MUSICAL SONIFICATION DESIGN

A Thesis

Submitted to the Faculty

in partial fulfillment of the requirements for the

degree of

Master of Arts

in

ELECTRO-ACOUSTIC MUSIC

by

Edward Childs

DARTMOUTH COLLEGE

Hanover, New Hampshire

May 27, 2003

Examining Committee:

(Chair) Larry Polansky

Petr Janata

Eric Lyon

Carol Folt
Dean of Graduate Studies

Abstract

A working definition of sonification, together with its history, purpose, usefulness and relationship to musical composition are presented. Assessment of the potential of sonification is made based on music cognition and brain science research. Guidelines for musical sonification design are suggested, as informed by analogy with graphical design, music cognition and musical composition. Several case studies are provided, both of musical compositions which draw heavily on the sonification of data as a working procedure, and of sonifications of complex numerical processes, historical meteorological data and financial data which are designed to provide insight into these data relationships, but are not necessarily musical compositions. As an exercise, musical compositions were prepared from all of these sonification projects.

Acknowledgments

The author expresses thanks to Professor Jon Appleton, Director of the Electro-Acoustic Graduate Program at Dartmouth, for his insight and guidance throughout the project. Professor Eric Lyon introduced me to the field of sonification and provided encouragement and hand-holding during my early efforts with electro-acoustic music composition, as well as always providing insight and inspiration. Professor Larry Polansky provided generous and insightful advice, especially during the write-up stage of this work. Professor Charles Dodge listened to all my sonifications, gave useful feedback and took a great interest in the project. Dr. Petr Janata gave extensive guidance, references and great advice on the implications of music cognition on sonification design. The assistance of the Graduate Office, the John Sloan Dickey Foundation for International Understanding, and the Bregman Electronic Music Studio are gratefully acknowledged, in the funding of a visit to the Helsinki University of Technology in Otaniemi, Finland, in April and May of 2002, and travel to the ICAD 2002 conference in Kyoto, Japan, during July of 2002. Thanks are also due to Professor Matti Karjalainen and Dr. Ville Pulkki for their generous hospitality and assistance during my one month visit to Finland. My son, Benjamin Childs, gave generously of his time and advice on the numerous occasions I tied myself up in knots over computational dilemmas. Alexandra Holz provided expert editorial assistance.

This work is dedicated to all, both family and friends, who provided support, assistance, inspiration and love over my three year journey at Dartmouth.

Contents

List of Tables

List of Figures

1 Introduction and History

Sonification (which will be defined below) is an interdisciplinary topic spanning the arts (music, graphic design), the social sciences (psychology) and the natural sciences (engineering and computer science). This thesis is concerned with musical sonification design, that is, sonification in which the procedures of musical composition are integral to the design process. The above-mentioned disciplines will be "visited," insofar as they are relevant to the topic.

In Chapter 1, sonification will be defined, together with its context, and a brief history will be provided. The field of Auditory Display will be introduced, and some practical applications will be described. While the word sonification, in the context of Auditory Display has only been in use for about ten years, the underlying concepts have been used by composers since the time of Pythagoras and possibly earlier. It is therefore entirely appropriate to make detailed reference to musical composition when discussing sonification.

In Chapter 2, some theoretical aspects of musical sonification design will be presented, drawing upon the disciplines of music composition, music cognition and (by analogy) graphic design. It will be argued that music is a subset of sonification.

In Chapter 3, the composition *Achorripsis* by Iannis Xenakis will be analyzed in depth as a sonification of probability distributions [1]. The purpose will be to explore the boundary between artistic choice and determinism, and the relevance of his artistic choices to sonification design. Several other compositions will be briefly described.

In Chapter 4, some sonifications of other researchers will be discussed, technique and methodology will be addressed, followed by a presentation of several of my sonification projects involving computational fluid dynamics data [2], meteorological data [3] and financial data. Several musical compositions arising from these sonifications will be described.

In Chapter 5, some works in progress will be discussed and some thoughts regarding the future of sonification research will be presented.

1.1 What Is Sonification?

Webster's Unabridged Dictionary [4] defines sonification as "Act or process of producing sound, as the stridulation of insects." Stridulation is in turn defined as "To make a shrill or creaking noise, or a musical sound". Insects such as katydids, crickets, grasshoppers or cicadas stridulate by scraping together file-like structures on their wings. The word is generally considered to be fairly obscure, and, for example, only appears in the Unabridged Edition of Webster's.

In the past ten years (and dating, perhaps, even earlier), the word has acquired a new meaning in the realms of computer music, computer science, and most precisely in the emerging field of Auditory Display.

For example, sonification has been defined as:

> A mapping of numerically represented relations in some domain under study to relations in an acoustic domain for the purposes of interpreting, understanding, or communicating relations in the domain under study. *Carla Scaletti* [5]

As part of a 1997 report to the National Science Foundation, members of the International Community for Auditory Display defined sonification as:

> The transformation of data relations into perceived relations in an acoustic signal for the purposes of facilitating communication or interpretation. [6]

Sonification involves the study of relations in a data set and their transformation into sonic relations that may be perceived by a listener. The framework of these sound relations is often intended to be musical, or derived primarily from musical ideas. That is, the most well known and easily accessible system for the organization of sound, and the articulation of relationships among sounds, besides speech, is music. Music encompasses a wide variety of sound relationships: large and small scale temporal relationships, melodic, motivic, harmonic, polyphonic, spatial, to name a few. Although obviously the perception of music, even the definition of music itself, is to a large extent culturally determined and highly subjective, it does constitute what might be called an abstract language. In that way, it can lend itself to sonification.

The essential task of sonification, sometimes known as the "mapping problem," is to design the transformation such that the relations in the data are manifested in corresponding musical (sound) relations.

Once a mapping has been determined, sonifications are usually accomplished by acquiring the data in digital format, via a computer, and routing that data to a "sonification engine" which generates the musical sounds. Recording data from a digital thermometer, and mapping its temperature to the pitch of a computer-synthesized tone, would be one very simple example of a sonification (and one in which the relationships are very simple).

An intriguing aspect of sonification is the possibility of exposing patterns or trends, in data, through sound, which were not previously perceivable through visual or analytical examination. This, in fact, is one of the most interesting motivations for exploring this field. Two possibilities might exist for this realization:

1. Musical structures are intrinsic to the data.

2. The "mapping" scheme opens a musical window into data which has no apparent musicality.

In the first case, the data, when sonified in a straightforward manner, would yield musical results with little intervention or manipulation on the part of the designer. For example, some experiments with chaotic attractors yield interesting, and quite musical sounds [7]. In the second, the designer would "find" a mapping which would open or unlock relationships in the data for the listener.

1.1.1 A Brief History

Excellent in depth accounts of the history of sonification have been provided by Kramer [8] and [6]. Early examples which predate the use of the term "sonification" include the "Geiger counter, sonar, the auditory thermometer and numerous medical and cockpit auditory displays" [6]. In the early days of computing, a cheap AM radio placed on the mainframe near the CPU emitted a characteristic sound (presumably caused by electromagnetic interference) if the machine was stuck in an infinite loop. In the early days of garbage collection on a DEC PDP-6, in order to monitor CPU demand, the garbage collection register was connected to a digital-to-analog converter (DAC), which emitted bursts of noise to indicate activity [8].

While not strictly speaking a sonification, Morse code is a system of mapping the letters of the alphabet to rhythmic beeps of fixed pitch sinusoidal tones: dots (short) and dashes (long). Morse code operators who became proficient at decoding learned to pick out common words, and even sentences, from the recognition of larger scale rhythmic patterns. In a sense, the very simple but consistent mapping of letters to musical rhythms laid the ground work for the perception of more complex constructs from those basic building blocks. The Morse code experience illustrates the enormous human potential for non-speech auditory pattern recognition.

There are some devices which by nature of their construction, emit sounds in the audible range. While the sounds are not intentional, or designed, they are sometimes highly useful for monitoring and troubleshooting. The most common example would be the automobile engine. Most drivers establish a baseline, familiar soundscape associated with their car. If that baseline is altered or interrupted, even in a subtle way, the attention of the driver is engaged. Sometimes this type of sound monitoring is far more effective than the use of visual gauges or other devices. The sounds made by industrial equipment are frequently recorded or amplified since they provide a more effective, earlier warning of potential problems. In fact, humans are hypersensitive to environmental sonic deviations: “Things that go bump in the night.”

During a recent visit to a pizza place, while waiting for my credit card transaction to be completed, the clerk noticed without looking at the printer that the transaction had not gone through properly. He commented that credit card transactions had a certain sound which could be distinguished from other printing actions, such as the placement of a take-out order.

Experience of these self-sonifying devices point to the potential for adding sound to otherwise silent, or entirely visual processes. Early experiments with sonification arose from several technological advances or limitations:

1. Personal computers which could generate sophisticated sounds (1980s).

2. Significant gains in the speed of information flow (i.e. the Internet), and increased access to computational (number-crunching) resources.

3. Limitations in the human perception of highly complex, or simply voluminous data by visual or analytical means.

4. Emerging research in auditory cognition, which suggested a significant potential for sonification.

Early work in sonification was initiated at the Santa Fe Institute in the late 1980s by Gregory Kramer as part of their “Science of Complexity” research program. The purpose of the program was to explore the potential of auditory display as a tool for analyzing or finding interesting patterns or information in large data sets. This exploratory work was not unlike early computer music experiments with the Julia set and other fractals, the game of life and other math processes. In the early stages of his work, Kramer expected to find significant work, or at least a research community. Finding none, he convened the first auditory display conference in 1992 [8]. The International Community for Auditory Display, a not-for-profit corporation, was founded in 1996, and holds annual conferences every summer. Contributors to the conference typically include experimental psychologists, acousticians, computer scientists, engineers, physicists, artists and musicians.

1.2 Sonification as a Musical and Compositional Tool

The idea of using numerical relationships not originating in music to shape musical composition is undoubtedly ancient. Even the simple canon and more notably, the mensuration canon have been viewed as mathematical constructs. `http://www.humboldt.edu/~pzc1/mens/`. Significant musicological research has been devoted to the use of Fibonacci numbers and the Golden Mean in the works of Bach, Beethoven, et al. `http://www.mcs.surrey.ac.uk/Personal/R.Knott/Fibonacci/fibInArt.html#music`. It is widely thought that composers such as Bartok and Debussy made conscious use of the Golden Mean ratio in their music. According to Dr. Sever Tipei, “Bartok intentionally writes melodies which contain only intervals whose sizes can be expressed in Fibonacci numbers of semitones." Xenakis made use of the ratio in his first major work *Metastasis*. `http://web.hep.uiuc.edu/home/karliner/golden.html`.

Music of the past fifty years has seen an explosion in the use of many kinds of data and numerical systems in composition. Algorithmic music is a significant category in this genre. An algorithm is a procedure, mathematical or otherwise, that is iterative or repetitive in nature and usually works towards a specific goal. For example, an algorithm for removing all the firewood from my garage would be:

1. Place logs on wheelbarrow.
2. Move wheelbarrow to desired new location for logs.
3. Dump logs.
4. Return wheelbarrow to garage.
5. Repeat from (1.) until all logs have been removed from garage.

A good example of algorithmic music is *Chord Catalog* by Tom Johnson, in which all of 8178 possible chords in one octave are played in succession, starting with 2 note chords, over the course of approximately two hours. In *Rational Melodies* [9], and *Self-Similar Melodies* [10], Johnson maps numerical and (sometimes enunciated) enumerative processes directly to music with as little translation or interpretation as possible.

Algorithms are often associated with computer programming, since the computer excels at performing repetitive or iterative tasks. An example is the tri-diagonal matrix algorithm, which is a system for inverting (solving) a matrix which has specific properties. Since algorithms are executed over time, they lend themselves well to use in composition. The use of algorithms in computer-generated music is sometimes described as "computer music." Computer music has typically used random number generators, probability distributions, Markov chains, cellular automata such as "The Game of Life," chaos theory, fractals, Mandelbrot sets `http://www.ibiblio.org/emusic-l/back-issues/vol058/issue01.txt`, to name a few.

Another related genre is to take data from some non-musical source, e.g. star charts, as used in *Atlas Eclipticalis*, an early orchestral work by John Cage, and transform it into musical notes. Cage mapped the horizontal distance between stars to time, and the vertical distance to pitch.

It does not necessarily follow that the use of algorithms, or data in musical composition, is sonification. If we are to follow the definitions in Section 1.1 the following points must be addressed:

- Relations must exist in the data or numerical procedure.
- These must be mapped so that the relations can be perceived by a listener.
- The intent is to communicate something about, or interpret something in the data set.

