Pie Chart Sonification

Keith M. Franklin and Jonathan C. Roberts
University of Kent, Computing laboratory,
Canterbury, England, UK.
{kmf1@kent.ac.uk, j.c.roberts@kent.ac.uk}

Abstract

Different acoustic variables such as pitch, volume, timbre and position can be used to represent quantitative, qualitative and categorical aspects of the information. Such sonifications are particularly useful for those with visual impairments; they are also beneficial in circumstances where visual representations would be impossible to use or to enrich a graphical realization. We demonstrate methods of representing an audible pie chart representation such that the hearer understands the information through an equivalent representation. We implement and evaluate five designs. In each the user is positioned at the center of the chart and perceives the information through positional sound sources.

Keywords--- **Sonification, Visualization, Charts**.

1. Introduction

Sound is very important in our every day life; often we use it without consciously knowing that we are. For example, when we walk down a meandering corridor we use sound made by the oncoming people to avoid collisions, alternatively at a music concert given by an amateur group we may particularly notice mistakes the amateurs have made, or if someone drops some coins then we may quickly perceive the value of the dropped coinage. Indeed, over the years sound has be used to represent more information-rich phenomenon, from representing errors by alarms (such as used in computer interfaces), through sorting algorithms [1] to more recently visualizing the web [2] or sonifying well-logs from oil and gas exploration [3]. Sonification may be readily used to allow the user to view patterns within the data. In particular, sound is especially important for people with visual impairments. Thus, it is important to find ways to represent information that is accessible to this community.

There are other reasons why charts and diagrams should be represented by sound, in addition to making the representations accessible to partial or non-sighted users. For example, there may be situations where the user cannot view a screen because they are monitoring something else (e.g. in a machine room where the engineer is constantly monitoring the material being cut and machined), or sonification may be useful to represent information where only very small screens are available. Non-visual visualizations, particularly pie-chart sonification tender a range of particular challenges such as how to represent the concepts of two-dimensional charts into sound space, how to effectively map the values to sound, whether users can actually perceive the information and, in particular, how accurate do users perceive the information.

In this paper, we present and evaluate some novel ways to represent pie charts using sound. Similarly to representing data by graphics (using the retinal variables of color, size, orientation, symbol etc) there are many sound attributes that can be used to represent the data. Thus, the challenge is how to use these variables to effectively map the data into audible parameters. We are particularly interested in non-speech mappings, certainly there are advantages in representing the information by speech, but such transformations loose spatial illustration and thus can miss-represent or allow the listener to miss-out on the richness of the underlying information. Also, some may argue that pie charts are not very good at effectively displaying the underlying data; but, pie charts are often used, appear in many texts and thus we believe should exist in an audible form.

Obviously, a pie chart made from sound cannot be identical to its visual counterpart, but inspiration can come from the visual pie chart itself. This idea of generating an equivalent representation is supported by current teaching methodologies; for example, blind and visually impaired users feel visual representations of diagrams using swell paper [4]. Using this method they develop a similar mental model to normal sighted users. Thus, by creating a sonic pie chart that is based on the visual representation, users will be able to draw on their previous experience and knowledge of pie charts.

2. Background

There are various papers that demonstrate chart sonification, for example, Ramloll [5] and Bonebright [6] represent line graphs and earlier work by Brown [1], when he represents the state of the sorting algorithms, could be classified as bar-chart sonification. However, we have to look at the haptic literature to find a non-visual pie chart representation. Indeed, there are a reasonable number of papers on visualizing charts through haptics. Many of the methods single point devices such as the Phantom ™ force feedback joystick [7]. For instance, Yu et al [8] developed haptic pie charts using a Logitech™ Wingman Force Feedback mouse which applies forces to mimic ridges round each pie-segment. In this example the user can feel these boundaries as they explore the sections of the chart, however, it is difficult to perceive and judge an angle from touch especially and again the user views the chart through a single point of contact. Other researchers have sonified more abstract data such as well-logs [3] and web browsing using WebSound [2].

