volatile market indices. The same system has also been tested in two experiments carried out at the Department of Psychological and Brain Sciences at Dartmouth College. In the first experiment, subjects performed "change of direction" monitoring tasks of varying difficulty with and without auditory display. The results indicated a significant increase in accuracy when the auditory display was used. In the second experiment, subjects performed the same monitoring task with and without auditory display but were given a second, "number-matching" task which forced them to direct their visual attention away from the "change of direction" task from time to time. The auditory display increased accuracy more dramatically than in the first experiment, since the subjects were able to rely on the sonification to perform the "change of direction" monitoring task when they were distracted with the "number-matching" task.

## 1. INTRODUCTION

The requirement to monitor large numbers of data streams on visual displays is increasingly common in a variety of work environments: financial trading, electricity grid network monitoring, air traffic control, operating and emergency room, security and safety, to name a few. Frequently, the user needs to monitor the upward or downward movement of a variety of numerical indicators, such as the Dow Jones Industrial Average, patient blood pressure, percent load on a number of transmission lines, etc. In many cases, the user must perform a critical action if any of the indicators change by a certain amount, or approach an upper or lower target. In the visual mode, the user must constantly scan one or more displays to check changes in these indicators. User fatigue and other distractions often lead to inconsistent monitoring and failure to perform the required action in a timely manner.

Financial traders on electronic floors use anywhere from two to fifteen screens to monitor market indicators, electronic trading platforms and proprietary spreadsheets. Proprietary traders typically make three to four hundred trades in one day, requiring split-second decisions. Failure to respond in a timely manner (usually within two to three seconds) to an actionable condition may result in increased risk or a lost opportunity.

Prior work in the sonification of financial data has primarily been with historical data sets, in which data from one or more years is "played back" in a matter of minutes. Richard Voss was possibly the first researcher to experiment with the sonification of historical financial data (the price of IBM stock) [1]. He was primarily interested in exploring the "musical" characteristics of the rising and falling price of the stock over several decades. Kramer experimented with multidimensional displays of historical stock market indices [2]. Neuhoff, et al. [3] studied perceptual distortions when multidimensional stock market data is mapped to multiple acoustical attributes. Berger, et al. used filtered noise to sonify historical data [4]. Nesbitt and Barrass developed a multimodal display for a "depth of market" trading strategy [5]. In the present work, the sonification of various commercially available real-time data streams was carried out as an add-on to existing conventional visual displays (such as Bloomberg terminals).

The sonification of real-time data has been developed in other applications, for example, the pulse oximeter with signal sonification [6], the sonification of traffic on internet web servers [7][8][9], and the sonification of software bugs and program execution [10], to name a few.

There have been some multimodal studies in which auditory or haptic displays have been introduced to address the deficiencies in visual-only displays (e.g. to enhance the perception of space in GIS navigation systems) [11].

## 2. FINANCIAL DATA SONIFICATION SYSTEM

Following interviews with financial traders in varied environments (portfolio management, hedge fund, municipal bonds, government bonds, options, etc.) a prototype sonification system was developed to assist in the monitoring of various market instruments, including bid-ask stacks in electronic trading platforms.

Based on this research, four generic data behavior patterns were identified and sonified.

### 2.1. Relative Movement

The trader is aware of the general position of an index based on some benchmark, such as the opening price or the previous day's close, but desires an auditory display of relative movement (uptick or downtick). The behavior is displayed using two staccato notes of different pitches. If the pitch of the second note is higher than that of the first, an uptick has occurred, if lower, a downtick. Different indices may be distinguished by instrument (bassoon, piano, etc.), register and stereo location [relative.mp3].

The trader may adjust the size of the tick (significant movement) which will trigger the auditory display, the volume, tempo and stereo location. The significant movement is usually set to be small (one or two points).

The relative movement sonification is typically used in intensive proprietary trading situations in which small movements in market indicators require immediate action.

