

The Audio Abacus: Representing Numerical Values with Nonspeech Sound for the Visually Impaired

Bruce N. Walker, Jeffrey Lindsay, and Justin Godfrey

Sonification Lab, School of Psychology

Georgia Institute of Technology,

654 Cherry Street

Atlanta, GA, USA 30332-0170

404 894 8265

bruce.walker@psych.gatech.edu

ABSTRACT

Point estimation is a relatively unexplored facet of sonification. We present a new computer application, the Audio Abacus, designed to transform numbers into tones following the analogy of an abacus. As this is an entirely novel approach to sonifying exact data values, we have begun a systematic line of investigation into the application settings that work most effectively. Results are presented for an initial study. Users were able to perform relatively well with very little practice or training, boding well for this type of display. Further investigations are planned. This could prove to be very useful for visually impaired individuals given the common nature of numerical data in everyday settings.

Categories and Subject Descriptors

J.0. [Computer Applications]: General

General Terms

Design, Experimentation, Human Factors.

Keywords

Data sonification, auditory display, visually impaired users, value estimation.

1. INTRODUCTION

In the broadest sense, many auditory displays have been designed to present information about numbers, and the relations between those numbers, to a listener. Such sonifications [1] have been used in cases where the primary task of the listener is to extract meaning from the auditory display, and also when the display is used in conjunction with a visual or other type of task. In either case (unimodal or multimodal), there are at least two distinct and basic information extraction tasks that the listener can perform with the auditory display: *trend analysis* and *point estimation*.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ASSETS'04, October 18-20, Atlanta, Georgia, USA.

Copyright 2004 ACM 1-58113-911-X/04/0010...\$5.00.

Trend analysis is the task of determining patterns in the data set over the course of several data points, such as determining if the price of a stock is rising or falling. Point estimation is the task of determining as nearly as possible the exact value of a data stream at a specific point in time or space. For example, a broker might want to know the specific price of a stock on a particular date. Both tasks are important, though they may be differentially important in particular task settings. An auditory display designer must assess the needs of the user, and, following such a task analysis, create a display that supports the correct goals.

Trend analysis and point estimation are both tasks that occur frequently during everyday activities. For sighted individuals often a graphical or numeric display is used to convey trend information or specific values. Examples range from the weather report to stock market data, from digital clocks to the latest employment statistics. However, for those with visual impairments such displays can be difficult or impossible to use. In science-related fields of employment, the situation is even more difficult, since nearly all commercially available scientific software is exclusively visual in the display of data.

Sound is being considered in scientific software, but there are many ways auditory displays could be used to enhance the daily lives of both sighted and visually impaired persons. Consider the checkout line at the grocery store, where many cash registers have some type of simple tone to signal a successful scan/entry of an item to be purchased. If a sighted individual wishes to check the price of an item as it is scanned, all they need do is glance at the screen on the register. For a visually impaired individual, however, there is no way to verify the price of an item short of a verbal inquiry. What if instead of a simple tone each time an item is scanned, there were a sound that indicated the exact price of the item? This would allow the shopper to hear the prices of items as they are scanned and also allow the cashier to have a nonvisual verification that items have been priced correctly. Most sonifications are not created with the goal of supporting point estimation tasks. For this reason it is vital that approaches that facilitate this task, such as the Audio Abacus, be developed and investigated.

Many sonifications use changing data values to drive changes in pitch, loudness, tempo, or other sound qualities [e.g., 2, 3-6]. Some sonifications have used a granular approach, wherein a display is built up from small elements, ranging from brief tones to sampled words to frog calls [e.g., 7, 8, 9]. There has, understandably, been a growing interest in determining how best to design such displays to support a range of analysis tasks. For example, Flowers and Hauer [10]

found that important characteristics of data, such as slope, shape, and level of a graph were perceptually salient when sonified, and that people could interpret the trends or shapes of line graphs where each data point was represented with a musical note. Flowers, Buhman, and Turnage [11] found that auditory scatter plots are effective at conveying magnitude and sign of correlations. Individuals have also been able to match auditory graphs with two data series to visual graphs [12]. Brown and Brewster [13] suggest that individuals can interpret and draw at least two data series from an auditory graph. Barrass, [14], Walker, [15], and others have looked at the appropriate choice of mapping, scaling, and polarity, to further enhance the effectiveness of auditory displays. Nevertheless, most of the displays have been designed in a fundamentally similar manner: a time-series of sounds represents a time series of data. While this basic approach has been shown to be very effective for a great variety of trend analysis tasks, we contend that it is not ideal for point estimation.