Pie charts themselves display two components: a list of categories that have values. Each value is displayed as a percentage of the total, which are represented by angles of a circle. Thus, it is easy to see that many sonification designs could be inspired from this visual form, such as positioning sounds at the center of each pie-part, or representing the percentages by tone duration. Before we present our designs, we include some background information on both auditory displays and sound parameters.

2.1. Auditory Displays

We use two main technologies to represent the sound: headphones and speaker systems. Both systems can play back high fidelity sounds over a huge range of frequencies and volumes, but both have a limited number of physical sound sources. Thus, "virtual auditory displays" are used to recreate auditory signals such that the sounds are perceived from various locations. Most headphone-based systems use Head-Related Transfer Functions (HRTFs) to simulate the sound spatialization, other systems may use multiple speakers and vary the sound generated from each speaker to imitate sound from different locations [9]. The way we hear is very personable and depends on anatomical measurements, and so the HRTFs can be personalized [10].

Headphones have the advantage that the sounds do not disturb other people and provide an enclosed environment such that the user is not distracted by external noises. However, in certain circumstances headphones cannot be used and also some companies ban headphone usage for safety reasons as they inhibit interpersonal communication. In comparison, speaker systems allow larger groups of people to have the same experience and simultaneous communication is possible, but users may be distracted or miss-understand the sonification as a result of background noise. Thus, the developer needs to consider questions including: how noisy is the surrounding environment, how long they will be required to listen. For example, Helle et al [11] demonstrated that users turned off sounds (on mobile phone menus) as users found them inappropriate or annoying. Similarly, Sikora et al [12] noted that certain sounds (such as natural music) could be listened for longer periods than earcons or other mappings.

2.2. Sound Parameters

When developing sonification systems, the developer needs to consider the most effective way to map the data values into sound parameters. These parameters include pitch, volume/loudness, timbre, location, rhythm, duration and tempo [13]. Systems often utilize a single tone that is adapted to represent the values (e.g. [1,2,3]). However these approaches can have drawbacks, for instance, the relationship between the sound and the data is not necessarily intuitive, as noted by Hermann et al [14]. Moreover, Walker [15] (in one experiment) mapped dollars to pitch and observed that blind users perceived a negative relation whereas sighted users had a positive correlation to the monetary value.

Sound position is another parameter (incidentally omitted [13]). In fact, we use positional sound to represent aspects of the pie chart (see section 3). Although, positional sound is not often used in sonification research like [16] shows how the location of a sound can be successfully used to convey navigational information. However, the use of position, like any other perceptual variable, is not without its limitations. For example, in the experiments by Holland [16] users (when using headphones) found it difficult to distinguish between sounds placed in front and those from the same position behind. Moreover, it is difficult to accurately determine the position of the sounds. This error is known as the Minimum Audible Angle (MAA) [17,18], which represents the smallest angular separation of sound sources that is discernible by a listener. The actual value of the MAA, in fact, depends on the signal frequency and the orientation round the user. Figure 1 shows a chart of some MAA values for sounds at different pitches and at different locations around the head in the azimuth plane.

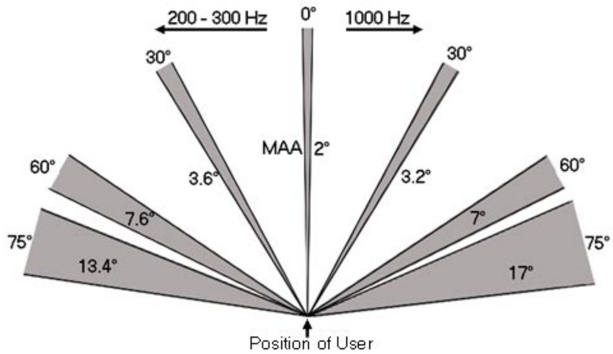

Figure 1 This diagram shows how the Minimum Audible Angle (as depicted by the inner values) changes depending on both the different frequencies used (left and right halves of the diagram) and the orientation around the user.