### 2.2. Absolute Movement

The trader wishes to track the absolute movement of an index based on a reference point such as the opening price or the previous day's close, over the course of a trading day. The behavior is displayed using three staccato notes. The first note is a reference pitch (always the same) and represents the reference point. The interval between the pitch of the third note and the first note represents the difference between the current price and the reference point, based on the significant movement set by the user. For example, if the significant movement is set to 25 points, then if the index is up 75 points from the reference, the interval between the third and first note would be up a minor third. The interval between the second note and first note represents the difference between the price and the reference the last time a sonification occurred. To continue the previous example, if the index had reached 100 points from the reference at the previous sonification, the interval of the second note would be up a major third from the first note. (It follows that the third note will now be one half step below the second, to indicate a downtick). As with relative movement, different indices are distinguished by instrument, register and stereo location. [absolute.mp3]

The absolute movement sonification is typically used in less intensive situations in which the trader wishes to be informed of broader movements in market indicators over the course of a day.

### 2.3. Approach to a Target

In addition to tracking the absolute movement of an index or equity, the trader wishes to follow the approach of the price to a technical, such as a 30, 50 or 200 day moving average, a Bollinger band or some other user-determined target, such as a change of 2% from the opening price. The behavior is displayed using the three staccato notes of the Absolute Movement and the addition of a fourth note. The interval between the fourth note and the first, reference note represents the difference between the target and the reference (one half step = one significant movement, defined the same as for the Absolute Movement). If the current price is at some distance from the target, the fourth note does not sound. When the current price approaches the target within a user-defined threshold, the fourth note sounds. As the current price gets closer to the target, the duration and loudness of the fourth note increases. When the current price reaches the target, the fourth note sounds at maximum loudness and at twice the maximum duration. [target.mp3]

### 2.4. Bid-Ask Stack

Bid-Ask stacks are common in the trading of most securities. The "bid" is the amount someone is willing to pay for a given security. The "ask" is the amount at which someone is willing to sell the security. When the (lower) bid and the (higher) ask converge to the same price, the trade occurs. The Bid-Ask Stack is a column of bids and asks ranked in order of the highest bid, and the lowest ask. The columns also include the size of the bid and ask, sometimes known as the "show".

In fast-paced interactive trading situations, the trader wishes to know immediately if there is a large show in any of the securities being traded. The Bid-Ask Stack is displayed using a two-note *tremolando* followed by a long tone. If the data represents a bid, the *tremolando* begins on the lower of the two notes and ends (with the long tone) on the lower note. The converse is true if the data represents an ask. The difference in pitch between the two notes is proportional to the difference between the bid and the ask (the "spread"). The length of the *tremolando* is proportional to the show. The display will occur only when the show exceeds a threshold set by the trader. The identification of the security, as with the previous sonifications is by the instrument and the stereo location [bidask.mp3].

The Bid-Ask Stack sonification displays four pieces of information:

1. Which security.
2. Whether it is a bid, or an ask.
3. The spread between the bid and the ask.
4. The show.

In the interactive electronic trading environment, a large bid may be "answered" by a corresponding ask, and the spread may narrow as the bid and ask prices are adjusted. These trends may be followed through the auditory display with decreased need to consult the visual display constantly.

## 3. PROTOTYPE PROGRAM RESULTS

The prototype program has been running since December of 2002. Repeated visits have been made to the prototype sites to solicit feedback from the traders and to make improvements to the software.

All prototypes have been in continuous use since the installations. None of them have been turned off. Users cited various advantages to the auditory display:

- Easy to learn, thanks to reinforcement from the accompanying visual displays. Only minimal training is necessary.
- Quantitative information on the movement of market indices could be obtained from the auditory display without having to consult the visual display.
- The ability to set thresholds and significant movements gave traders control over when the displays sounded and led to more relaxed and consistent monitoring, even when distractions occurred.
- The immediacy of the auditory display was at times more effective than the visual in communicating market anomalies.

One trader commented that executing trades via keyboard strokes in response to the auditory displays was like playing a musical instrument in reverse.

## 4. EXPERIMENTAL PROGRAM

In an effort to derive some quantitative measure of the efficacy of the auditory display, a series of experiments have been planned. In the first set of two experiments, the Relative Movement sonification was tested in a controlled environment. The subjects were presented with visual displays similar to a market data terminal such as Bloomberg. They were asked to perform certain monitoring tasks with and without auditory display.