1.1 Improving Point Estimation

Smith and Walker [16] have pointed out some sound design techniques that can be used to improve point estimation. Adding additional context sounds that perform like axes and tickmarks in an auditory graph can improve performance; this approach has been used in sonification designs usually because the added sounds seemed to help, and not for any theoretical reason. While Tufte [17] has developed several guidelines for displaying points on a visual graphs, it is still unclear which, if any, of these principles apply to sonification concepts.

Training listeners how to do the complex sub-steps (e.g., interpolation) required in a point estimation task is another way to improve performance [16]. However, there are limits to the performance that can be obtained in any of these examples. According to Quinones and Ehrenstein [18], training is defined as attempting to change individuals in a way that is consistent with task requirements, and as a way of applying principles of human learning and skill acquisition. Training to improve performance on sonification-based tasks is a relatively new area of research. Upson [19] saw small improvements in students' mathematical abilities after using sonification exercises for educational purposes, after having participants go through a short tutorial on his SoundGrid program. Peres and Lane [13] conducted a study in which participants were asked to identify a box plot from a multiple choice selection after hearing a sonified representation of the visual graphs. However, performance on the task did not increase with practice despite a 15 minute training session and 50 experimental trials. Was this lack of performance increase due to a real limitation on human abilities to perform a point estimation task? One of the challenges pointed out by Peres and Lane [13] and others is that getting good numerical resolution is hard because of the relatively large just noticeable differences (JNDs; and perhaps more importantly, the *reliably* noticeable differences) in the auditory system. That is, if one uses, for example, MIDI notes to represent changes in a value, there is a limit to the actual numbers one can represent, as well as a limit to the pitch recognition abilities of the listener. To represent 100 separate values, 100 perceptually distinct pitches need to be available. In practice this is very untenable. We suggest that there must be other, categorically different ways to represent




Figure 1. Sample abacus, showing each wire with different beads. In this version, counting is aided by the beads in the upper register to indicate groups of 5. Simpler versions just have 9 beads on a single wire for each unit or digit.

exact data values, and present one example of a different approach to this thorny problem.

1.2 The Abacus

In order to overcome the need to have a great number of sounds to represent a large range of data values, we took inspiration from the concept behind the abacus. The abacus (see Figure 1) is an ancient counting device (though still used today in many places!) that has a frame that holds wires, on which beads can move freely. In the simplest form, each wire holds a set number of beads (e.g., 9), and the beads on each wire represent different units. For example, the beads on the right-most wire each represent one item (scoop of grain, horse, etc.), the next wire represents tens of units, then 100s, 1000s, and so on. To start (i.e., to zero the abacus), all beads are slid to one side of the abacus. That is, the absence of any beads on the “counting” side of the abacus indicates zero in that unit. To record two units, two beads from the right-most wire are moved to the opposite side of the frame. Moving two of the beads on the second wire across would indicate the addition of 20, for a total of 22. In actual practice, most “modern” abaci use a more sophisticated counting scheme, but the simpler version described here is sufficient for our purposes. Building on this approach, we decided to use a series of sounds to represent the different units, with the pitch of each sound representing a value from 0 to 9. In this way, only ten distinct pitches need to be used to represent value within a unit, and with just four separate sounds one can represent the range of 0-9999. With this concept, we have created the Audio Abacus.

2. The Audio Abacus

The Audio Abacus is a Java program designed to allow the sonification of discrete data points, for example “582”. The basic idea is that in order to represent an exact numerical value, a series of brief sounds (in this example, three sounds) are played in succession. Each of the sounds can, itself, be one of ten notes produced by a MIDI instrument. For example, if zero is mapped to MIDI note 60, then “1” could be mapped to note 61, “2” to note 62, and so on up to note 69. To “play” the number 582 the program would play a three-note series composed of MIDI notes 65, 68, and 62. The first sound obviously represents the 100s digit, the second sound represents the 10s, and the third sound represents the




Figure 2. Screen capture of the Audio Abacus. Generates sounds from data values. See text for description.