Moreover, just as a listener to music would perceive individual sounds as one whole, so the user perceives the sonification parameters holistically. Thus, although these parameters (as listed above) may represent individual values from the data it is quite difficult to perceive them independently. For instance, users may perceive a signal to be moving (due to an increase in volume), whereas in fact the position remains constant. This perceived dependency of the sounds might misrepresent the information. Such effects are understood in the visual domain; for instance, we perceive intensity (brightness) in relation to surrounding objects that can make objects appear lighter or darker than they are, an effect known as Simultaneous Brightness Contrast.

In addition to these simple parameters, the user may use more complex sounds or group the sounds together in motifs; these may be categorized as auditory icons, earcons, musical and verbal sounds. Auditory icons are real-life sounds, such as the sound of breaking glass or the sound of traffic going past [19], and are often used where their meaning can be implied. Earcons, in contrast, are rhythmic sequences of pitches that are unique and identifiable [20] that commonly has no implicit meaning. The user will need to learn their meaning but they are generally considered easier to listen to for a longer duration than auditory icons [12]. Musical sounds are short extracts from musical pieces such as a couple of seconds out of Beethoven's 5th symphony. These are good because the user is able to use musical sounds for longer than any other type of sound [12] and as there is no implied meaning associated: they can be mapped onto any function. Music has been used in the graphical interface for the visually impaired [21]. Verbal sounds are voice patterns that are either pre-recorded or dynamically generated. The meaning of the message is relatively unambiguous so no training is required (unless the language or accent is unfamiliar to the user), however, there are still problems in generating realistic voice patterns that are not too abrasive or distracting. In our designs we use simple orthogonal sound attributes to represent the pie chart information.

3. Different Designs and Implementation

There are many possible designs that could be invented to sonically represent a pie chart. We have developed five different representations. The concept behind our designs is to represent the information in a way that mimics a graphical pie chart. As the graphical pie chart is based on positional information so we use positional sounds in the sonified version. The user is placed at the center of the world and the sounds are generated around them. The orientation of the user is synchronized with the sound such that the user always faces towards 'zero percent' and sounds are emitted from nodes that are on the perimeter of the pie chart, which is now in the azimuth plane and at ear level (see Figure 2 for pictorial description). Thus, the user should be able to determine different 'virtual' sound sources and thus calculate a percentage value. Obviously, there are different ways to map the sound parameters within these constraints. One mapping could be to represent each pie-segment by a localized sound of a different timbre or pitch (each sound plays at the same time) the duration of which represents the percentage, however from early experiments it seemed difficult to perceive individual values; thus, in this paper we focus on methods that sound each pie-segment individually.

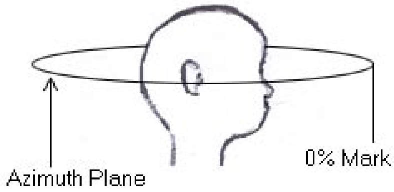

Figure 2 Pictorial description of the concept behind the first four representations, in that the user's head is placed at the center of all the sound sources in the azimuth plane.

Design 1

In this design we aim to signify the angle of the pie chart itself. Two sound sources are used to signify the start and end of the pie segment, each sound is played subsequent to the other. A small amount of silence is added before the next segment is sonified. It is important to note that the starting point of the next piece of the pie chart will be the ending point of the previous section. This adds a small amount of redundancy, which the user can use to confirm the size of the previous segment. The process of adding a gap and playing the next area of the information is repeated until all sections of the chart have been represented. For example, using this design to represent the pie chart shown in Figure 3, virtual speakers would be positioned at the locations 1-5. The user would hear sound from the virtual speakers in the following order: 1 then 2, silence, 2, 3, through to 5 then 1.