In the first experimental study, we simulated different levels of perceptual load by varying the number of locations that subjects had to monitor simultaneously in order to respond to target events. In addition we varied the auditory information subjects received about the occurrence and quality of significant events. Overall we found that auditory information significantly increased the proportion of correct detections and reduced the number of missed target events. Furthermore, a sonification strategy that provided information about the quality of a significant event further enhanced response accuracy.

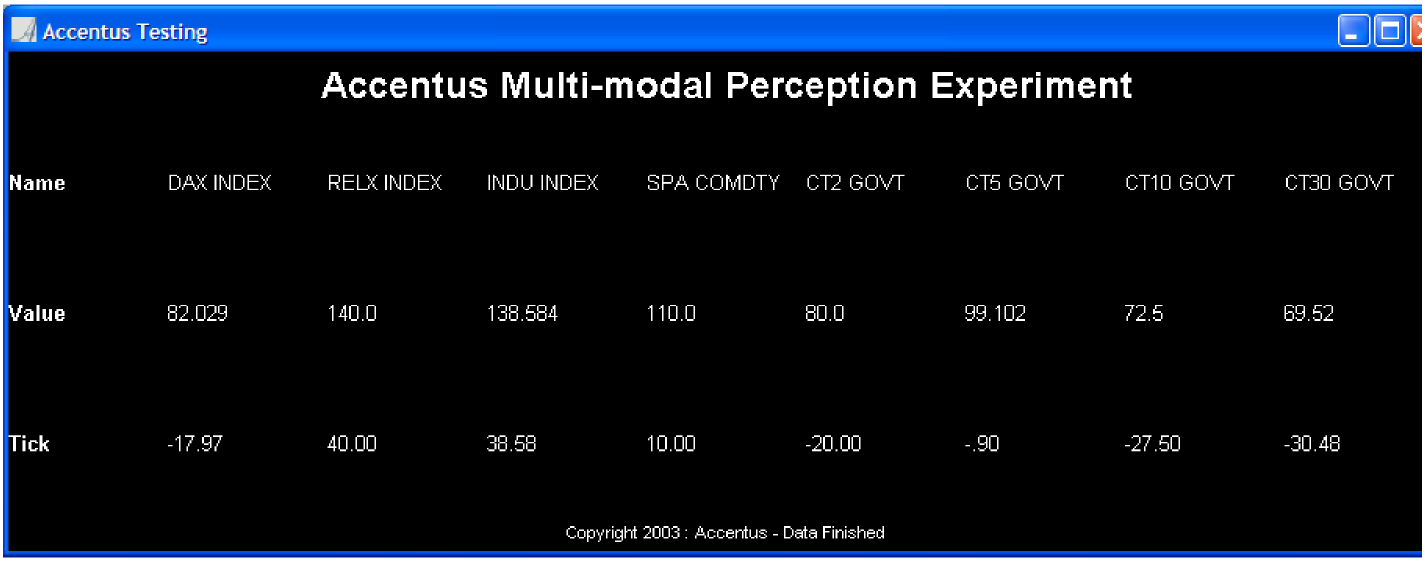


Figure 1. *Layout of the testing screen in 8-stream mode. Subjects had to attend to the bottom row labeled "Tick" and respond whenever the number changed upward or downward and landed on a multiple of ten..*

## 5. EXPERIMENT 1 METHODS

12 young adults from the Hanover, N.H. community participated in an experimental session lasting 90 minutes. None of the subjects were familiar with the experiment. Subjects were required to monitor changing numbers in either 2, 4, or 8 columns distributed across a window on a computer monitor (Figure 1). Whenever a "significant move" occurred, the subjects' task was to press a key corresponding to the column and direction of the significant move as quickly as possible. A move was considered significant whenever the number at one of the different locations moved up or down and fell on a multiple of ten. If the significant move was more positive than the previous number, then the subjects pressed a key on the top row of letters, whereas if the multiple of ten was more negative than the previous number, they pressed a key on the bottom row of letters. The numbers represented the "change on the day" and could therefore be positive or negative. Similar to a Bloomberg terminal, whenever a number changed, a large square centered on the number flashed red if the change was downward and green if the change was upward, irrespective of whether the change was a significant move.