1s. Given this basic scheme, there are a huge number of additional attribute settings that one can make, in theory. In order to experiment with this form of number presentation, and to see how successful it might be in any of its various incarnations, we created the Audio Abacus. The Abacus software has a multitude of settings that provide for a wide range of possible manipulations. Many of these settings are currently being studied and manipulated in order to determine their effect on users' ability to interpret sonified data. Once that is well understood, we can set defaults in the software to maximize performance right from the beginning, while still allowing the experienced user to adjust settings flexibly to suit her or his preferences or needs.

2.1 Pitch

In the example above, the values within each unit or digit were mapped onto sequential MIDI notes. Using musical notes has the benefit that the sounds are largely equated in terms of perceptual separation (notes 60 and 61 are the same distance apart in perceptual space as notes 63 and 64, so that the numbers "0" and "1" are heard as having the same separation as the numbers "3" and "4"). However, there is no reason that the numbers need to be exactly one note apart. They could be two or five notes apart, in order to make the tones (and therefore, the numbers) more distinguishable. In order to support this flexibly, a pair of sliders allows the minimum and maximum note numbers to range from MIDI note 0 to 127. The tones representing the digits 0 – 9 are equally spaced between the minimum and maximum pitch value. For the current study (described below), minimum pitch was set at note 40 and maximum at note 90.

2.2 Pan

In the basic three-note sequence described at the outset, all the notes could be played through one central speaker (or over headphones with pan set to the center). Alternatively, each sound could be played in different spatial locations. For example, the 100s could be to the left, the 10s in the center, and the 1s off to the right. In that way, the unit or digit would be doubly mapped to the order in the series (largest played first) and to the spatial location it appears to come from (largest to the left). To support this panning possibility, we again use two sliders to set the minimum and maximum "panning position". For example, a minimum pan value of 0 and a maximum of 127 would cause the digits of the number to be equally spaced through the listener's sound image, ranging from all the way to the left (-90°) to all the way to the right (+90°).

To further illustrate the combination of pitch and pan, consider the number "145" being sonified. The listener would hear the note 61 (representing "100") played on the left side of their sound image, the note 64 ("40") in the center, and the note 65 ("5") to their right. The user would be required to add up the digits (100 + 40 + 5) to comprehend the overall value of 145. Note, however, that for the current study, the pan function was not used.

2.3 Decimal pass

Representing the numbers 145 and 1.45 (e.g., \$1.45) is conceptually the same—each sound represents a unit or digit, but in the case of decimal digits the units are tenths

and hundredths. What does change is communicating to the user that there is a decimal, and indicating where it is in the sequence. First of all, a decimal point is represented by the sound of a high hat cymbal crash (percussive, largely non-pitched). The way in which a decimal point is represented within the number sequence can be changed using three radio buttons in the Audio Abacus. The radio buttons determine the relationship of the other tones to the cymbal crash. This is especially relevant when the spatial panning is being used.

The first setting, “Single Pass L -> R”, simply plays the series of tones from left to right. Therefore, if the number 145.34 is played, five digits are divided into the sound image, with a symbal crash between the third and fourth sounds in the series. If the sound is spatialized, the decimal occupies a spatial location in keeping with its order in the series (so in this case, there are actually six sounds, in six successive spatial locations).

The second setting, “Double Pass L -> R . L -> R”, plays all of the whole units in series, left to right, plays the decimal sound, and then returns to the beginning (left) of the sound image before the digits after the decimal point are played. For example, when the number 145.34 is played on this setting the digits “1”, “4”, and “5” will be divided into three equal sections of the listener’s sound image (also influenced by the Pan setting described above). After the decimal cymbal crashes, the sound image is “reset”, and the digits “3” and “4” will be divided into two equal sections of the user’s sound image.

The final setting, “Bounce Pass L -> R . R -> L” plays numbers to the left of the decimal point from left to right, similar to the first setting. However, the numbers to the right of the decimal point are played going from right to left across the sound image. Consider, again, the number 145.34. For the decimal digits “3” and “4”, the “3” would be played on the right of the sound image and “4” would be played on the left.