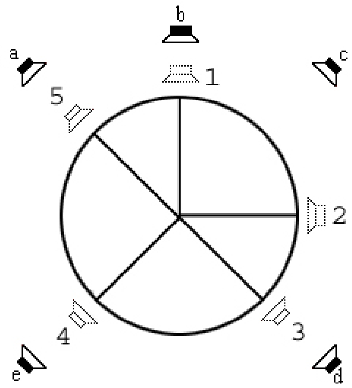

Figure 3 This diagram depicts a pie chart represented by virtual sound sources (indicated by '1 to 5'); the location of the physical speakers are indicated by 'a to e'.

This representation has been developed using VRML. Sounds are emitted from a series of sound nodes positioned along the perimeter at the boundary between each section of the pie chart. In this case simple tones are played to the user in the sequence as described above. Additionally, the categories could be represented using different pitches, however, the MAA error changes depending on pitch and thus there is a limit to the number of pitches that can be effectively used (see Figure 1). In this example we use used a multiple source speaker system (Creative Labs™ 5.1 Cambridge Soundworks).

Design 2

The second demonstration is a variation of the first; however in addition to playing sounds at the boundaries we also sound out intermediate (tick) marks between the start and end boundaries. This gives the user a clearer indication of movement around the circumference of the pie chart. This is a direct analogy with tick marks on the axis of a line graph. The inclusion of these extra sounds provides subsidiary information that can aid the user to better evaluate the size of the portion, which is especially beneficial when the sounds are emitted perpendicular to the user (at large MAA positions). This would also allow a larger range of pitches to be used to specify different categories. The idea of adding intermediate information is similar to the work of Rigas [21], who demonstrated that by giving intermediary information the user was able to answer questions about the presented information with a greater accuracy.

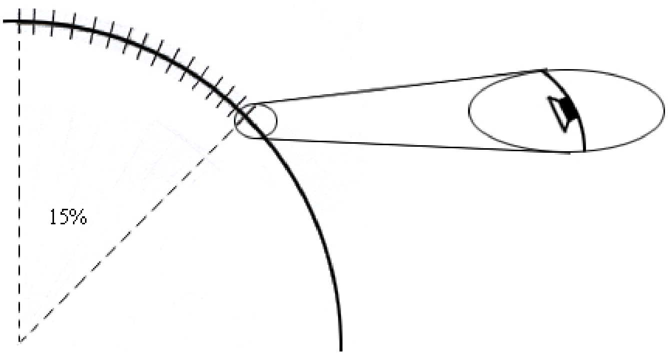

Figure 4 A section of a pie chart is shown with the location of each sound emitting nodes for that section, marked by a line crossing the perimeter line.

An example of part of this system is shown in Figure 4; in this case the size of the section is 15%. The user would first hear a sound coming directly from their right hand side followed by a series of sounds moving clockwise around the perimeter till it reaches the end boundary. The user will hear a total of 16 sounds, because a sound is played for each start and end boundary. This design was implemented in the same way as the first, with simple tones, in a VRML world and with a surround sound speaker system.

Design 3

The third model normalizes all the starting points so that each section of the pie is presented with the same reference point (directly in front of the user). This is advantageous; because the MAA is lowered for all small sections (being displayed in an area, which the user has a higher accuracy in estimating positions), thus different segments of the pie chart should also be easier to compare. Figure 5 shows how segments of the pie chart (depicted on the left) are sonified using the three virtual speakers (shown on the right). In this example, the user would hear sounds from the speakers as follows: 1 then 2, a short silence, 1 then 3, another short silence and finally 1 then 2.