Two factors were manipulated during the experiment: 1) The number of information streams (2, 4, or 8) that had to be simultaneously monitored, and 2) the type of information the subjects received. In the "No Sound" condition, subjects saw only a flash at the appropriate location, but had to look to the location to make a significant move judgment. In the "Beeps" condition, subjects heard two beeps of the same pitch separated by a brief (ca. 250 ms) interval whenever there was a significant move. Thus, the beeps provided information that a response was required but provided no information about the correct column or the direction of the change. In the "Sonification" condition, a pair or notes was presented whenever a significant move occurred. The timbre (instrumental quality) of the notes represented a specific column. For instance, a low bell sound represented numbers in the leftmost column, whereas a piano sound represented the 3rd column from the right. Furthermore, the 2nd of the two notes was either higher or lower than the 1st note to indicate upward and downward movements, respectively. Thus, once they learned the mappings between timbre and spatial location in the Sonification condition, subjects could make a correct response without having to look at the screen.

During a practice run, the 9 conditions were presented in order of increasing complexity, starting with "No Sounds, 2 streams" and progressing through "Sonification, 8 streams." Each condition lasted 3 minutes, and contained 45-55 significant moves each. The amount of time between significant moves ranged from 3.3 to 4.3 seconds (mean = 3.7 s). The average interval between all events was 285-387 ms (mean = 334 ms). Following the practice run, subjects repeated all 9 conditions in a pseudo-random order. All key presses were recorded to a data file for offline analysis. The data for one subject were discarded because the response data indicated that she was pressing invalid sets of keys in a majority of the conditions.

Detection accuracy was the primary performance criterion. The first event following a significant move was taken as the response to the significant event and classified as either "correct", meaning the correct stream and correct direction were selected, "correct stream but incorrect direction," or "incorrect stream." If no response was made prior to the next significant move, the response was classified as a "miss". The proportions of correct responses and misses were entered into repeated-measures analyses of variance (ANOVAs) with 2 factors (number of streams, and information condition). All significance values are reported with Greenhouse-Geisser corrections.

## 6. EXPERIMENT 1 RESULTS

Figure 2 illustrates that providing auditory information about the occurrence of significant events significantly increased response accuracy (averaged across the number of information streams) from 75% with no sound, to 84% with warning beeps, to 88% with 2-note sonification [$F(2,20)=10.56$, $p < 0.003$]. The difference in detection accuracy between Sonfication and Beeps in the 4 and 8 stream conditions was not significant [$F(1,10)=1.77$, n.s.].

The addition of sound significantly reduced the number of missed significant events [$F(2,20)=74.00$, $p<0.0001$] (Figure 3), though there was no difference between the simple alerting function of the pair of beeps and the 2-note sonification that provided information both about identity of the information stream and the direction of the significant move. A further analysis of errors that were committed pointed to a significant difference between the auditory conditions (Figure 4). Specifically, under conditions of higher visual load in the 4 and 8 stream conditions, 2-note sonification significantly reduced the number of errors about the direction of a significant move [$F(1,10)=5.13$, $p<0.05$]. Thus, warning beeps helped orient a subject to the correct information stream but left them more uncertain about the direction of a significant move. Because the direction information was incorporated into the 2-note sonification, the direction of the move was classified correctly more often.

## 7. EXPERIMENT 1 CONCLUSIONS

The results of this study indicate that adding auditory information to a visual monitoring scenario significantly enhances the detection of target events. At first glance, it may seem paradoxical that increasing the amount of information actually improves performance. Two factors likely contribute to the enhanced performance. First, a sound signals the occurrence of a significant event. Thus, the continually changing number and flashes don't have to be monitored continually in order to detect the occasional target event. Rather, the sound serves to focus attention at the necessary times. Second, the extra information is presented to a sensory modality that isn't operating at full capacity and is capable of monitoring the environment in parallel with the visual system.

Overall, the results showed only small, if any, differences between the Beeps and Sonification conditions. Thus, the performance enhancements appear to reflect the alerting properties of the sound. However, we did observe an important difference between the auditory conditions. Namely, the Sonification strategy, in which the two notes compactly conveyed both stream and direction-of-movement information, led to fewer errors (than in the Beeps condition) in judging the direction of movements when 4 or 8 streams had to be monitored simultaneously. This difference suggests that we have identified a point at which the benefits of simply alerting to the occurrence of target events have been maximized. In other words, as the complexity of the monitoring environment increases, simply knowing that a significant event occurred is not enough. Instead, it is also necessary to know what information stream the relevant event belongs to, along with information about the change that the event represents. In the context of our experiment, the 2-note sonification satisfies the latter requirements.