Sonification is more than using numbers as a compositional tool to generate notes. The "mapping" of data, or numbers to sound is designed so as to express, communicate, or interpret, in the acoustical domain, the relations which already exist in the data domain. The primary mode of accomplishing this is to draw on some existing knowledge of the relations, or organization, that characterizes music. Most music reflects some level of concept, decision, or organization (even granted that it should be as inclusive as possible), and in designing a sonification, the intent is to draw upon different techniques of forming musical concepts, making musical/artistic decisions, and organizing the various materials, whether they be ambient noise, computer-synthesized sounds, or notated music performed by live musicians. For example, if the data contains a relationship of "distance," the sonification should also communicate "distance" to the listener (and there may be a variety of ways to do this – harmonic, intervalic, registral, timbral, spatial, dynamic distance, etc.).

1.3 Practical Uses of Sonification in Auditory Display

Sonification forms part of the more general field of Auditory Display which is defined as "the use of non-speech sound to convey information."

The field is formally divided into "Auditory Icons," "Audification" and "Sonification".

1.3.1 Auditory Icons

The use of different sounds in computing applications to communicate "success," "failure," "error," "You've got mail" and many other conditions is now commonplace. The development of auditory icons (sometimes called earcons) for the vision-impaired originated in the 1980s. Gaver [11] introduced a **SonicFinder** for Macintosh, in which samples of tapping, sliding and scraping noises were used in conjunction with actions taken by the user on the Macintosh Finder, such as opening and closing folders, or dragging and dropping. The **SonicFinder** was never incorporated into the Macintosh operating system because of memory constraints. More recently, mobile phones manufacturers [12] have used sound for the navigation of menus, due to the small size of the screens as well as safety issues. Usually, auditory icons are fixed soundbites which are always the same when invoked, and are not altered by instantaneous data values. Their association with an event is often metaphorical (as in the "knock-knock" sound when a "buddy" appears on-line in an instant messaging application). Automobiles also use “sound messages,” e.g. a beeping bell if keys are left in the ignition or the lights are left on.

1.3.2 Audification

Audification has been termed “0th Order Sonification” by Carla Scaletti [5]. In audification, data is converted directly to sound without mapping or translation of any kind. Data for many applications is recorded in digital format, which consists of a file, or stream, of (typically) short (i.e., 16 bit) integers, or “samples,” with some indication of the number of samples per second, or “sampling rate” at the head of the file. The sampling rate for audio files (i.e., compact disks) is usually 44.1 kHz. The appropriate sampling rate for a given application, following Nyquist, is usually taken to be twice the value of the highest frequency that needs to be captured (typically $\approx$ 20 kHz for audio). For other applications, data is often collected at much lower sampling rates, but can be “played back” at higher rates so that the characteristic frequencies fall within the audio range. The idea of audification is simply to convert existing data into sample (integer) format, and then either play the result back at a sampling rate suitable for hearing, or loop the data using wave table synthesis. The technique is most successful when the data contains significant periodic components, otherwise the audification will sound like noise.

Audification has been successfully applied to exploration seismology and planetary seismology (volcanic eruptions, earthquakes, large explosions) [13]. Dombois presented some fascinating results at the 7th and 8th International Conference on Auditory Display, [14] [15]. Dombois took data in the frequency range of 2.77×10^{-4} Hz (1 cycle per hour) to 40 Hz (a range of 17 octaves), and sped up the data into the audio range. On the one hand, he wanted to maximise the speed so as to be able to listen to weeks and months in a matter of minutes. On the other, he needed to retain the characteristics or “signature” of the seismological event. After experimentation, he ended up with a speed-up factor of 2,200 (allowing an hour of data to be heard in 1.6 seconds). In this case, the data is not really “mapped” or “transformed.”

A related area is physical modeling synthesis, in which the data obtained (a time-varying pressure field) is directly related to the sound. In physical modeling synthesis, the sound-generating mechanism of musical instruments, or other devices, is calculated from “first principles,” i.e., from the governing equations of physics. The result of these calculations is then “rendered” into sound. This rendering process does not involve mapping or transformation. The technique has been carried out to the point of simulating the varying pressure field in air in organ pipes and in the flute, using a technique called “computational fluid dynamics” (CFD). McCabe and Rangwalla [16] presented two examples of auditory display: the simulation of an artificial heart pump and rotor-stator interaction in turbomachinery. In the first, MIDI sound was used to enhance the post-processing of the artificial heart simulation, in particular to signal global changes in the system such as the opening and closing of a heart valve. In the second, the time-varying pressure field predicted by the model was rendered directly into sound. This was accomplished by tracking the time-varying normalized pressure at specific points in the flow field (in a sense, inserting microphones at these points) and scaling the pressure variation to short, signed integers, which then formed a wave table, see Figure 1. The simulated sound was then compared with known characteristics of the actual sound, and conclusions then drawn about the validity and accuracy of the computational fluid dynamics model.

1.3.3 Sonification

The field of sonification is relatively new, but has rapidly gained importance in recent years for a number of important reasons, including:

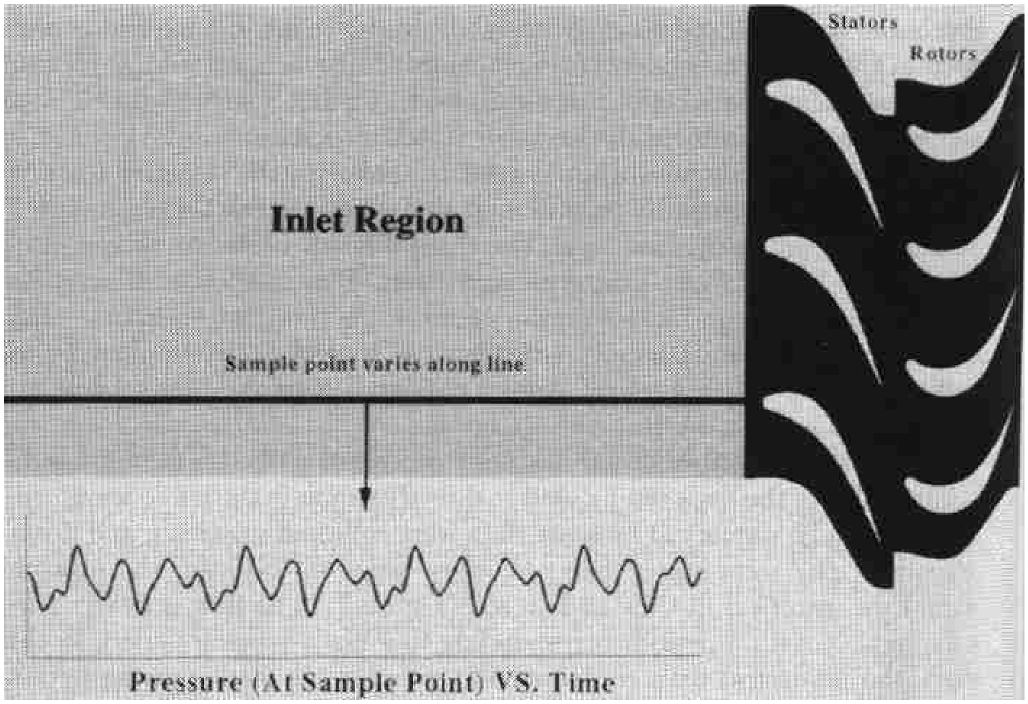

Figure 1: *Rotor-Stator Interaction* [16]

1. The overuse of visual displays in real-time data-intensive high stress situations.

2. The availability of real-time sound generation capabilities on standard desktop, laptop and mobile computing systems.

3. The body of psychological evidence, which suggests that the human ear is superior to, or more appropriate than the eye for processing certain kinds of information (especially "live" information which changes over time).

For centuries, printed visual displays have been used for displaying information in the form of bar charts, pie charts and graphs [17]. Visual displays on computers have been developed and used to a high degree. This reliance on visual display has forced the use of several screens in data intensive situations, in an attempt to display as much information simultaneously as possible. The user of such displays is often required to monitor a condition or number which constantly changes with time. Visual monitoring is tiring and stressful because it requires staring at a fixed place for long periods of time. The use of sound in computers has become increasingly sophisticated in recent years and it is now common for users to listen to high quality music over their computer audio systems. Processor speed is now sufficiently fast that musical sounds can be synthesized on standard computer hardware, using either digital audio synthesis or using MIDI. Based on these fairly recent capabilities, it is now possible to convert any kind of data (financial, medical, scientific, etc.), into a sound stream, using a technique called sonification. Sonification is the science and art of taking a data stream which has some significance (such as the real-time value of the Dow Jones Industrial Average) and mapping that data to musical sound so as to convey information to the listener. As a simple example, the value of the DJIA could be mapped to pitch and played on synthetic flute in such a way that if the index goes up or down, the "melody" rises or falls in pitch. Market volume could be used to control the loudness of the notes, or the amount of reverberation. In general, data can be mapped to musical pitch, loudness, rhythm, articulation, tone and other sonic parameters. It is thus possible to convey a substantial amount of information in sound. Several instruments of different tones, articulations (smooth or sharp) and spatial locations (stereo left-right) can be played simultaneously to display, e.g. several financial indices at the same time. Research in two related disciplines of music cognition (psychology) `http://www.music-cog.ohio-state.edu/Resources/resources.html` and human factors engineering has suggested that the use of sonification in the auditory display of real-time data may be far more effective than visual display. Music cognition research and our experience has shown that:

1. People become familiar with musical pieces through a process of passive/implicit learning (i.e. background listening).

2. Once the learning process has taken place, recognition of the same musical piece when it is played is very quick.

3. The style, composer or artist of similar musical pieces can be accurately guessed even if the pieces are being heard for the first time.

4. Once a piece or style is learned, the listener holds strong expectations when the music is played again. If something unexpected occurs (like a wrong note), this expectation violation provokes a strong reaction.

These research results have the following implications for the sonification of real-time data:

1. Monitoring data through sonification is feasible as a background task, while the user interacts with visual data, much in the same way that a person can listen to music in the background while working productively on other tasks.

2. Once the user is familiar with the sonification, patterns in the data (through their representation in a music-like sound) will be recognized quickly.

3. Similarities between patterns will be readily recognized.

4. Expectations of typical behavior will be formed, and any expectation violation will instantly attract the users attention.

The implication is that auditory monitoring of real-time data is likely to be less stressful, and possibly more effective, than visual monitoring. Auditory monitoring can be somewhat passive (like listening to background music), but unexpected or significant behavior in the data would still be recognized because of the strength of the expectation violation reaction. At the very least the data can be represented visually *and* auditorially together – and those who are more visual can monitor visually while those who are more auditory can monitor aurally (not necessarily substituting one for the other – but complementing and enhancing).

Human Factors Engineering, which looks at the design of human user interfaces [18], incorporates auditory cognition results in the design of multi-modal interfaces, comprising visual, auditory and haptic displays. These results have been applied extensively in medical and military applications, with the conclusion that sound provides an additional channel of information.

2 A Theory of Musical Sonification Design

Sonification, as defined in the context of auditory display, is generally intended for practical purposes. However, a methodology for designing "useful" sonifications has not yet been codified. A systematic account of sonification practice does not really exist since the field is so young. In this chapter, several approaches to a theory of musical sonification design will be suggested, as motivated by the following goals:

1. The sonification must be inviting, pleasant, even compelling when appropriate, otherwise the listener will turn it off and its practical value will be diminished.

2. The sonification must, as faithfully as possible, embody the data relationships and communicate them effectively to the listener.

The following disciplines would appear to have relevance to these goals:

1. Musical composition, especially with respect to "compelling, inviting, listenable" and the communication of relationships.

2. Music cognition, especially with respect to the perception of relationships.

3. Visual display, which may possess useful analogies to inform the design of auditory display.

The chapter will close with some reflections on what sort of data or situation lends itself to sonification, and what the process of designing a musical sonification is like.

2.1 Musical Composition

The central thesis of musical sonification design theory, and the reason that the word "musical" has been chosen in this context, is that in many cases the best and most useful sonifications will draw heavily upon the artistry of musical (particularly electro-acoustic) composition. This assertion has been made by some of the earliest workers in the field [19], however some have cautioned that a musical sonification might become "overloaded" with extraneous connotations which could lead to misinterpretation of the data. Fitch [20] has argued, for example, that if an interval-based sonification resulted in a progression from a minor third to a major third, the listener might experience "happy" emotions, which would be inappropriate if that sonification indicated an elevation in the blood pressure of a patient. Fitch's (medical) sonifications use for the most part FM synthesized sounds which mimic heartbeats, breathing, coughing, etc. He then "piggybacks" these synthesized sounds with pitch content, reverberation and modulation index to sonify other medical parameters.