Of course, many other possible settings could be implemented. These are just three that we thought were conceptually distinct, and therefore built into the software. For the current study, only the first setting, “Single Pass L -> R” was used.

2.4 Ratio

The amount of time that each sound in a series plays can also be adjusted on a per-digit basis. For example, the 100s unit sound could be 3 seconds long, with the 10s sound 2 seconds, and the 1s sound 1 second long. This might assist in distinguishing the different units (as the spatial separation does), and it might also allow for more processing time for the larger units (100s), which might be more relevant. However, if all the data values are changing between 194 and 199, then the largest unit is probably less interesting, so it need not play for so long, in relation to the other digits. In that case a sound length ratio of 1:1:1 or even 1:1:3 might be most appropriate. To support this flexibility, sliders allow the duration of each note in front of the decimal to be set relative to one another. For example, if the ratios were set to 3, 2, and 1, then in the number “145” the digit “1” would have the longest duration, “4” slightly less, and “1” would have the shortest duration. The precise length of one unit of time is determined by yet another slider, namely the Time Scale setting, described below. For the current study, all

of the ratios were set at 1 until further experiments call for manipulation of these values.

2.5 Time Scale

This slider changes the duration of one beat for the series of tones. Thus, if the 100s unit were set to have a ratio value of 3 beats, and each beat is set to 1 second, the 100s would play for 3 seconds. It ranges in value from 1/5 second to 5 seconds. For the current study, this slider was set at 1 second. Note that by having this seemingly complicated ratio setup, the relative lengths of the sounds can be maintained, but the whole series can be sped up or slowed down. This helps with early training trials being played slowly, and then later trials can be played quickly, without changing the relative sound lengths.

2.6 Decimal Beat

Allows a separation between the end of the tones representing the digits to the right of the decimal point and the start of the tones representing the digits to the left of the decimal point. For the current study this setting was left at 1.

2.7 Instrument panel

The default MIDI instrument used to play the notes is the piano, but the Abacus allows the user to choose any MIDI instrument for representing the digits. The instrument panel in the application shows the choices, and lets the user easily pick an instrument.

Because the program is written in Java, it can operate on any system with Java installed. The code itself resides in a Java Archive (JAR) file for simply portability and execution. Abacus imports numbers from text files, most commonly comma-delimited. Alternatively, users can also type numbers directly into the cells provided, to have them sonified. Clicking on a particular cell will play just that number, whereas clicking the “Play” button plays the whole data set, with each entry in the data set being turned into separate series of sounds. The series are then all played back to back.

3. EVALUATION

In order to determine if this is even a viable approach to sonifying exact data values, we have begun a series of systematic studies in which various parameters of the application are considered. Clearly this is an ongoing process, as a number of options can be modified. We present initial results of the first experiment where we asked participants to try to determine the prices of individual stocks from a fictitious data set.

3.1 Method

3.1.1 Participants

Thirty undergraduate participants (21 females and 9 males) took part in the study for course credit. The average age of participants was 19.6 years (SD = 1.3), and 20 had played a musical instrument regularly. Of those who had played a musical instrument, they had started playing at an average age of 9.6 years (SD = 2.7) and had been playing an average of 5.9 years (SD = 3.8), to include formal instruction (M = 4.9 years, SD = 3.0). Some of the instruments participants had played or were currently playing included piano, oboe, violin, french horn, and flute. All participants were right handed.




Figure 3. Percentage of successful trials sorted by absolute error value.

3.1.2 Materials / Apparatus

Participants were tested using a Dell Inspiron 8500 laptop computer, with a 2.20 GHz processor, 512 MB of RAM, and running Windows XP Professional with Java v1.4.2 installed. Stereo headphones (Sony MDR-7506) were used to play the tones.

3.1.3 Procedure

Participants were informed that they were taking part in an experiment designed to assess how well individuals could interpret simple sounds to deduce information about specific prices in the stock market. Subjects were first asked to calibrate the volume of the headphones, to ensure they could hear all the tones at a safe and comfortable level.