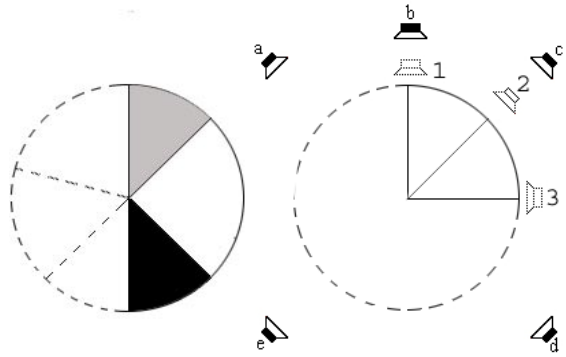

Figure 5 This diagram depicts design 3 and shows how the first three segments of the pie chart (left) are normalized so that the sound always starts at position 1. Additional virtual speakers are positioned to indicate the extent of the segments.

Design 4

This fourth representation is a combination of design 2 and 3. It combines the methods of sounding the 'tick' marks along with the new exploration technique of normalizing the start points. Thus, the user should benefit from both the additional information, when the angle between the boundaries is smaller than the MAA, and the two advantages of normalizing the start points.

Design 5

The final design takes additional inspiration from Morse code, where a combination of long and short beeps, known as dashes and dots respectively, construct individual letters. A value of 10% is indicated by a dash, whilst a dot symbolizes a value of 1%. Thus these two values can now represent any whole percentage value within a pie chart, for example if the section being displayed was 43% then the user would hear four dashes followed by three dots (see Figure 6 for full example). In addition to the duration of the sound we also change the pitch of these sounds to distinguish between each segment. This method allows users to calculate quantifiable values, and thus is advantageous over other diagrammatic sonifications such as [6,22,23].

This design was implemented in Java. The program looked at each section of the pie chart individually and breaks the size of the section into its smallest number of dashes and dots. When playing the sounds that represented a section, the program presented all of the dashes for a particular section followed by the dots. After this a short gap is inserted before the next segment is presented. We've used two different simple tones for the dashes and dots and because we were not using the position variable, this model can be presented on either headphones or speakers.

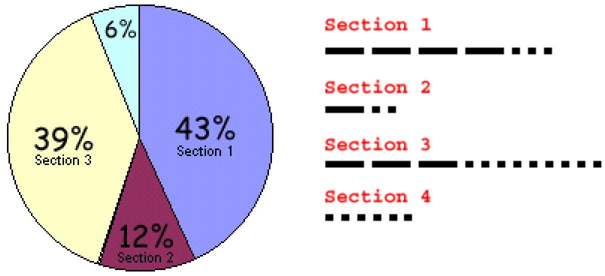

Figure 6 An example pie chart with the dashes and dots for each section listed on the right

4. Evaluation

We have evaluated each design using a group of volunteers. Each participant listened to a sonification, which was repeated as many times as required, they then attempted to match their understanding of the sonified version to a graphical pie chart (from a choice of 15 different charts). In order that the users did not get tired of the identification task, and to normalize the results, the tested three out of the five designs, which were selected at random.

Each volunteer was presented with four randomly selected pie charts for each design, which they had to identify. They were given a short training session before being tested where they heard different pie charts. Additionally, the test-subjects were both observed and questioned about different aspects of the design, such as how long did they feel that they could listen to this type of representation before becoming unproductive, and whether there were any parts of the sonification that were confusing. The results from this task were used to determine the usability of these designs.

5. Results and Discussion

From our experiments we observe that the users could understand the information, but the accuracy of the different models varied greatly. From Figure 7 we can see that design 5 was the most accurate and design 3 the least; however, the effectiveness and usability of the design should not be solely based on accuracy but other factors, such as the number of times the sonification was replayed (until the user believed they understood the sonification). Figure 8 shows averages of how many times the users replayed the five different models.

Based on these results, the user feedback and our observations we consider that Design 5 (Morse code variant) was the most effective representation, with 100% accuracy and a low average of replays. Design 2 (non-normalized with intermediary point information) is the second most usable since its average number of replays was lower than the other designs; and from a closer look at the results it seems that the lower accuracy was due to a miss-ordering of some of the segments, this could indicate a problem with the user recalling the visual diagram rather than misunderstanding the sonification itself. Design 4 was the next usable followed by 1 and then 3. Based on the current results designs 1 and 3 (indicating start and end points only) seem to be ineffective for accurate usable pie chart sonification.