In situations where recognizable acoustical metaphors or similes exist, Fitch's approach makes sense. However, it can also be argued that sonifications should draw upon the resources of shared cultural sonic experience, a major component of which is music. Recognition of musical intervals, for example, even among people who have no formal musical training, is a significant resource. The average person, for example, might not be able to sing, or respond to, a perfect fifth. However, the first two notes of the *Star Wars* theme (which form a perfect fifth) is, for better or for worse, part of the collective experience in most parts of the planet. Similarly, the first two notes of the song "Maria" from *West Side Story* by Leonard Bernstein, which form the interval of a tritone, or diminished fifth, would be perceived by most to be very different from the perfect fifth of *Star Wars*.

The comparison with "practical" music (e.g., movie music) which draws heavily on art/classical music to achieve its ends, is inescapable. Movie music, which most would consider essential to film, has now reached a high degree of sophistication, sometimes even drawing upon art music generally considered fairly inaccessible to the general public. For example, a film composer who had studied at Julliard during the time that serial music dominated compositional pedagogy related that he "escaped" to Hollywood upon graduation in order to work in an environment which welcomed tonal procedures. However, when asked to compose to a war movie sequence in which an airborne division landed via parachute and for some reason had no idea where they were, he found serial music to be effective in conveying that state of mind. Lack of a tonal center, at least in cultures whose popular music is dominated by "common practice" procedures, gives rise to a sense of "being lost" [21]. I recently took a CD of *Pierrot Lunaire* by Arnold Schoenberg to accompany a personal training session. My trainer, who has no formal musical training, immediately identified the style in the context of "scary" movie music.

Movie music, or for that matter any kind of program music which conveys emotions or narrative, e.g., the *Symphonie Fantastique*, or *Romeo et Juliette* by Hector Berlioz, could be described as a sonification of that emotion or narrative. In a sense, then, any kind of formal procedure in musical composition could also be thought of as a sonification. That is, any musical composition could be considered as the result of a process in which something (a concept, idea, algorithm, set of rules, a waterfall, Chopin's "anguish at hearing the news that Warsaw had fallen into the hands of the Russians" [22]) has been mapped into some series of instructions (machine or human) which generate sound.

It is probably fair to say then, that all music, in some sense, is a sonification. However, music is a subset of sonification. Many kinds of data sonification are arguably *not* music [23] in the sense that a “pure” sonification is intended for practical, rather than artistic purposes.

In “traditional” or “common-practice” composition, the “concept” or “set of rules” is an abstraction (e.g. sonata form), which can be described independently of a specific musical piece. However, the concept is usually a generalization of common musical practice and therefore arises from a musical context. In writing a sonata, the composer “maps” the abstract notion to her own specific thematic, harmonic and sectional scheme. The sonata is an interesting example since it is a larger scale “mapping.” Larger, romantic sonata movements (e.g., the first movement of Mahler’s 6th Symphony) can last as long as 30 minutes, but still convey the structure of exposition, development, recapitulation, etc. Those who are familiar with the form can identify it in pieces they have not previously heard or studied. The same can be said of the rondo form, theme and variations, passacaglia, etc.

These musical abstractions or constructs often have cyclical characteristics, involving a starting point, some sort of departure or journey, and a return. The return is sometimes striking because the composer generally sets up an expectation that the return will be like the beginning, but then introduces changes. Cycles are a type of pattern and pervade science, engineering, social behavior and human experience. It is intriguing to speculate that such patterns could be elucidated by mapping larger scale data trends to, e.g., a sonata form in such a way as to facilitate the perception of cycles.

As already mentioned in Chapter 1, the source material for the mapping process in more recent music has been broadened considerably to other disciplines, such as science, mathematics and engineering which traditionally have had little connection with music (other than in the realm of acoustics, the physics of sounds and musical instruments, audio engineering, etc.). Composers have experimented with the use of data from a great many sources to shape their musical compositions in some way.

It is significant to note that in cases where the abstraction of a concept from musical practice has occurred in the twentieth century, the concepts "look" more mathematical. Serialism à la Schoenberg could be considered as a sonification of a magic square (matrix with the 48 permutations of a tone row, 12 transpositions, P, R, I, RI). Babbitt's pitch class notation, which, e.g. replaces the nomenclature of "minor third", "perfect fifth", "major sixth" with the numbers 3, 7 and 9, respectively is another, very mathematical construct.

In order that musical sonification design be informed by musical composition, it is necessary to identify the boundaries between the determinism of a mathematical formula or data set, and the intentions of the composer. It is presumably the intervention, or choices made by the composer, which reflect her aesthetic values and distinguish "music" as a specific case sonification. These interventions or choices can provide valuable insight for the designer of sonifications.

In Chapter 3 I will give detailed examples of some of these boundaries, or interventions, in specific musical compositions. By way of introduction, some choice/intervention frameworks are listed here.

2.1.1 Algorithmic Choices

In algorithmic pieces, the music is determined by some sort of process which is set in motion by the composer and takes its course. The artistic input is, presumably, the choice of the algorithm, its implementation, and its mapping to sound. And beyond that, adjustments made afterward if the composer does not like the result. Boulez, for example used the deterministic procedures of total serialism in the composition of some of his early pieces, such as *Piano Structures*. However, when the procedures resulted in something he did not like, he would change it [24]. Sometimes an algorithm is used in a programmatic sense. In Honegger's composition *Pacific 231*, the composer used a mathematical acceleration algorithm, however, the musical implementation definitely gives the impression of an accelerating locomotive. The sonification designer must thus be prepared (and possess the skills) to modify a scheme for aesthetic reasons, or to shape the scheme to be suggestive of the sonified data context when appropriate.

2.1.2 Conceptual Pieces

Conceptual pieces are a more general category than algorithmic pieces. All algorithms are concepts of some sort, but the converse is definitely not the case. Sometimes the concept is musical, e.g., go up, go down (such as *The Piano Study in Mixed Accents* by Ruth Crawford Seeger) or get louder, get softer, (such as *Having Never Written a Note for Percussion*, by James Tenney see Fig. 2), but it need not be. The composition *MK Ultra* is described by the composer Christopher Penrose as being based on the concept of mind control. It is a musical exploration of illegal drug-induced mind control experiments carried out in the US in the 1950s, their effect on the subjects of these experiments, and the broader social implications [26]. *MK Ultra* is, however, not programmatic. Nonetheless, every section of the piece presents a different angle on mind control. A concept is a more overarching, large scale compositional framework, whereas an algorithm is a finer-grained set of instructions of how to get from A to B. In data sonification, concept is critical, to establish character and context.

2.1.3 Restriction of Materials

Building a substantial piece up from very simple or limited materials is a compositional practice of long standing. In *Art of the Fugue*, J. S. Bach used one subject to compose four simple fugues, three counterfugues, four double and

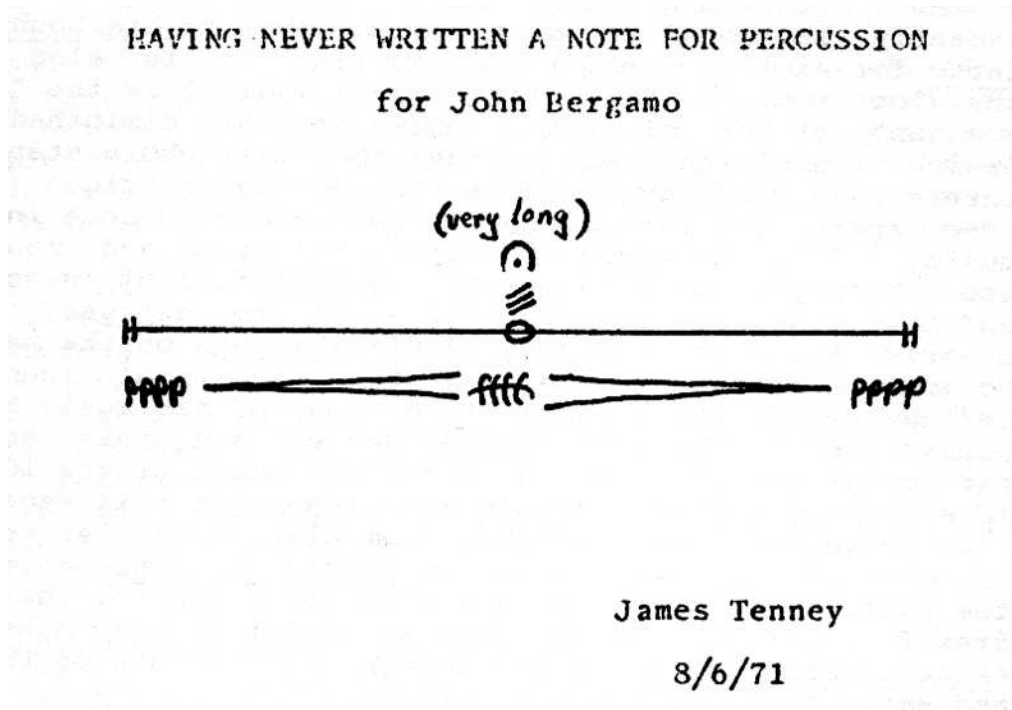

Figure 2: *Having Never Written a Note for Percussion [25]*

triple fugues, four canons, three mirror fugues and one quadruple fugue (for a total of nineteen)! Many sonata movements composed by Franz Joseph Haydn use the same thematic material for the first and second subject. The entire final movement of the *4th Symphony* by Johannes Brahms is a passacaglia, based on the same harmonic progression of 8 chords. Twentieth century minimalism is another embodiment of this practice. Electro-acoustic composers will often "make an entire piece from a one second sound file." Such a statement is invariably self-congratulatory, because the achievement requires a significant degree of inventiveness. However the result, when successful, is often coherent, elegant or compelling. There are strong arguments for restricting materials in data sonification. Since the intent of sonification is to communicate data relationships, musical overload must be avoided. That is, the relationships must be conveyed with the least musical content possible. If the musical aspects of the sonification predominate, the data relationships will be masked.

2.1.4 Self-Imposed Rules

In Western European music from approximately 1700 to 1900, compositional rules of harmony, voice-leading and form were the "common practice." The originality or inventiveness of the composer generally arose from finding new solutions, extending the rules, or breaking the rules. In the twentieth century, the practice of composing to rules continued, except that a significant part of the compositional effort lay in inventing new rules. The practice is related to restriction of materials, (see Section 2.1.3), however here the emphasis is on consistency, i.e., "sticking to the rules." In data sonification, the mapping scheme forms a large part of the rules. Finding a suitable mapping, and applying it rigorously is critical to the coherence of the sonification.

2.2 The Importance of Music Cognition to Sonification

The perception of the listener is an obvious consideration in musical sonification design. Within experimental psychology and brain science, music cognition is undoubtedly the most relevant discipline. A general review of this field is beyond the scope of this thesis. The investigation here will then be limited to the consideration of the following questions, especially with respect to listeners who have had "no formal musical training": How do the results and conclusions of music cognition experiments inform:

1. The potential for practical uses of sonification?

2. Sonification design?

2.2.1 Potential for Practical Uses of Sonification

Recent research in music cognition has explored recognition, learning and expectancy in populations which include both musicians and non-musicians.

According to Bharucha, et al.[27]:

> People without formal musical training recognize easily and quickly musical pieces or styles with which they are familiar. In the first few hundred milliseconds, timbre and voice quality can be used to infer a musical style, a performer or even a specific piece or song [28]. Over longer time spans, structural cues (tonality, modality, pitch contour, meter, rhythm) and lyrics can further facilitate recognition. The ability to recognize such complex auditory patterns without formal or explicit training presupposes a process by which the brain has adapted to or learned these patterns by sustained exposure.

The brain activity associated with such recognition can be inferred by monitoring blood flow changes using the techniques of positron emission tomography (PET) and functional magnetic resonance imaging (fMRI). The assumption is that as these cognitive tasks engage large groups of neurons, blood flow increases in order to meet the demand for more oxygen [29].

Various neural network models have been proposed and implemented to model this "passive" or "implicit" learning process. For example, Tillman, Bharucha and Bigand have proposed "a connectionist model that (a) simulates the implicit learning of pitch structures in Western harmony and (b) accounts for a range of empirical findings in music perception." [30] The characteristics of this type of learning are

1. It takes place passively, by mere exposure.
2. The subject cannot necessarily verbalize what she has learned.

Other researchers have proposed similar neural nets for learning, e.g., rhythm [31].

A key aspect of music cognition is the concept of an expectancy of something that will happen in the future, based on knowledge acquired in the past. This phenomenon is familiar to all who enjoy listening to music. Once a "melody" or "theme" has been perceived while listening to a piece, we naturally expect it to repeat. If it is repeated, but with a change, our attention is aroused. Many composers have manipulated this expectancy violation to achieve effect or aesthetic goals in their music.