Next, each participant was taught how the sounds would represent information and allowed practice trials to ensure they understood how the study would operate. Subjects were told to first listen to 10 tones being played, representing the digits 0 to 9, with the added instruction that as pitch increased, so did the value of the digit the tone represented. Once the participant was confident with this concept, simple numbers such as 10, 30, 100, and 1000 were played. After the simple numbers were understood, the participant was taught how decimal points were represented, and then listened to a couple of numbers containing decimals. After completing the learning phase, each participant was allowed five practice trials, during which time they attempted to write down the number they thought the practice tones represented. Once being told the correct answer, a subject was allowed to listen the tones again to attempt to recalibrate their listening.

Finally, the participant moved into performing the task of interest. Five blocks of ten numbers each had been randomly generated prior to the experiment, and each subject would listen to the same 50 numbers. Each participant was instructed that the first four blocks would have either one, two, three, or four digits in front of the decimal point, with all trials in a block having the same number of digits.

However, the fifth block would contain numbers with one, two, three, or four digits in front of the decimal point. The experimenter played each tone only once, after which the participant wrote down the number that they felt the sounds represented. When ready to move on, the participant signaled the experimenter to move on to the next trial.

Once participants had completed all five blocks, any questions they had were addressed, the experimenter asked them what they thought of the study, and thanked them before they left.

3.2 Results

Initially all participants' responses were coded as either successful or unsuccessful based on whether their answer contained the same number of digits as the number presented by the Audio Abacus. The unsuccessful responses were then excluded. It is important to note that unsuccessful responses accounted for only around 10 percent of all responses across all blocks.

The data were then analyzed by comparing the actual number presented by the Abacus to participants' responses. The primary result of interest was the absolute error across all values. This error was calculated by subtracting each digit in the actual number from the corresponding digit in the participants' response and then taking the absolute value of this difference (e.g. participant's response=4, actual response=3; $3-4=-1$; absolute error=1).

Plotting absolute error reveals that 80 percent of participants' responses were within two of the actual value, and 90 percent of all participants' responses were within three (Figure 3).

Taken together, these data support the fact that even with only minimal training participants are able to determine the number of digits in the sonified number, and are able to discern the actual digits presented by the abacus with reasonable certainty.

4. GENERAL DISCUSSION

There are two key points the data described here make evident. The first is that users are able to successfully use an interface such as the Audio Abacus to perform a point estimation task, with only very limited training or introduction. This bodes well for the use of such an alternative display method in cases where exact data values are needed. Given the multitude of display options provided by the program, the data presented here barely scratches the surface of the myriad possible interface configurations for the Audio Abacus.

It is likely the case that different tasks and users will require unique combinations of settings. For example, blind science students who are experienced with the Audio Abacus may use a particular ratio setting and increase the overall playback speed. Casual users, such as customers in a grocery store, may need a slower speed and different pan settings. The software is set up to encourage experimentation by users and other researchers. It is exciting to look forward to a whole line of investigation aimed at uncovering what works best with this new interaction method. Other important avenues of investigation include the effect of practice and training and the effect of real world contexts on Abacus users' performance.

Of course, the auditory display of quantitative information is not only of interest to the visually impaired. Potential users range from school children learning math in a more dynamic and multimodal way, to fighter pilots who need to know their weapons status and altitude without interrupting the multiple speech channels they must already monitor. In recognition of this, our investigations usually include—and in this particular study began with—sighted participants. Certainly, however, it will be the next step in the present project to involve visually impaired and blind listeners to ensure that the concept, and the implementation details, of the Audio Abacus are appropriate and effective for all potential users. Our previous research has typically found that there are many similarities, but sometimes some subtle differences between the way sighted and blind listeners interpret what they hear in an auditory display [e.g., 20]. Understanding this, and making displays that are both useful to, and usable by a broad range of users is an effective design path that benefits all involved.

The second important point to be made is the general need for more research into sonification as it relates to point estimation tasks. There exists a large body of research that clearly demonstrates the efficacy of the use of sounds to convey information about data trends [e.g. 10, 11]. Indeed, there is a whole community of researchers in that area (International Community for Auditory Display, www.icad.org). However, the present work indicates that new approaches, “thinking outside the box” so to speak, have the potential to expand the range of effective sonification to include more nontraditional tasks. By no means a definitive statement, the data presented here are merely the tip of the research iceberg that is presently growing. Careful and creative ways of augmenting a person's interaction with the world can clearly reap major benefits, both in an academic and theoretical sense, and in a practical life-changing way.