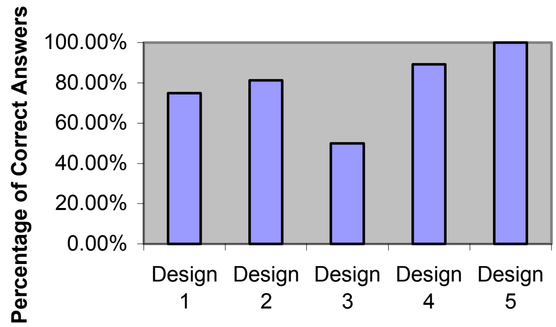

Figure 7 Chart showing the average accuracy of each of the designs

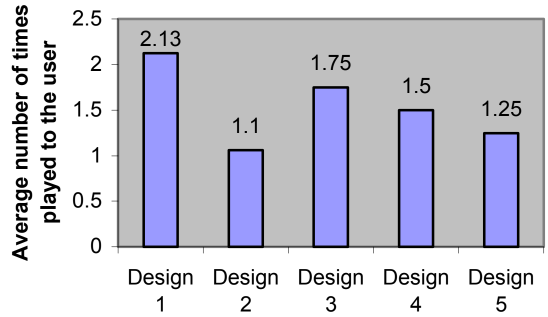

Figure 8 Chart showing the number of times that the sonification was played to the users for each design

Moreover, we originally believed that the first four models would have a higher accuracy than the results depict and that these (position based) designs would be better than the fifth Morse code version. However this lower accuracy might be a result of how the VRML browser implements sound spatialization. Furthermore, VRML sometimes creates an extra noise (a clipping sound) when changing from one sound source to another. We are continuing evaluations of these representations, especially to see if this is a limitation of VRML, and will investigate other software environments such as OpenAL.

6. Conclusions and Future Work

We conclude that usable pie charts sonifications are possible and can be created via different mappings. We also propose (like Rigas [21]) that by sonifying additional intermediary points (similar to tick marks on an axis) the user can more accurately identify the value.

Although, we initially believed that normalizing the starting points would lead to higher accuracy (design 3 and 4) our results refute this notion. Thus, from this we conclude that maybe the reference point (the starting point) should be positioned in an area where localization is poor (e.g. at the North West position) rather than the within the most sensitive area (that is straight ahead).

Finally, Design 5 (the Morse code version) could be useful to depict other chart types, such as bar charts. This also has the benefit that it will not require the user to have perfect pitch, to estimate the pitch to estimate a quantitative value, like current methods do.

We plan to further expand this research by (1) enriching the sound representation, so that more information is presented about the structure of the data within the same time. (2) Investigating new exploration techniques so that users can select different sections of the pie chart that they wish to listen to and receive this information in a parallel format. (3) Looking at possible ways to decrease the time taken to represent the data in the current serial format. (4) Integrating the technique used in design 5 into the other representations. (5) Investigate the usability and usefulness of these new systems with some context such as using real life data.

Acknowledgements

We would like to thank the many people that took part in the testing of these representations, as well as Clive Miller from the Royal National Institute of the Blind, for his feedback on our work.