Experimental studies of expectation violation in both language and music are common. Both language and music possess a syntax that is implicitly learned. When events occur which are difficult to integrate into the learned syntax, the expectation violation gives rise to an event-related potential (ERP) that can be measured by the electroencephalogram (EEG), recorded from electrodes placed on the scalp in various locations. By virtue of location, latency and polarity, the signature of these EEG responses can be characterized [29]. For example, the P600 is a positive voltage peak which occurs approximately 600 ms after the stimulus of a subject by words that are "difficult to integrate structurally into meaningful sentences." [32]. Similar responses have been measured in a musical context, when an "unexpected" or "out of place" target chord is presented to a subject within a familiar chord progression. For example at a position in the progression that the tonic in the key of C major (I) was expected, a nearby key (E♭ major) or a distant key (E major) was introduced. These unexpected target chords were found to elicit the P600 potential. Other auditory potentials, such as the N100, have also been measured [29]. Similar studies have been performed for deviant notes in learned melodic sequences [33]. There are even potentials associated with sensory deviants outside of the attention window (i.e., a reaction to something different that happens, even if you are not paying attention), such as the mismatch negativity (MMN) [34] Janata has shown that potentials are elicited by the *absence* of an expected note in a melodic sequence [29]. These results support the notion of auditory imaging, or audiation, in which a person "hears in their head" sound or musical sequences learned on previous occasions.

The above results, together with on-going research, suggest that people without musical training could, through the same passive/implicit process, "learn" relations in well-sonified data and respond to expectation violations. Potentially, these relations could be as subtle and refined as those which exist in the finest music, and could exist over a range of time scales (note to note, or from the beginning to the end of the Wagner Ring cycle). Sonification would thus provide an entirely new way for people to relate to data (as compared with visual), both in the attitude of perception (passive), and the nature of the relations. The potential is especially compelling in the case of data which changes over time so quickly that visual monitoring and interaction is unwieldy and erratic.

There is also a body of literature, e.g., [35] and [36] which suggests that a human being's interaction with sound is different from the visual and can have a more profound, lasting and potentially healing (or potentially damaging/disturbing) effect. There exist certain innate organizational principles in humans which resonate and respond in certain ways to sounds and music. For example, many people will unconsciously, audibly or inaudibly sing along with a melody that they hear and like, or will move their bodies to compelling rhythms. It could be argued that some composers consciously exploit these innate principles. For example, they know what kinds of rhythms and tempi are "compelling" and choose either to use them, or withhold them. It is known that in "house music," DJ's are aware of specific (and quite narrow) tempo ranges to which the people respond. Ballet dancers are extremely sensitive to tempi and describe a good conductor as one who "puts the music under their feet."

Another aspect of this profundity, which may only be known from anecdotal experience, is that advanced Alzheimer patients often retain memories of music they have known after all other memories are gone.

Yet another aspect is reflected in writings from other cultures, in which hearing appears to play a more significant role. For example, records of the Battle of the Little Bighorn, in which the Native American Sioux killed 263 soldiers including General Custer's entire command, state that before the battle, the Sioux leader Sitting Bull had a vision of white cavalry and soldiers falling down, as a voice said, "I will give you these because they have no ears." [37]

2.2.2 Implications for Sonification Design

Given some evidence that certain aspects of Western music are passively learned and retained by the general public, including non-musicians, it makes sense to design practical sonifications which would leverage off of that pre-existing knowledge. While the expectation violation has been experimentally measured for harmonic progressions, melody and timbre, it is likely that many other aspects of Western music are learned and expected in many contexts.

On the other hand, the sonification must not be embellished to the point that attention is focussed on the music rather than the information. Thus the sonification design must tread a fine line, using, where possible, recognizable musical constructs which may be learned and assimilated quickly, while avoiding over-decoration of these constructs with extraneous material which is not essential to communicating the information. The musical aspect of the sonification must not be so suggestive that it might give rise to an expectation in the listener which is wholly unrelated to the sonification.

Beyond the choice of materials, one question that commonly arises in connection with the sonification of complex or voluminous data is "How many data streams can be sonified simultaneously?" The answer to that question needs to be addressed from the point of view of *auditory stream segregation* and *attention*.

The former (*auditory stream segregation*) is the study of how people parse an auditory scene into groups or sources which are perceived to be distinct and is based on Gestalt principles developed from visual perception [38] [39]. Since, however, the sonification designer constructs the auditory scene in the first place, stream segregation should not be an issue. If the designer wished to sonify, say, the increase or decrease, over time, of ten variables, she could provide sonic updates for each variable that were sufficiently intermittent in time, and sufficiently separated in space and timbre that it was possible for the listener to identify the variables to which the sounds referred, whether the variable had increased or decreased, and by how much.

The latter (*attention*) is a more general psychological term dealing with the question "How many things can a person attend to *at the same time*?" [40] The attention limitation would then "kick in" at the point that, to use the example in the previous paragraph, the number of variables being monitored were increased to the point that the sonic updates became so frequent that they began to be perceived as concurrent, or could not be identified or understood quickly enough before the next one started. Psychologists generally agree that attention is limited and has some relationship to concurrent tasks, modality and practice. It is possible, for example, for multiple sonic updates to begin to be perceived as a Gestalt, especially if the data exhibit repetitive or similar behavior which results in a sequence, and the listener has become experienced in recognizing these patterns. The limits of attention do not appear to be quantifiable in the general case [40].

The previous two paragraphs assume that the intent of the sonification is to track several variables *separately*. The question becomes far more complex when the intent is to perceive higher order data relationships in which the data to sound mappings become interwoven. The question may be related to the perception of polyphonic music, in which we may choose to focus on a single part (selective), or just take in everything together (global). Recent research has suggested that this distinction may be reflected in brain activity [41].

In the design of complex, polyphonic sonifications, it is likely that the designer will draw on the highly developed craft of orchestration [42] to achieve goals of transparency, clarity and listenability.

2.3 Visual Display Analogies

The purpose of this thesis is to explore what might constitute a "good" sonification, recognizing that "good" may mean "artistic", and therefore that the sonification would be considered a "good" or "bad" piece, or it may mean "useful," and be judged according to how well the sound communicates the relations in the data to the listener. My belief is that, since music could be considered as a way of organizing sound events, an artistic sonification will do a good job, among other things, at organizing the sound so that the data relations emerge clearly. This notion will be explored at length in the Chapter 3, Section 3.1 study of *Achorripsis* by Iannis Xenakis.

Stated another way, music could be considered as a large amount of data, arranged to present coherent sound structures which convey meaning far beyond the sum of the individual parts, or data streams. While sonification is usually considered as a representation of one or more data streams simultaneously, it could also provide a soundscape more like the coalescence of multiple auditory streams too numerous to attend to individually (e.g., the sound of an orchestra) which would be intended to convey a more profound, global "sense" of the data.

One avenue to evaluating auditory displays is to study the principles of what makes a "good" visual display and see if these adapt to auditory display. Another way would be to forget about existing analogies (greater numbers are higher on a chart, therefore their sonification should be higher in pitch), and to build a set of sonification tools based on what we know about human reactions to music and other sounds. (Blind people obviously rely more on sound and their sense of "up" and "down" is from their ears rather than eyes).

2.3.1 Graphical Practice

Visual displays for the representation of data, in the form of x, y plots were published as early as the 10th century [17], p. 28. The first known bar chart was published in 1788 by William Playfair, a Scottish political economist [17], p. 33. In his book, *The Visual Display of Quantitative Information*, Edward R. Tufte [17] examines graphical practice from its early days until the present. In the first part of his book "Graphical Practice", Tufte presents many graphical examples, both good and bad, according to principles he formulates. These principles, which appear to be based on practical, rather than artistic considerations, are presented here in summary form, and their possible adaptation to the auditory domain will be discussed.

2.3.2 Principles of Graphical Excellence

Tufte's principles of graphical excellence are as follows [17], p. 51:

1. "Graphical excellence is the well-designed presentation of interesting data – a matter of *substance*, of *statistics* and of *design*."

2. "Graphical excellence consists of complex ideas communicated with clarity, precision and efficiency."

3. "Graphical excellence is that which gives to the viewer the greatest number of ideas in the shortest time with the least ink in the smallest space."

These principles can be adapted to auditory display, except that in the third principle, the concept of time, ink and space must be rethought. By time, Tufte refers to the amount of time a viewer needs to ponder a visual display before she "gets" the point. In auditory display, Tufte's "time" would probably translate to "repeated listenings." In most auditory displays, there is a training component, in which the user must listen to a sonification over a period of time to "hear" how the data been transformed to sound. So excellence here would imply a sonification that is easy to learn. Then, "ink" could simply translate to "sound," in the simplistic sense that "ink" is calculated by the area on the page that it takes up, and "sound" could be calculated by some combination of duration and energy (the integral of instantaeous power over the duration). Sonic design should be transparent and not overloaded with extraneous elements. "Space" in the visual world would usually translate to "time" in the auditory domain. In the sonification of real-time data, Tufte's principle of "least ink" would presumably favor sparseness in the sound. However, sonic sparseness may militate against intelligibility and coherence.

An interesting angle on time arises from the related field of audification, in which data at lower sampling rates is played back at higher rates so that the characteristic frequencies fall within the audio range. See the discussion in Section 1.3.2.

2.3.3 Principles of Graphical Integrity

Tufte's review of graphical presentations, especially those which are politically motivated, shows how major distortions of the truth can occur. His principles of graphical integrity are [17], p. 77:

1. "The representation of numbers, as physically measured on the surface of the graphic itself, should be directly proportional to the numerical quantities represented."

2. “Clear, detailed and thorough labeling should be used to defeat graphical distortion and ambiguity. Write out explanations of the data on the graphic itself. Label important events in the data.”

3. ”Show data variation, not design variation.”

4. “In time-series displays of money, deflated and standardized units of monetary measurement are nearly always better than nominal units.”

5. “The number of information-carrying (variable) dimensions depicted should not exceed the number of dimensions in the data.”

6. “Graphics must not quote data out of context.”

These principles raise some important sonification issues: first, the issue of proportionality and scaling. Here, Tufte argues that, e.g., the ratio of height between two bars of a chart should be the same as the ratio of the data itself. If, for example, the price of a stock is mapped to volume, should an increase by a factor of two in the price map to doubling of the volume? If mapped to pitch, should it be a doubling of the frequency (octave)? There may be strong arguments in auditory perception theory which require some adjustment and further exploration. In [43], Polansky describes a live interactive piece called *Three Studies* in which he effectively sonifies “distance” in three ways, using, in turn, melody, rhythm and distance.

The second issue is labeling. Some sort of voiceover or agogic accent might be used, as long as it does not distract the user from the information in the sound itself. Ideally this principle would point to a sonification that needs no explanation (such as douglas repetto’s *sine family tree*, `http://music.dartmouth.edu/~douglas/portfolio/sinefamilytree.html`). Second best might be a clear sonification interface, in which the user understands exactly at glance which data is being mapped to sound and how. Perhaps the best interface would be one in which the users sets up her own mapping strategy.

The third issue is variation. Here, the composer of the sonification must resist the temptation to artistically embellish. For example, *Climate Symphony* by Marty Quinn, which maps ellipticity of the earth's orbit, obliquity, ocean circulation, precession, subprecession, ice sheet movement and its subharmonic, solar variability, volcanic activity and time to various musical patterns and constructs `http://www.quinnarts.com/srl/gisp2/climatesymphony3.html`. Most of these mappings are to elaborately decorated melodic sequences, and while the piece is enjoyable, the data is masked by the embellishment.

The fourth and sixth principles really have more to do with the correct reporting of financial data over long time periods, or just general honesty in selecting data in the first place. They are not, *per se*, graphical or sonification issues.

The fifth principle opens up a fascinating difference between the visual and the sonic. In the visual world there are three interrelated concepts: length, area and volume (all of which can be represented, using perspective where necessary, on the page of a book, or the screen of a computer or television. Tufte is referring to the misuse of *area* or *volume* to represent the variation in only one dimension of information. That is, if one uses the area of a square to represent dollars, and doubling the dollars leads to doubling the length of the side of the square, the area of the square will quadruple, distorting the data by a factor of two. It is better, according to Tufte, to use a single space dimension for a one dimensional variable.

In the sound world, there are many "dimensions," which may be assigned independently, including pitch, duration, volume, "timbre," envelope (which is part of "timbre") and time. (However, these sound dimensions are not *perceived* independently). As an aside, definition of the term “timbre,” and the determination of the number of its dimensions, is somewhat controversial, and a full discussion is outside the scope of this thesis. The word will be used, loosely, to refer to the (often) spectral characteristics of a “voice” which, for example, distinguish a trumpet from a violin. To apply Tufte’s principle, there would be no reason, e.g., to map the number of lemons imported from Colombia every month to both pitch and volume. However a distinct timbre for lemons from Colombia might be needed if the number of lemons from Venezuela is also tracked.