5. REFERENCES

- [1] G. Kramer, B. N. Walker, T. Bonebright, P. Cook, J. Flowers, N. Miner, J. Neuhoff, R. Bargar, S. Barrass, J. Berger, G. Evreinov, W. T. Fitch, M. Gröhn, S. Handel, H. Kaper, H. Levkowitz, S. Lodha, B. Shinn-Cunningham, M. Simoni, and S. Tipei, "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," International Community for Auditory Display (ICAD), Santa Fe, NM 1999.
- [2] B. N. Walker and G. Kramer, "Mappings and metaphors in auditory displays: an experimental assessment," presented at 3rd International Conference on Auditory Display (ICAD96), Palo Alto, CA, 1996.
- [3] J. A. Gardner, R. Lundquist, and S. Sahyun, "TRIANGLE: A practical application of non-speech audio for imparting information," presented at International Conference on Auditory Display, Palo Alto, CA, 1996.
- [4] T. T. Hermann, P. Meinicke, and H. Ritter, "Principle curve sonification," presented at International Conference on Auditory Display, Atlanta, GA, 2000.
- [5] B. Sturm, "Sonification of particle systems via de Broglie's hypothesis," presented at International Conference on Auditory Display, Atlanta, GA, 2000.
- [6] J. H. Flowers, L. E. Whitwer, D. Grafel, C. and C. A. Kotan, "Sonification of daily weather records: Issues of perception, attention, and memory in design choices," presented at International Conference on Auditory Display, Espoo, Finland, 2001.
- [7] K. V. Nesbitt and S. Barrass, "Evaluation of a multimodal sonification and visualization of depth of market stock data.," presented at International Conference on Auditory Display, Kyoto, Japan, 2002.
- [8] N. E. Miner and T. P. Caudell, "Using wavelets to synthesize stochastic-based sounds for immersive virtual environments," presented at International Conference on Auditory Display, Palo Alto, CA, 1997.
- [9] H. Petrie and S. Morley, "The use of non-speech sounds in non-visual interfaces to the MS-Windows GUI for blind computer users," presented at International Conference on Auditory Display, Glasgow, UK, 1998.
- [10] J. H. Flowers and T. A. Hauer, "Musical versus visual graphs: Cross-modal equivalence in perception of time series data," *Human Factors*, vol. 37, pp. 553 - 569, 1995.
- [11] J. H. Flowers, D. C. Burham, and K. D. Turnage, "Cross-modal equivalence of visual and auditory scatterplots for exploring bivariate data samples," *Human Factors*, vol. 39, pp. 341 - 351, 1997.
- [12] T. L. Bonebright, M. A. Nees, T. T. Connerley, and G. R. McGain, "Testing the effectiveness of sonified graphs for education: A programmatic research project," presented at International Conference on Auditory Display, Espoo, Finland, 2001.
- [13] C. P. Peres and D. M. Lane, "Sonification of statistical graphs," presented at International Conference on Auditory Display, Boston, MA, 2003.

- [14] S. Barrass, "Auditory information design," Australian National University, 1997.
- [15] B. N. Walker, "Magnitude estimation of conceptual data dimensions for use in sonification," in Psychology Department. Houston, TX: Rice University, 2000.
- [16] D. R. Smith and B. N. Walker, "Tick-marks, axes, and label: The effects of adding context to auditory graphs," presented at International Conference on Auditory Display, Kyoto, Japan, 2002.
- [17] E. R. Tufte, *Envisioning Information*. Cheshire, CT: Graphics Press, 2001.
- [18] M. A. Quinones and A. Ehrenstein, "Psychological perspectives on training in organizations," in *Training for a Rapidly Changing Workplace: Applications of Psychological Research*, M. A. Quinones and A. Ehrenstein, Eds. Washington, DC: American Psychological Association, 1997.
- [19] R. Upson, "Educational sonification exercises: Pathways for mathematics and musical achievement," presented at International Conference on Auditory Display, Kyoto, Japan, 2002.
- [20] B. N. Walker and D. M. Lane, "Psychophysical scaling of sonification mappings: A comparison of sighted and visually impaired listeners," presented at International Conference on Auditory Display, Espoo, Finland, 2001.