References

1. Marc H. Brown and John Hershberger, Color and Sound in Algorithm Animation, IEEE Computer, volume 25, number 12, pages 52 – 63, 1992
2. Lori Stefano Petrucci, Eric Harth, Patrick Roth, André Assimacopoulos, Thierry Pun. WebSound: a Generic Web Sonification Tool, and its Application to an Auditory Web Browser for Blind and Visually Impaired Users. ICAD 2000, Atlanta, April, 2000
3. Barrass S. and Zehner B. (2000) Responsive Sonification of Well-logs, Proceedings of the International Conference on Auditory Display ICAD 2000, Atlanta, April 2000
4. Thomas Way, Automatic visual to tactile transalation, part1: Human factors, access methods and image manipulation, 1996
5. Rameshsharma Ramloll, Wai Yu and Stephen Brewster, Constructing Sonified Haptic Line Graphs for the Blind Student: First Steps, Proceedings of ACM Assets, Arlington, 2000
6. Testing The Effectiveness of Sonified Graphs For Education: A Programmatic Research Project, Proceedings of the International Conference on Auditory Displays, Finland, July, 2001
7. Jonathan C Roberts, Keith M Franklin and Jonathan Cullinane, Virtual Haptic Exploration Visualization of Line Graphs and Charts, Proceedings of The Engineering Reality of Virtual Reality, SPIE volume 4660B, pages 401 – 410, San Jose, January, 2002
8. Wai Yu, Douglas Reid and Stephen Brewster, Web-Based Multimodel Graphs for Visually Impaired People, In 1st Cambridge Workshop on Universal Access and Assistive Technology (CWUAAT), Cambridge, UK, 2002
9. Shinn-Cunningham, BG (1998). “Applications of virtual auditory displays,” 20th Ann Conf IEEE Eng Med Biol Soc, Hong Kong, China, pages 1105-1108
10. D. Zotkin, R. Duraiswami and L.S. Davis, "Rendering Localized Spatial Audio in a VirtualAuditory Space," University of Maryland, Computer Science, Technical report. CS-TR-4348. 2002
11. Seppo Helle, Gregory Laplatre, Juha Marila and Pauli Laine, Menu Sonification In A Mobile Phone – A Prototype Study, Proceedings of the International Conference on Auditory Display, Finland, July, 2001
12. Cynthia A. Sikora, Linda Roberts and La Tondra Murray, Musical vs. Real World Feedback Signals, CHI ’95, May, 1995
13. Tara M. Madhyastha and Daniel A. Reed, Data Sonification: Do You See What I Hear?, IEEE Software, volume 12, number 2, pages 45 – 56, 1995
14. T. Hermann and H. Ritter. Listen to your Data: Model-Based Sonification for Data Analysis. In M. R. Syed, editor, Advances in intelligent computing and multimedia systems. Int. for Advanced Studies in System Research and Cybernetics, 1999.
15. Bruce N Walker, David M Lane, Psychophysical Scaling of Sonification Mappings: A Comparision*(sic)* of Visually Impaired and Sighted Listeners, Proceedings of the International Conference on Auditory Display, Finland, July, 2001
16. Simon Holland, David R Mouse, Audio GPS: spatial audio in a minimal attention interface, Proceedings of the Third International Workshop on Human Computer Interaction with Mobile Devices, Lille, France, September, 2001
17. Brian C. J. Moore, An Introduction to Psychology of Hearing, 4th Edition, Academic Press, 1997
18. A W Mills, On The Minimum Audible Angle, The Journal Of The Acoustical Society of America, volume 30, pages 237 – 246, 1958
19. William W Gaver, The Sonicfinder: An Interface that Uses Auditory Icons, Proceedings of Human Computer Interaction, volume 4, number 1, pages 67 – 94, 1989
20. M. Blattner, D. Sumikawa and R Greenberg, Earcons and icons: Their structure and common design principles, Human Computer Interaction, pages 11 – 44, 1989
21. Dimirios I. Rigas and James L Alty, The Use of Music in a Graphical Interface for the Visually Impaired, IFIP TC13 International Conference on Human Computer Interaction (Interact ’97), pages 228 – 235, 1997
22. Bruce N Walker, Gregory Kramer and David M Lane, Psychophysical Scaling of Sonification Mappings, Proceedings of the International Conference on Auditory Displays, 2000
23. Michael Gilfix and Prof. Alva Couch, Peep (The Network Auralizer): Monitoring Your Network With Sound, Proceedings of USENIX, pages 109 – 117, December, 2000