2.3.4 Sources of Integrity and Excellence

Tufte argues that many published graphics are mediocre [17], p. 87. “Graphical competence demands three quite different skills: the substantive, statistical and artistic. Yet now most graphical work, particularly at news publications, is under the direction of but a single expertise – the artistic.” So also, the auditory display artist should possess the ability to understand the substance (i.e., the real story that the data is telling), have the quantitative skills to deploy faithful statistics, and sophistication in rendering these components into sound.

2.3.5 Theory of Data Graphics

In the second section of his book, Tufte presents his theory of how the principles in the first section (see Section 2.3.1) can be realized.

2.3.6 Data-Ink and Graphical Redesign

His recommendations for realizing the third principle of graphical excellence (see Section 2.3.2) [17], p. 105, are:

1. “Above all else show the data.”

2. “Maximize the data-ink ratio.”

3. “Erase non-data-ink.”

4. “Erase redundant data-ink.”

5. “Revise and edit.”

The data-ink ratio R_{di} is defined as [17], p. 93:

$$R_{di} = \frac{di}{ti}, \tag{2.1}$$

where di is the amount of ink used for data and ti is the total amount of ink used to print the graphic.

These recommendations illustrate some basic differences in visual and auditory displays. Tufte is advising the graphical designer to reduce the amount of ink used for displaying axes, tick marks and grid lines. These marks are used to orient the viewer to the range and quantitative value of the data, and are all perceived at the same time. The eye can line up data points of interest with tick marks on the axes. Excessive use of these reference markings, however, may overwhelm the data.

Auditory displays usually show trends or complex patterns which evolve over time. Tufte's graphical axes, then, might translate, e.g. to some sort of reference pitch(s) (in one timbre) for pitch-mapped data (in some other timbre) so that the listener could detect, at least qualitatively, when the data was above or below (or in between) the reference(s). As referenced earlier, Polansky [43] explores the use of melodic, rhythmic and harmonic sonifications of distance/reference.

The sound-data ratio could presumably be reduced by displaying reference tones as short rhythmic bursts and data as a continuously varying line.

The use of reference durations or envelopes may also be possible. Reference volumes and timbres may be far more difficult to achieve.

Tufte's general recommendations for simplicity, sparseness and continual revision may clearly be applied to auditory displays. However, there are other considerations in sound design, such as pulse, number of melodic/motivic lines, consonance, register, etc. for which sparseness is not necessarily applicable.

2.3.7 Chartjunk: Vibrations, Grids and Ducks

Tufte's overall recommendation here is [17], p. 121:

1. "Forgo chartjunk, including moiré vibration, the grid and the duck."

The moiré vibration, the grid and the duck are artifacts of visual displays which attract attention to themselves, instead of the data. There would definitely be analogies to auditory displays:

1. Any "busyness" in the sound display which compels the listener's attention.

2. The overuse of quotations or references to existing music which unduly preoccupy the listener.

An example from Tufte's book of an overly worked up bar chart is shown as Fig. 3. Excessive use of crosshatching results in visual "vibrations." A simpler presentation of the same data in tabular form actually provides more information (see Fig. 4).

2.3.8 Redesign

The remainder of Tufte's book deals with the innovative redesign of standard graphical devices such as the bar chart, the scatter plot, the flow chart. While these devices have no direct analog in auditory display, most of the innovations are motivated by the desire to simplify and remove extraneous ink, or to make a graphical marking serve more than one function. Since the field of auditory

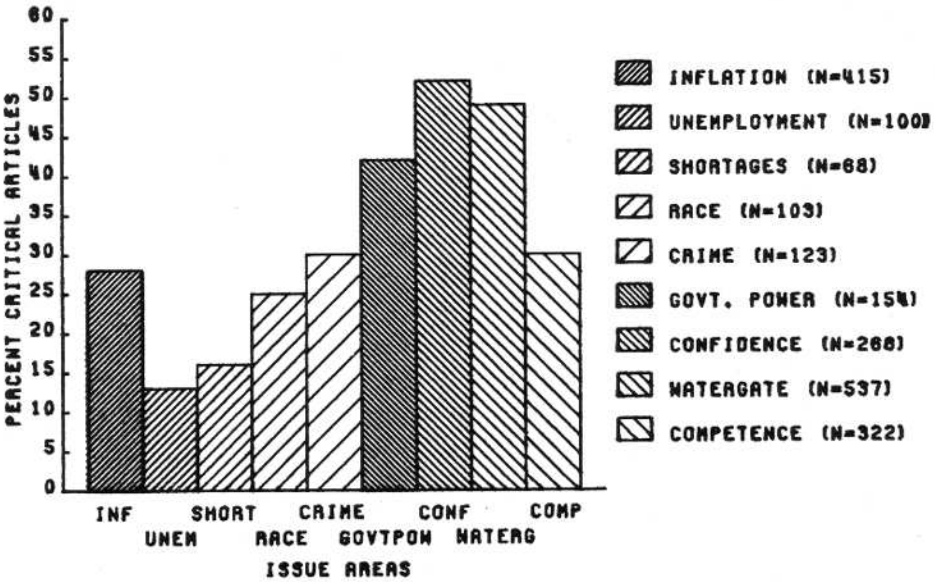

Figure 3: *Bar Chart with Vibrations [17]*

Content and tone of front-page articles in 94 U.S. newspapers, October and November, 1974	Number of articles	Percent of articles with negative criticism of specific person or policy
Watergate: defendants and prosecutors, Ford's pardon of Nixon	537	49%
Inflation, high cost of living	415	28%
Government competence: costs, quality, salaries of public employees	322	30%
Confidence in government: power of special interests, trust in political leaders, dishonesty in politics	266	52%
Government power: regulation of business, secrecy, control of CIA and FBI	154	42%
Crime	123	30%
Race	103	25%
Unemployment	100	13%
Shortages: energy, food	68	16%

Figure 4: *Alternative to Bar Chart [17]*

display is so young, there may as yet not be any clichés like bar charts and pie charts, however the principles of sparseness and simplicity still apply.

One advantage of following Tufte's principles is that the graphic may be reduced in size and still be readable, allowing the deployment of multiple copies of a similar display to convey multivariate trends. This leads to another interesting concept, that of data density. Tufte gives an example of a square 80×80 grid imposed on one square inch, which comprises a total of 25,281 distinctions ([17], p. 161), close to the visual limit for density of data. Currently, the highest data densities are achieved in cartography.

It would be extremely interesting to determine the limit of perception for data density in sound, if this has not already been done, bearing in mind that the metric for finest distinct gradation is sound must be measured differently for time, pitch, loudness, timbral and structural differences.

2.4 What Is Sonification Good for?

Under what circumstances, compared to visual data displays, are sonifications useful? I would suggest that there may be three general scenarios in which there is significant benefit from sonification. The first is any situation in which time plays some significant role in the data. Examples of the first criterion would be

- The monitoring of real-time data, in which the user is attempting to attend to rapidly changing data on multiple visual displays. Visual monitoring requires the user to stare at a fixed place on the screen and there are significant limitations as to how many different quantities can be monitored simultaneously.
- "Listening" to large, time-dependent historical data sets, in which acceleration of the data results in audible sound.

- Scanning voluminous data sets for patterns or characteristics which have not yielded to visual or analytical analysis in a reasonable amount of time.

The second scenario would include situations where vision is needed for some other function (e.g., a surgeon who must focus visually on a procedure and cannot "look" at a screen for vital patient information). The use of a Geiger counter was tested on a group of subjects with auditory display only, combination auditory/visual display and visual display only. The subjects were far more efficient locating simulated radioactive material with the auditory display only configuration [44].

The third scenario would be the need to analyze, or detect small or subtle changes in data which are difficult to perceive visually. The human auditory system is generally sensitive to small pitch differences in sustained tones, pops, clicks or distortions in musical recordings, small tempo changes, etc.

Sonification is not generally useful in the representation of data whose structure or relationship is immediately apparent from an elegant visual display, such as a scatter plot. For example, the "Sonic Scatter Plot" of Madhyastha and Reed [45] takes longer to listen to, and is more difficult to understand, than its visual counterpart shown in Fig. 5.

2.5 The Sonification Design Process

Having worked on several sonification projects in the past three years, I would offer the following outline of the design process:

1. Understand as much about the data as possible. The more "domain expertise" one can gain, before designing the mapping, the better, since the real intent is to communicate the relationships in the data. If these relationships are not understood, the sonification is likely to be ineffective. This "background research" step is often performed by composers in the

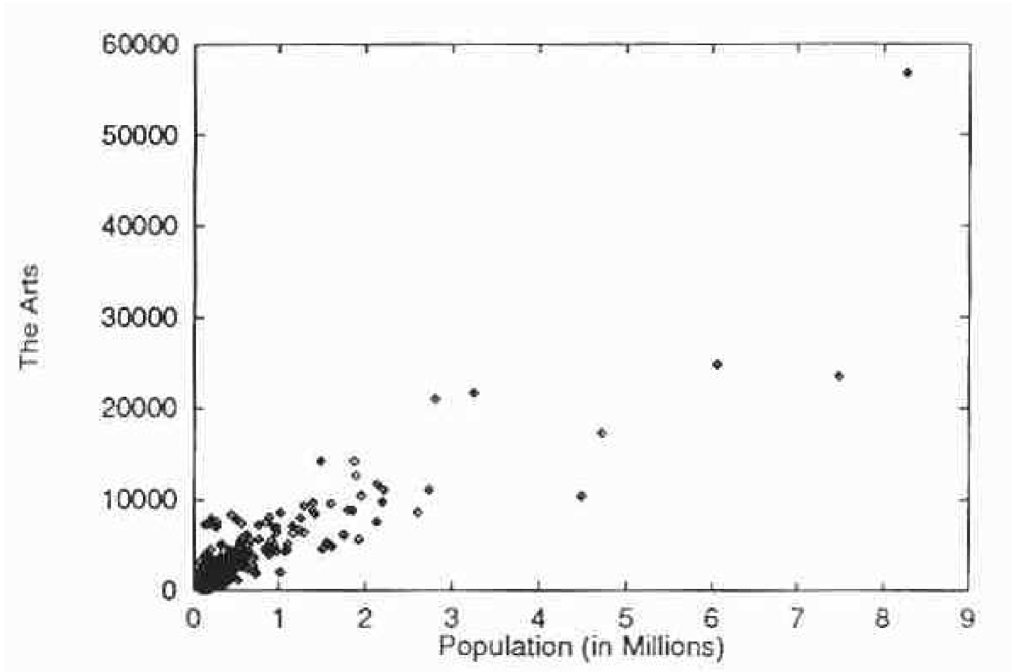

Figure 5: *Scatter Plot of Arts Population Index* [45]

planning stage of a new piece.

2. Organize the data as much as possible. If one data quantity is known to relate to any other data quantities (e.g., have a fixed ratio), the mapping should exploit this known relationship. If the number of variables determining a function can be reduced by dimensional analysis, these dimensionless variables should be used in the mapping. Bounding of all data sets should be performed, and normalization should be considered.

3. Use conceptually consistent mappings. In general, one should map time to time and physical location to auditory space. Largeness might be mapped to an auditory variable that is related to the size of an acoustic space (i.e. reverberation parameters).

4. An interpretation of the data, where possible, should be used in the mapping. For example, if some aspect of the flow of air is being sonified,

perhaps the sonification should sound like air moving. This step is greatly aided by artistic imagination, and closely related to musical composition.

5. Implement the mapping and use one's ears to critique the result. In this step it is essential to focus on "what do you hear?" The coherence of the mapping, and its implementation, can only be confirmed by detailed listening, followed by adjustment. This step is part of the working experience of nearly every composer of electro-acoustic music.

6. Test the sonification as widely as possible. Use as many data examples as possible to see if variations in the data scenarios can be effectively heard and are correctly reflected in the sonification.

3 Musical Case Studies

In this chapter, several musical compositions in which sonification is a key component will be analyzed, the primary purpose being to determine:

1. How did they go about it?

2. How faithful were they to their sonification scheme?

3. Where are the boundaries between artistic choice and determinism from the mapping scheme?

Achorripsis by Iannis Xenakis was chosen first because it is a landmark piece in the use of probability distributions to shape musical composition. Xenakis did not have access to a computer so performed all calculations by hand. It is fairly well documented in *Formalized Music* [46] and thus lends itself to detailed analysis and verification. Furthermore, except for work done in parallel by another researcher, which takes a different approach [47], no detailed analysis of the work has been published. I was primarily concerned with teasing out the boundaries between artistic choice and determinism, feeling that the understanding of these choices could provide enormous insight for sonification design.