Our monitoring interface was an extremely simplified version of a financial trader's data monitoring environment. A more realistic interface would consist of more densely packed information streams, multiple windows on a single monitor, and several monitors. However, even under our simplified conditions we found that strategic auditory information assists the subject.

Finally, we observed enhanced performance with auditory information following a single 45 minute training session. In the Sonification condition, subjects had to learn the associations between timbres and spatial locations and translate perceptual judgments about changes in pitch to appropriate finger movements. None of the subjects were familiar with the stimulus/response mappings prior to the experiment. Their superior performance in the auditory conditions indicates that novices are able to quickly harness the utility of auditory information. We expect that with further training in the Sonification conditions, subjects' performance will improve further.

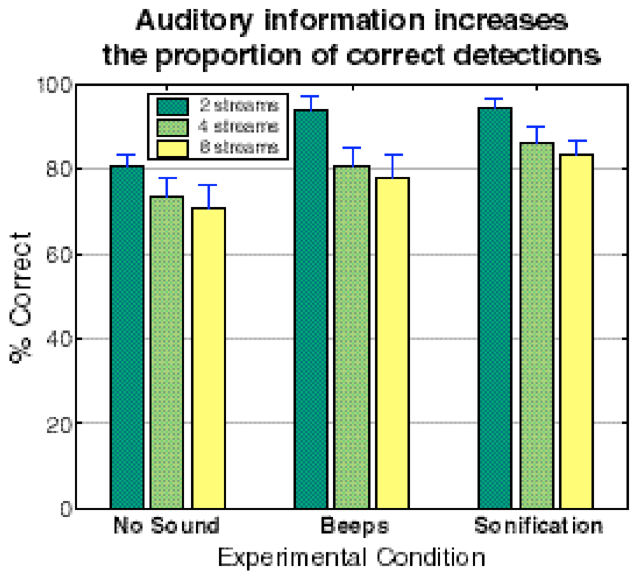


Figure 2. *Detection accuracy as a function of the number of information streams that must be monitored and the type of information provided to the subject. In this and other figures, the error bars represent 1 standard error of the mean.*

## 8. EXPERIMENT 2 METHODS

In the second experiment, the methodology was the same as for the first experiment, except that the "change direction" task to detect ten point movements was relegated to a secondary status, and a primary visual task was added. In the financial trading environment, a visual monitoring task such as that of the first experiment is complicated by the fact that the trader must simultaneously monitor other conditions on other visual displays, place orders on the telephone, confer with colleagues, etc. In such cases, the auditory display is expected to increase the monitoring accuracy.

In the second experiment, see Fig. 5, the subjects were asked to monitor the four large numbers displayed below the original two to eight stream monitoring task. If any two of the four numbers match (e.g. the first and the fourth in Fig. 5), the user is required to press the spacebar.

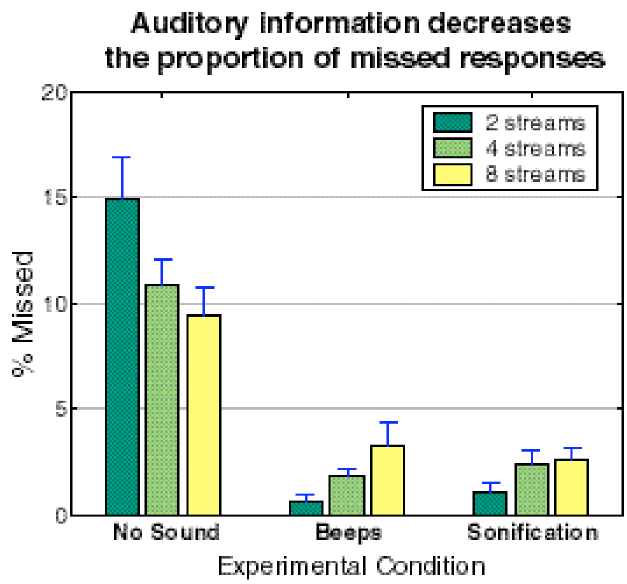


Figure 3. *Proportion of missed responses as a function of the number of streams and information provided.*

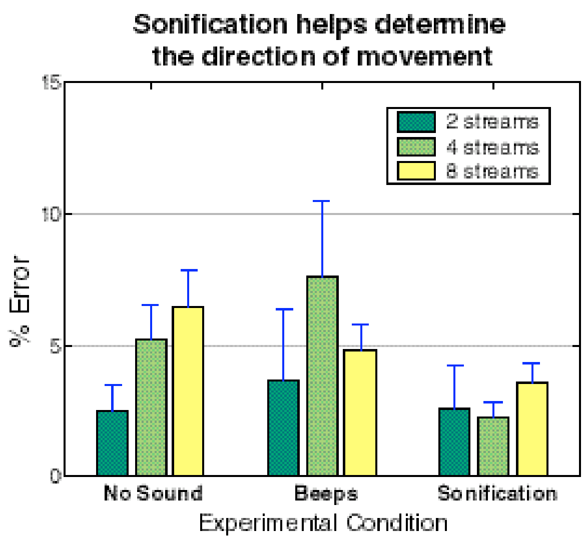


Figure 4. *Percentage of incorrect responses in which the correct stream was selected but the judgment about the direction of the significant movement was incorrect.*

## 9. EXPERIMENT 2 RESULTS

In Experiment 2, we tested the hypothesis that providing timbral and pitch cues would facilitate performance on the "change direction" task (as in Experiment 1), despite the addition of second task in which subjects detected matching pairs of numbers from among 4 numbers arrayed at the bottom of the screen. Figure 6 shows the target detection accuracy in both tasks as a function of the number of data streams that subjects had to monitor and the auditory information provided to them.

Repeated-measures analysis of variance (ANOVA) was used to test of the main effects of task type, sonification type, and number of data streams. The main effect of task type was highly significant [F(1,11)=119.76, p<0.0001], indicating that overall target detection performance was better in the "change direction" task. Of primary interest was the significant main effect of sonification type [F(2,22)=10.05, p<0.0008]. This effect indicates that accuracy (combined across both tasks) differs as a function of the auditory information provided to the subject. Figure 6 shows that average performance was highest in both tasks when both the data stream and direction of change were denoted by sound. This result was further tested through a series of pairwise comparisons. The two-note sonification resulted in a significant 13.9% increase in detection accuracy relative to the silent condition [t(22)=3.99, p<0.001], and a 7.6% increase relative to the warning condition [t(22)=2.14, p<0.05]. A 6.3% increase in accuracy with warning sounds relative to silence showed a trend toward significance [t(22)=1.77, p<0.10].

Complete auditory information led not only to a significant performance increase in the "change direction" task, but also to a significant 8.2% accuracy increase in the simultaneously performed "matching" task [t(22)=2.34, p<0.05, relative to silence]. Simple warning sounds did not yield a significant performance enhancement on the matching task [t(22)=1.45, n.s.].

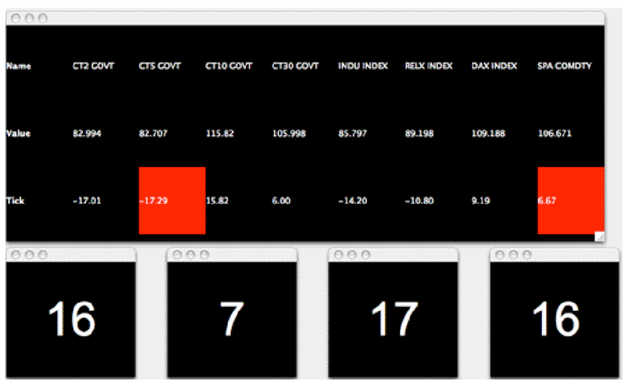


Figure 5. *Interface for Experiment 2 in which a distracting, number-matching task is added (the four large numbers in the lower section of the display).*

Examining the types of errors made in the change direction task further highlighted the informational utility of the 2-note sonification. Figure 7 shows that the percentage of missed response doubles when no auditory information is provided. Without auditory information, the percentage of missed target events was 19%. In the 2-note and warning sound conditions, this percentage decreased significantly to 6.2% [t(22)=-5.18, p<0.0001] and 8.7%

respectively [t(2))=-4.10, p<0.0005]. The 2-note and warning sound conditions did not differ in their efficacy of preventing missed responses [t(22)=-0.98, n.s.].