Earth's Magnetic Field by Charles Dodge is a landmark piece in the use of raw scientific data to generate pitches in a systematic manner. While the term sonification was not in use at the time it was composed, the piece is probably the first electro-acoustic sonification.

The Red and White Cows by Daniel Goode is drawn from a tradition of algorithmic, enumerative pieces composed by Goode, Tom Johnson and others.

Lottery by Nick Didkovsky is a unique and highly unusual real-time sonification of a social interaction.

3.1 Achorripsis

3.1.1 Purpose

The purpose of this section is to explore the boundaries between "traditional" composing practice, and the application of mathematical formulae derived from science and technology, in the composition *Achorripsis* by the Greek-French composer Iannis Xenakis. The motivations for this exploration are:

1. To determine how rigorously Xenakis applied his formulations to generating specific notes. Did he really use all those complicated formulae to write his music?

2. To assess whether or not the scientific concepts underlying *Achorripsis* transfer effectively to sound. Do we really "hear" his piece (at least in part) as a "sonification" of probability distributions?

While Xenakis did not have sonification in mind, his artistic choices in rendering mathematical formulations into musical events (time, space, timbre, *glissando* speed) provide useful contributions to the "mapping problem" in three significant ways:

1. He pushes the limit of loading the ear with multiple formulations simultaneously.

2. His mapping of "velocity" to string *glissando* speed provides a useful method of working with a vector quantity with magnitude and direction.

3. His artistic renderings, i.e., "musifications" of these distributions, invite the question, in general, as to whether musical/artistic sonifications are

more intelligible to the human ear than sonifications prepared without any musical "filtering" or constraints (e.g., that they could be notated and performed by musicians).

3.1.2 Biographical

Xenakis was professionally involved with three distinct disciplines: music, architecture, and science and mathematics. In 1976, he received a "Doctorat d'État" from the Sorbonne [48], for his contributions in these three fields. According to James Harley [49], he studied civil engineering at the Athens Polytechnic and later worked in Paris as an engineering assistant for Le Corbusier, who was so impressed with Xenakis' work that he delegated architectural projects to him. He went on to design the Phillips Pavilion at the 1958 World's Fair in Brussels. His short tape piece, *Concrète PH*, was used as an interlude between performances of Varèse's *Poème Électronique* in the Pavilion. As a composer, Xenakis studied extensively with Olivier Messaien, who encouraged Xenakis to use his mathematical and engineering background in composition, and composed full time from 1960 [49]. His science and mathematical work went far beyond civil engineering into the kinetic theory of gases, probability theory and computer science. Many of his compositions were implemented through the use of computer programs. In the 1970s, Xenakis invented the UPIC [50] system which allows the user to create graphical designs on a tablet and have them rendered directly into sound. His last composition *O-Mega* was premiered in November 1997 [49] and he died on February 4, 2001.

3.1.3 Influences

Xenakis was the first 20th century composer to do significant "crossover" work between these three disciplines. For example, in *Metastasis* (1954), his first major *succès du scandale* [49], he used computations based on the mathematical permutations of intervals [51], pp. 72-73. The piece, for string orchestra, the first to employ complete *divisi* in the strings, consists almost entirely of string *glissandi* generated from elegant line drawings. Later (1958), he used these musical drawings to design the World's Fair Phillips Pavilion.

Pithoprakta (1955-56) was the first composition in which Xenakis used probability distributions as a tool. It was, in part, his response to dodecaphonic music. In his article, "The Crisis of Serial Music" [52] [46], he wrote:

> Linear polyphony destroys itself by its very complexity; what one hears is in reality nothing but a mass of notes in various registers. The enormous complexity prevents the audience from following the intertwining of the lines and has as its macroscopic effect an irrational and fortuitous dispersion of sounds over the whole extent of the sonic spectrum. There is consequently a contradiction between the polyphonic linear system and the heard result, which is surface or mass. This contradiction inherent in polyphony will disappear when the independence of sounds is total. In fact, when linear combinations and their polyphonic superpositions no longer operate, what will count will be the statistical mean of isolated states and of transformations of sonic components at a given moment. The macroscopic effect can then be controlled by the mean of the movements of elements we select. The result is the introduction of the notion of probability, which implies, in this particular case, combinatory calculus. Here, in a few words, is the possible escape route from the "linear category" in musical thought.

Xenakis' goal was to compose music whose underlying structure, though unusual, could be clearly heard. He felt that the linear polyphony of serial music had become so complex that it could no longer be heard as independent lines. He invented the idea of using statistical distributions as the building blocks for his music.

His interest in probability was also inspired by the sounds he heard while fighting in the Greek resistance [46]:

> Everyone has observed the sonic phenomena of a political crowd of dozens or hundreds and thousands of people. The human river shouts a slogan in a uniform rhythm. Then another slogan springs from the head of the demonstration; it spreads towards the tail, replacing the first. A wave of transition thus passes from the head to the tail. The clamor fills the city, and the inhibiting force of voice and rhythm reaches a climax. It is an event of great power and beauty in its ferocity. Then the impact between the demonstrators and the enemy occurs. The perfect rhythm of the last slogan breaks up in a huge cluster of chaotic shouts, which also spreads to the tail. Imagine, in addition, the reports of dozens of machine guns and the whistle of bullets adding their punctuations to this total disorder. The crowd is then rapidly dispersed, and after sonic and visual hell follows a detonating calm, full of despair, dust and death. The statistical laws of these events, separated from their political or moral context, are the same as those of the cicadas or the rain. They are the laws of the passage from complete order to total disorder in a continuous or explosive manner. They are stochastic laws.

3.1.4 Criticism

His inventions and music are controversial. Some critics suggest that his extensive writings on his own musics, full of numbers and complex equations, are intentionally obscure.

> As far as Xenakis is concerned, let me emphasize at once that I'd be much more interested in his research if he hadn't set out so obviously to reduce its accessibility and its credibility in a manner which is immediately apparent as soon as you open his book on formal musics. *Pierre Schaeffer* [53]

Schaeffer also accused Xenakis of delegating compositional decisions to a computer [53].

Others criticize his deployment of mathematical equations to analyze music.

> One need only peruse the the theoretical writings of such polemicists as Allen Forte and Iannis Xenakis to understand how the misapplication of scientific methods of inquiry can result in opaque and ultimately meaningless music "theory." *Jon Appleton* [54]

Many felt that his music did not "sound" as though it were derived from mathematical equations [55].

> The seeming contradiction between Xenakis' intellectualized starting point for pieces and the primeval energy of the music is often questioned. *BBC Online*

Harley [56] notes the difficulties Xenakis' music poses for the musical analyst who does not possess an engineering or mathematical background:

> The theoretical formulations Xenakis put forward in his book *Formalized Music* are such that the reader could easily conclude that further work would require advanced training in mathematics, something most music students do not receive as part of their musical training. Furthermore, even if the formulas, graphs and matrices do

not discourage the potential scholar, it soon becomes apparent that to trace the specific processes by which a musical result was achieved would necessitate having access to a great deal of data used by the composer. Stochastic processes, after all, by their very nature are indeterminate; a specific outcome can only be predicted within a limited degree of probability without having access to all of the numerical inputs and constraints on the mathematical functions.

3.1.5 Sonification

In the past decade, researchers from a variety of backgrounds have been producing sounds based on mathematical equations, computational processes or data from experiments, simulations and other monitoring situations. The procedure is to take data with non-musical significance (e.g., a patient's blood pressure) and to "map" that data into musical sound (e.g., a trumpet whose pitch goes up and down according to the instantaneous value of the blood pressure). The goal is to enable the listener to monitor data by ear as well as by eye, and to use his/her auditory recognition skills to perceive patterns not readily apparent from visual data displays. Much of the work in this field has been presented at the *International Conference on Auditory Display* [57], whose first conference took place in 1992.

A successful sonification depends on a good "mapping," and Xenakis' musical realizations of various probability distribution may be considered as significant, pioneering contributions to sonification (even though, of course, Xenakis' goals were artistic rather than practical).

3.1.6 Performances

His second major composition (not counting some earlier, withdrawn material) *Achorripsis* (Greek for "jets of sound"), composed in 1956-57, was first performed in Buenos Aires in August 1958 under the direction of Herman Scherchen, who, until his death in 1963, championed Xenakis' music [51] [58]. The work received further performances in 1959 in Europe (to mostly scandalous reaction), and in the early 1960s in America under the direction of Gunther Schuller, Lukas Foss and Leonard Bernstein [51]. *Achorripsis* had a major success during the first all-Xenakis festival at the Salle Gaveau in Paris in 1965, performed by the Ensemble de Musique Contemporaine under the direction of Konstantin Simonovitch from which the only extant recording of the piece was made [58]. The vinyl LP release has the curious flaw that the first few minutes of the track labeled *Achorripsis* is some other piece. The recording was never transferred to CD.

3.1.7 Background

Achorripsis broke new ground in at least two significant ways. Firstly, Xenakis made extensive use of four distinct probability distributions to define overall structure, time between musical events, intervals between successive pitches and *glissando* speed. Secondly, his conception of the compositional process was similar to a flow chart that would be used to prepare for a computer programming project. In fact, during 1957, Xenakis proposed to Pierre Schaeffer, the director of the Groupe de Recherches Musicales, that the studio undertake significant pioneering work in computer music. Schaeffer, who was more interested in recording and manipulating natural sounds, rejected the proposal. Xenakis finally gained access to an IBM 7090 computer at the IBM headquarters in Paris beginning in 1961, thanks to the intervention of Herman Scherchen [59]. He then implemented the procedures he had outlined for *Achorripsis* in the composition of *ST/10*, a work for 10 instruments first performed in 1962, and, subsequently, several other works composed with the FORTRAN IV program *FREE STOCHASTIC MUSIC* [46].

3.1.8 Process

The process for *Achorripsis* and later computer compositions was as follows [46], p. 22:

1. Initial conceptions.
2. Definition of sound entities (instrumental, electronic or other material).
3. Definition of transformation (macrocomposition: kinds of operations 2 above would undergo and their arrangement in succession or simultaneity).
4. Microcomposition (choice and detailed fixing of relations).
5. Sequential programming of 3 and 4 (schema and pattern of entire work).
6. Implementation of calculations.
7. Final symbolic result of the programming (notation).
8. Sound realization.

In *Achorripsis*, this was accomplished entirely by hand.

3.1.9 Documentation

Most steps of this procedure are discussed in detail in *Formalized Music* [46]. Xenakis provides numerous equations, together with sample calculations, and a detailed discussion of one of the more complex sections of his score. In this section Xenakis' equations and sample calculations will be closely compared with the score in an attempt to disentangle his own artistic decisions from those dictated by his theories, and to determine the nature of the relationship between his numbers and his notes.

3.1.10 Analysis

Achorripsis is constructed on several levels in an architectural fashion, starting with the large-scale concept of the piece and moving downward, level by level, to the fine details. These levels (as indicated in the Section headings, below) correspond closely with the steps in the Xenakis procedure outlined in Section 3.1.8. At each level or step, it is possible to determine:

1. The artistic decisions.
2. The theory and calculations.
3. The actual results.

3.1.11 Level 1 - Conception and Sound Entities

Artistic Decisions

The overall scheme for *Achorripsis* is shown in Fig. 6, and consists of a matrix of 28 columns (representing time blocks) and 7 rows (representing timbral classes of instruments). These timbral classes are:

1. *Flute* (Piccolo, E♭ Clarinet, Bass Clarinet)
2. *Oboe* (Oboe, Bassoon, Contrabassoon)
3. *String glissando* (Violin, Cello, Bass)
4. *Percussion* (Xylophone, Wood Block, Bass Drum)
5. *Pizzicato* (Violin, Cello, Bass)
6. *Brass* (2 Trumpets, Trombone)
7. *String arco* (Violin, Cello, Bass)

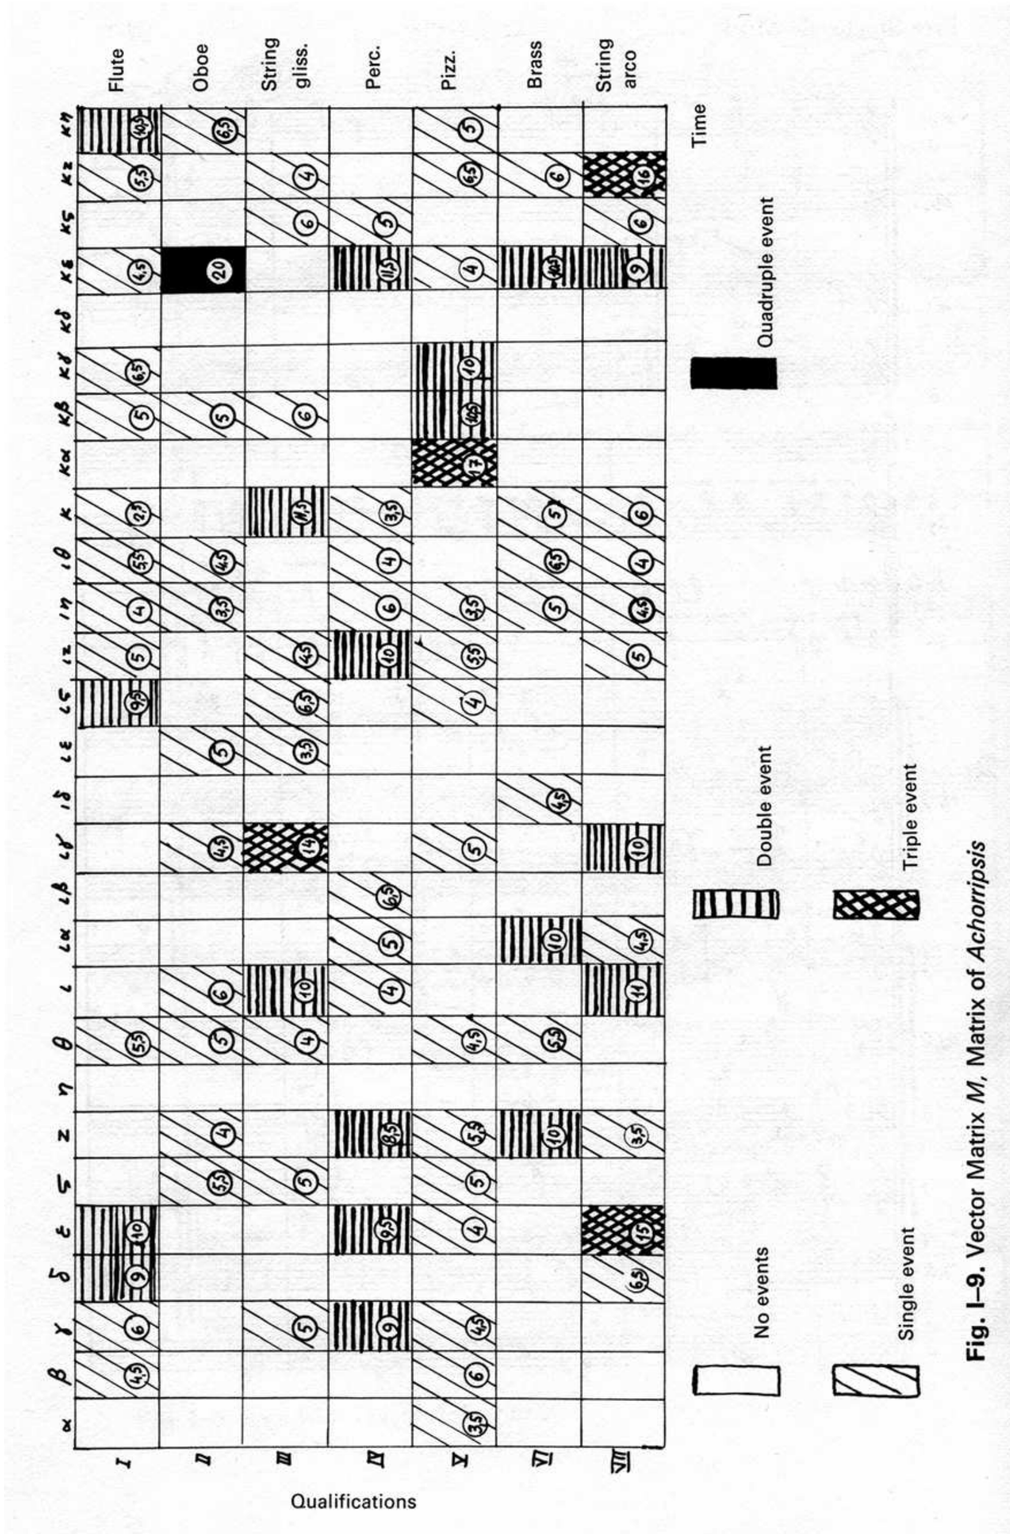

Figure 6: *The Matrix of Achorripsis* [46]

where the *italicized* entries are the *names* of the timbral classes and the parenthesized instruments are those that make up that class. In all, there are 3 Violins, 3 Cellos and 3 Basses, and all strings move back and forth over the course of the piece between *glissando*, *pizzicato* and *arco* passages.

The total length of the piece is set to be 7 minutes, which means that each of the 28 columns lasts 15 seconds. Each of the 28 time blocks of 15 seconds is set to be 6.5 measures in length, in which the time signature is $\frac{2}{2}$ with half note = MM 52. Thus each measure (two half notes) lasts $\frac{120}{52}$ seconds, and there are 182 measures in the score.

Theory and Calculations

No probabalistic theories are invoked at this level.

Actual Results

On this level, the layout of the score matches exactly with the scheme set out as Fig. 6. However, there are different timbral variations specified in the score which are not announced in his scheme and appear to play no part in his theories:

1. The string *arco* passages all consist of short notes played either all down-bow, all up-bow or as harmonics.

2. The brass all apply mutes from the 14th to the 21st time block.

3. The final string *glissando* passages are *sul ponticello* and *tremolo.*

The deployment of these additional timbral colorations appear to have no place at all in his announced compositional scheme, and therefore are freely applied as a matter of artistic choice.

3.1.12 Level 2 - Macrocomposition I

Artistic Decisions

In the next level/step, Xenakis decides how to allocate musical events to the $7 \times 28 = 196$ cells of the matrix. To do this, he starts with the assumption that the average number of events/cell $\lambda = 0.6$. (At this level, no specification has been made as to what constitutes an "event").

Theory and Calculations

Xenakis invokes the Poisson probability distribution, which is used for situations in which one wants to estimate how many instances of a particular event will occur in a given time or space, when the average behavior is known. For example, how many subway travelers in NYC per minute are likely to swipe their Metrocards in a particular turnstile, from 10 AM to 11 AM, on a particular day, at the Hunter College Station on the 6 line, if it is known, that, on average, a given number of people pass, per hour, per turnstile?

In the case of *Achorripsis*, given the artistic choice that the average number of events per cell is 0.6, what is the probability that in any given cell there will be 0, 1, 2, 3, 4 or 5 events occurring? (Here, the time length of each cell is 15 seconds, and the "space" is the location of the cell within the 7 timbral instrument classes.)

This may be estimated using Poisson's formula:

$$P_k = \frac{\lambda^k}{k!} e^{-\lambda} \tag{3.2}$$

where k is the number of events ($k = 0, 1, 2, 3, 4, 5$ in this situation), e is the base of natural logarithms ($e = 2.71828\ldots$) and $k!$ (k factorial) for $k! = 5! = 5 \cdot 4 \cdot 3 \cdot 2 \cdot 1$. Equation 3.2 is valid as long as $\lambda < 7$ [60]. By definition $0! = 1$. For example, the probability of 0 events occuring in a cell is:

$$P_0 = \frac{0.6^0}{0!}e^{-0.6} = 0.5488 \tag{3.3}$$

from which we see that in slightly over half of the cells, no events will be occurring: $196 \times 0.5488 = 107$. Equation 3.3 can be calculated easily today on any calculator which has the exponential function. Xenakis probably used tables. Applying this same procedure for $k = 1, \ldots, 5$, we find that the number of cells n_k in which k events occur is:

$$\begin{aligned}
n_1 = P_1 \times N &= 0.3293 \times 196 = 65, \\
n_2 = P_2 \times N &= 0.0988 \times 196 = 19, \\
n_3 = P_3 \times N &= 0.0198 \times 196 = 4, \\
n_4 = P_4 \times N &= 0.0030 \times 196 = 1, \\
n_5 = P_5 \times N &= 0.0004 \times 196 = 0.
\end{aligned} \tag{3.4}$$

where N is the total number of cells in the matrix, i.e., $N = 196$.

Actual Results

As one can see from Fig. 6, Xenakis has implemented his calculations exactly. (Just count the number of single events, etc. and compare with Equation 3.4). Note that he refers to two events as a double event, etc. Only one quadruple event occurs, and there are no quintuple events. The resulting matrix is quite sparse, with nothing occurring in over half of the cells. This is a direct result of his decision, in Section 3.1.12, to set the average number of events per cell λ to 0.6.

3.1.13 Level 3 - Macrocomposition II

Artistic Decisions

Xenakis imposes an additional constraint on the distribution of the various event classes among his 196 cells. He decrees that the frequencies of zero, single, double, triple and quadruple events be statistically distributed amongst the 28 time blocks in accordance with Poisson's Law. Thus, the new "unit" or "cell" is now the time block.

Theory and Calculations

Since there are a total of 65 single events distributed over 28 cells, the average number of single events per cell is now $65/28 = 2.32$, which becomes the new λ in the reapplication of Poisson's Law, so that the probability P_0 of no single events occurring in a cell (time block) is:

$$P_0 = \frac{2.32^0}{0!}e^{-2.32} = 0.09827 \tag{3.5}$$

so since there are 28 time blocks, the number $t_{0,single}$ in which no single events occur is $28 \times 0.09827 = 3$. We may now calculate the number of time blocks $t_{k,single}$ in which k single events occurs, $k = 1, \ldots, 7$:

$$\begin{aligned}
t_{1,single} = P_1 \times T &= 0.22799 \times 28 = 6, \\
t_{2,single} = P_2 \times T &= 0.26447 \times 28 = 8, \\
t_{3,single} = P_3 \times T &= 0.20453 \times 28 = 5, \\
t_{4,single} = P_4 \times T &= 0.11862 \times 28 = 3, \\
t_{5,single} = P_5 \times T &= 0.05504 \times 28 = 2, \\
t_{6,single} = P_6 \times T &= 0.02128 \times 28 = 1, \\
t_{7,single} = P_7 \times T &= 0.00705 \times 28 = 0.
\end{aligned} \tag{3.6}$$

where $T = 28$ is the total number of time blocks.

Now, we know that from Equation 3.4 there are 19 double events to be distributed over 28 time blocks. We reapply Poisson's equation with $\lambda = 19/28 = 0.67857$ in order to obtain the number of time blocks $t_{k,double}$ in which k double events occurs ($k = 0, \ldots, 4$)and finally obtain:

$$\begin{aligned} t_{0,double} &= 14, \\ t_{1,double} &= 10, \\ t_{2,double} &= 3, \\ t_{3,double} &= 1, \\ t_{4,double} &= 0. \end{aligned} \tag{3.7}$$

Calculations for the relatively small number of triple and quadruple events per time block follow exactly the same procedure, and we finally obtain:

$$t_{0,triple} = 24,$$

$$t_{1,triple} = 4, \tag{3.8}$$

$$t_{2,triple} = 0. \tag{3.9}$$

There is only one instance of a quadruple event. The "Top Level" and "Time Block" organizational levels have recently been analyzed in the context of Game Theory [47].

Actual Results

The constraints regarding the distribution of event classes in Fig. 6 are now established. Note, however, that the final distribution of these events is left to the composer. In other words, both the rows and the columns are interchangeable. For example, Xenakis could have chosen to place the single quadruple event in any timbral class, or any time block. Xenakis addresses this issue as follows [46]:

> We notice that the rows are interchangeable (= interchangeable timbres). So are the columns. This leads us to admit that the determinism of the matrix is weak in part, and that it serves chiefly as a basis for thought – for thought which manipulates frequencies of events of all kinds. The true work of molding sound consists of distributing the clouds in the two-dimensional space of the matrix, and of anticipating a priori all the sonic encounters before the calculation of details, eliminating prejudicial positions. It is a work of patient research which exploits all the creative faculties instantaneously. This matrix is like a game of chess for a single player who must follow certain rules of the game for a prize for which he himself is the judge. This game matrix has no unique strategy. It is not even possible to disentangle any balanced goals. It is very general and incalculable by pure reason.

Subject to the "rules of the game," Xenakis finally obtains the matrix of Fig. 6. How well does he follow it in his score? Leaving aside some confusion caused by a measure numbering error in the score, Xenakis places the various event classes in the correct cell, according to his matrix. However, there are some interesting "adjustments" that he makes:

1. The 6.5 measure time blocks do not begin and end rigidly. There is a lot of elision and overlap. However, this makes sense, both musically and stochastically. First, composers often elide phrases or provide transitions from one section to the next, and overlap entrances of a new instrumental

group with the fade-out of a previous group. Second, probability distributions themselves are not rigid, but curved (see Fig. 7). These distributions, at their extreme ends, tend towards zero, but never reach zero. There is always a possibility, albeit remote, that something will happen.