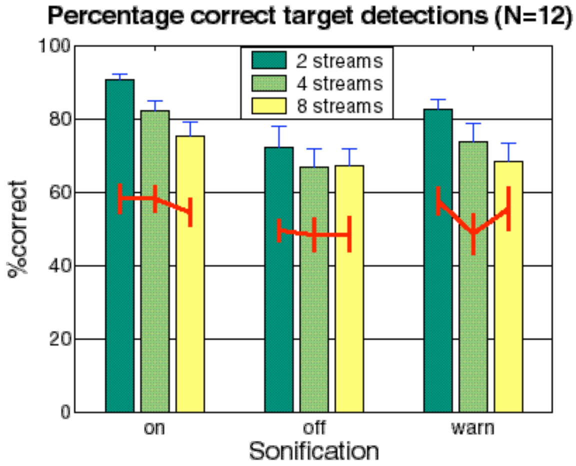


Figure 6. *Target detection accuracy as a function of the number of data streams and auditory information. In the 'on' condition, the data stream was represented by the timbre of 2 consecutive notes. The direction of the change to which subjects had to respond was indicated by the difference in pitch height of the 2nd note relative to the first. In the 'off' condition, subjects received no auditory information. In the 'warn' condition, subjects heard two consecutive notes at the same pitch whenever a significant change occurred. The timbre did not provide a cue as to the data stream in which the change occurred. The red lines indicate performance on the "number matching" task which oriented the subjects' visual attention to the bottom of the screen. All error bars are 1 standard error of the mean.*

Figure 8 shows that directional information in the 2-note sonification significantly facilitates correct identification of the direction of the change. While the warning sounds alert a subject to the occurrence of a target event, they nonetheless require the subject to orient to the visual display in order to determine the direction of the change. The lack of the directional information in the sound results in a significant 6% increase in errors [t(22)=4.89, p<0.0001]. In fact, the warning sound condition yielded significantly more errors (3.2%) of this type than the silent condition [t(22)=2.65, p<0.02]. Thus, while warning sounds cause fewer target events to be missed, and do facilitate overall target detection accuracy, they leave the subject less certain about how to respond. Aside from facilitating detection of target events, the 2-note sonification also results in 2.8% fewer directional errors than the silent condition [t(22)=-2.32, p<0.05].

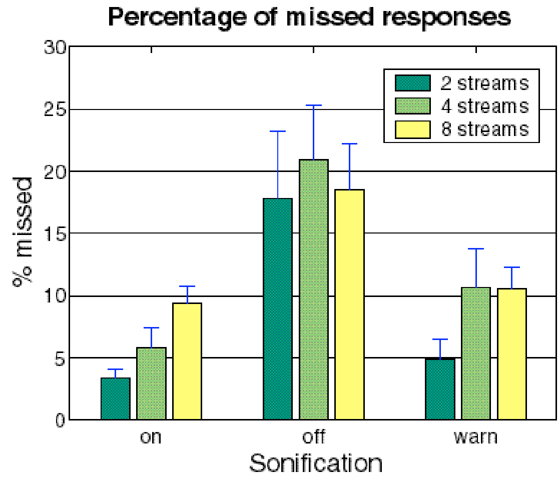


Figure 7. *Auditory information reduces the percentage of missed target events in the "change direction" task. See Figure 6 for an explanation of the Sonification conditions.*

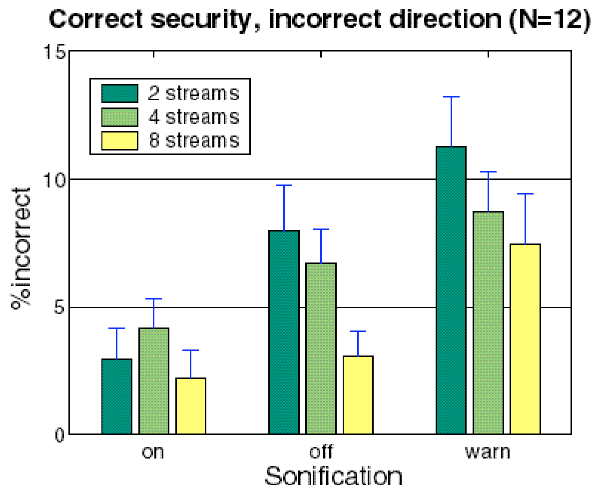


Figure 8. *The 2-note sonification scheme significantly reduces the number of errors committed in identifying the direction of a target change. See Figure 6 for an explanation of the Sonification conditions.*

## 10. EXPERIMENT 2 CONCLUSIONS

The 2-note sonification strategy, which identifies both the data stream and the direction of change in the target event, facilitates performance of both tasks in a dual-task visual monitoring situation relative to conditions in which subjects receive no auditory input or only warning sounds that signal the occurrence of a target event. The less than perfect performance on both tasks, even in the 2-note sonification condition, indicates that capacity of the visual attention system has been exceeded, and detection of target events in one task comes at the cost of missing target events in the other task.

Our data indicates that the overall attentional capacity has not been exceeded, however. The performance increases in the change direction task suggest that subjects extract the meaning of auditory input that is designed to help them perform this task. By relying on their audition to perform this task, they free their visual attention to perform the visual target detection task that is unaccompanied by auditory cues.

## 11. FURTHER WORK

A series of follow-on experiments is planned to test the efficacy of the remaining three generic data behavior pattern sonifications outlined in Section 2, in controlled environments similar to those described in the first two experiments.

Furthermore, we plan to increase the complexity of the visual display further in order to dissociate the relative merits of auditory alerts and information-containing sonification strategies.

Additionally, an experimental program with trainee financial traders in more realistic environments is also planned.

The ultimate test of a sonification system intended for commercial use is the acceptance of the user community. To date the feedback from this growing community is very encouraging.

## 12. ACKNOWLEDGMENTS

The assistance of Benjamin E. Childs in programming the user interface for the experiments is gratefully acknowledged.

## 13. REFERENCES

[1] http://web.gc.cuny.edu/sciart/0102/fractals.html

[2] G. Kramer, "Some Organizing Principles for Representing Data with Sound," in *Auditory Display,* Addison Wesley, MA, USA, 1994.

[3] J. G. Neuhoff, G. Kramer and J. Wayand, "Pitch and loudness interact in auditory displays: Can the data get lost in the map?", *J. Exp. Psychol.-Appl,* vol 8, no. 1, pp. 17-25, Mar. 2002

[4] http://www-ccrma.stanford.edu/groups/soni/index.html

[5] K. V. Nesbitt and S. Barrass, "Evaluation of a Multimodal Sonification and Visualization of Depth of Stock Market Data", in *Proceedings of the 2002 International Conference on Auditory Display,* Kyoto, Japan, July 2-5, 2002, pp. 233-238

[6] J. L. Reuss, "Pulse Oximeter with Signal Sonification", *United States Patent No. 6,449,501,* Sep. 10, 2002.

[7] M. Barra, et al., "Personal Webmelody: Customized Sonification of Web Servers", in *Proceedings of the 2001 International Conference on Auditory Display,* Helsinki, Finland, July 29 – August 1, 2001, pp. 1 – 9.

[8] M. H. Hansen and B. Rubin, "Babble Online: Applying Statistics and Design to Sonify the Internet", in *Proceedings of the 2001 International Conference on Auditory Display,* Helsinki, Finland, July 29 – August 1, 2001, pp. 10 – 15.

[9] M. Gilfix and A. Crouch, "Peep (The Network Auralizer): Analyzing your Network with Sound" in *2000 LISA XIV,* New Orleans, USA, December 3 – 8, 2000, pp. 109 - 118

[10] P. Vickers and J. L. Alty, "Using music to communicate computing information", *Interact. Comput.,* vol. 14, no. 5, pp. 435-456, Oct. 2002.

[11] W. Jeong and M. Gluck, "Multimodal Geographic Information Systems: Adding Haptic and Auditory Display," *Journal of the American Society for Information Science and Technology,* vol. 54, no.3, pp. 229-242, 2003