2. The columns η and $\kappa\delta$, in which no events are supposed to occur, do **not** consist of measures of rest. There are very few notes, but some activity, which is again allowed for in the "edges" of the probability distribution which lie below the single event.

3. Xenakis makes one radical departure from his scheme in the final time block $\kappa\eta$, in the *String Glissando* cell, which is supposed to have no activity. The violins and basses have contary motion *glissandi* with *tremolo*. The cello has a *tremolo* on its open A string. All strings play *sul ponticello*. This configuration has not previously been heard and presumably was added as an effective ending.

3.1.14 Level 4 - Macrocomposition III

Artistic Decisions

Xenakis now defines the "single event" as consisting of a sound density $\delta = 5$ sounds/measure, which corresponds to about 2.2 sounds/sec at MM 26 [46]. Thus, in a "single event" cell (consisting of 6.5 measures), there are on average 5 distinct entrances per measure. Xenakis notes that the upper limit for a normal orchestra is 10 sounds/sec, which, at the tempo of *Achorripsis* would correspond to 23 sounds or entrances per measure.

The "double event" would have a $\delta = 10$ sounds/measure, and so on, up to the "quadruple event", with $\delta = 20$ sounds/measure.

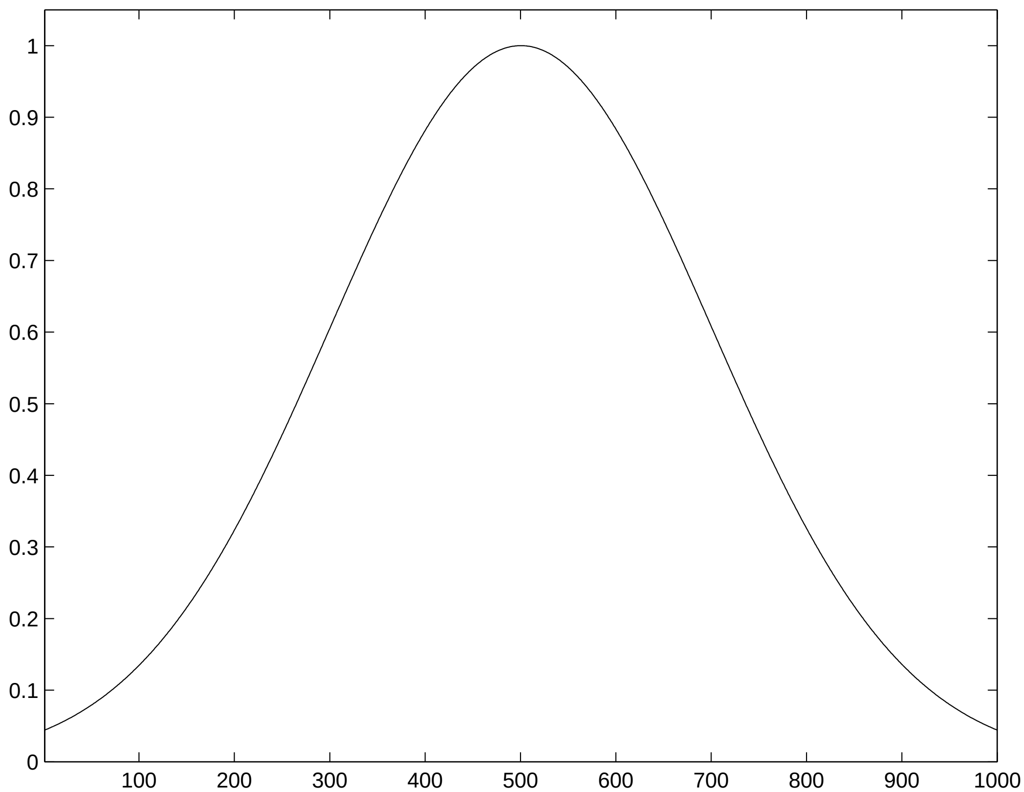

Figure 7: *A Gaussian Distribution*

Theory and Calculations

Having chosen an appropriate value of δ for the various event classes, no further theory or calculations are provided.

Actual Results

Without stating his procedure, Xenakis allows δ for a single event to range over $2.5 < \delta < 6.5$, distributed per the following histogram:

```
15 |                                         x
   |                                         x
   |                                         x
   |                                         x
   |                                         x
10 |                                 x   x   x
   |                                 x   x   x       x
   |                                 x   x   x       x
   |                                 x   x   x   x   x   x
   |                             x   x   x   x   x   x   x
 5 |                             x   x   x   x   x   x   x
   |                             x   x   x   x   x   x   x
   |                             x   x   x   x   x   x   x
   |                             x   x   x   x   x   x   x
   |                     x       x   x   x   x   x   x   x
 0  ----------------------------------------------------------
    0.0 0.5 1.0 1.5 2.0 2.5 3.0 3.5 4.0 4.5 5.0 5.5 6.0 6.5 7.0
```

where the horizontal axis represents values of δ and the vertical axis is the number of occurrences of a given value of δ (e.g. there is one occurrence of a single event cell in which $\delta = 2.5$). The average δ for the 65 single events is $4.94 \approx 5$, however Xenakis never states the parameters of the statistical distribution used. The values of δ chosen for each single event may be seen as a circled number in each single event cell in Figure 6. A similar undocumented procedure is used to choose values of δ for the double, triple and quadruple events.

3.1.15 Level 5 - Microcomposition/Implementation

At this stage, Xenakis turns his focus to the generation of events at the note level in each of the 196 cells of *Achorripsis*. He states the theoretical basis for the calculation of:

1. The time between successive events (i.e., notes).
2. The interval between successive pitches.
3. The "speed" of the *glissandi* in the string *glissando* cells.

Note that he does not address other aspects of the score such as:

1. The starting pitches for each instrument in each cell.
2. The duration of each note.
3. Dynamics.
4. Articulation. (String *arco* passages have some *staccato* notes, brass and woodwind have no articulation, and there are no accents).
5. The timbral choices mentioned in Section 3.1.11.

Artistic Decisions

Xenakis chooses the following statistical distributions:

1. The *exponential distribution* is used to govern the time between successive events.
2. The *linear distribution* is used to govern the intervals between successive pitches.
3. The *normal distribution* is used to govern *glissando* "speed."

Theory and Calculations

Exponential Distribution Squibbs has provided an excellent overview of Xenakis' general use of these statistical distibutions in his Ph.D. thesis [59], and has simplified some of the notation found in Xenakis [46]. Squibbs' versions will be used in this paper. The distribution for the time between successive notes is then:

$$P_i = e^{-\delta iv}(1 - e^{-\delta v}) \tag{3.10}$$

for $i = 0, 1, 2, \ldots$, v is the size of the time range and P_i is the probability that the time between events will fall within the given time range iv. Xenakis chooses a time range of 0.1 measure, which would be $\frac{12}{52}$ seconds, see Section 3.1.11. δ is the same as defined in Section 3.1.14.

As an example of this calculation, consider cell III, ιz, for which Xenakis provides sample calculations starting on p. 34 of [46]. This cell is a string *glissando*, single event cell, with $\delta = 4.5$ sounds per measure. Since every cell has 6.5 measures, there are $4.5 \times 6.5 \approx 29$ sound events (entrances), and therefore 28 time intervals between sound events. The probability P_0 that there will be time intervals in the range 0 - 0.1 measure (or between 0 and $\frac{12}{52}$ seconds) is:

$$P_0 = e^{-4.5\times0\times0.1}(1 - e^{-4.5\times0.1}) = 0.362 \tag{3.11}$$

Since there are 28 time intervals between sound events in this cell, there should be $0.362 \times 28 \approx 10$ intervals $\frac{12}{52}$ seconds or less. This calculation is the same as the first entry in the *Table of Durations*, [46], p. 34. Note that there is an obvious typographical error on p. 35, in which he states that $dx = 1/12.415 = 0.805$, which should be $= 0.0805$. (Thanks to the simplification of the Xenakis formulation by Squibbs, [59], p. 82, this number is not actually needed).

The reader may notice a similarity between Equations 3.10 and 3.2, in that both involve the exponential function $e^{-something}$. As Xenakis points out ([46], p. 324), Equation 3.10 is the product of two Poisson distributions: the probability that there will be no events in the time interval iv, times the probability that there will be one event in the time interval v. If one considers δ to be an average linear point density (the average number of points on a line, where the line represents time) then if these points are distributed along the line in a non-uniform manner, there will have to be there more shorter intervals than longer intervals between these points.

Linear Distribution Squibbs [59] provides a simplified version of Xenakis' formulation of the linear distribution, which governs the intervals between successive pitches in *Achorripsis*. This type of distribution is often applied to non-temporal situations, e.g., to determine if there is a correlation between some group of Mathematical Achievment Test Scores x whose average is $\overline{x}$ and the corresponding group Final Calculus Grades y whose average is $\overline{y}$ in a population of (say) 10 students [60]. One technique for assessing the correlation is to calculate the **Pearson product moment coefficient of correlation** r between these groups of values x and y as follows ([60], p. 448):

$$r = \frac{\sum_1^{10}(x-\overline{x})(y-\overline{y})}{\sqrt{\sum_1^{10}(x-\overline{x})^2}\sqrt{\sum_1^{10}(y-\overline{y})^2} \tag{3.12}$$

In Equation 3.12, if $r = 1$, the correlation is "perfect"and positive, that is, all points fall on a straight line with a positive slope. If $r = -1$ the correlation is also "perfect" but negative: the points all fall on a straight line with negative slope. A value or $r = 0$ indicates that there is no linear correlation. The closer $|r|$ approaches 1, the better the linear correlation.

To constrain the size the intervals between the pitches of successive string *glissando* entrances in Cell III ιz, Xenakis uses (in the Squibbs simplification,

[59], p. 86):

$$P_i = \frac{2}{n+1}(1 - \frac{i}{n}) \tag{3.13}$$

where $i = 0, 1, 2, \ldots, n$. $n = \frac{g}{v}$, where g is maximum interval size and v is the interval increment to be used in preparing the table of probabilities P_i. In Cell III ιz, Xenakis uses $v = 4.5$ semitones, and $g = 80$ semitones (the range of the orchestral strings, the lowest note being E1 - contrabass to C8 - violin, where C4 is middle C). Therefore $n = \frac{80}{4.5} \approx 18$. For example, the probability P_0 that the interval between successive pitches falls in the range 0 to 4.5 semitones is:

$$P_0 = \frac{2}{19}(1 - \frac{0}{18}) \approx 0.105 \tag{3.14}$$

Since there are $4.5 \times 6.5 \approx 29$ events in Cell III ιz, there will be $0.105 \times 29 \approx 3$ occurrences of intervals in the range 0 - 4.5 semitones. (There is actually a slight inconsistency here: if there are 29 *glissando* events then there can only be 28 intervals between these events. This will be discussed further).

Xenakis implies that the above values for g and v may apply only to this particular cell in the piece, for several reasons:

1. He may choose to restrict the pitch range g of the passage to a number smaller than the full range of the instrument family.

2. He may need to adjust g to the ranges of the other instrument families.

3. He may wish to correlate the pitch interval increment v with the sound density δ.

He addresses these issues as follows [46]:

> We shall not speak of the means of verification of liaisons and correlations between the various values used. It would be too long, complex, and tedious.

He does, however refer to Equation 3.12, with the implication that while he will make adjustments to g and v throughout the piece, there will still be a linear probability relationship between interval size, and the number of occurrences (the larger the interval, the lower the probability it will occur).

Normal Distribution Xenakis' use of *glissando* strings occurred first in *Metastasis* (*arco*) and next in *Pithoprakta* (*pizzicato*). In his analysis of *Pithoprakta*, he relates the distribution of "speeds" of the *glissandi* (change in pitch df divided by time increment dt, $\frac{df}{dt}$) to the distribution "speeds" of gas molecules, as derived by Maxwell/Boltzmann. He carries over the analogy to *Achorripsis*, thereby, in a sense, "mapping" the concept of velocity to string *glissandi*. It turns out that the distribution of speeds in a gas follows the Gaussian or Normal distribution, which is slightly more complicated mathematically than the Poisson, Exponential or Linear distributions. First, the probability density function $f(v)$ for the existence of a speed v is:

$$f(v) = \frac{2}{\alpha\sqrt{\pi}} e^{\frac{-v^2}{\alpha^2}}, \tag{3.15}$$

where α is defined as the "quadratic mean of all possible values of v" [46], p. 32, and is related to the temperature of the gas. Equation 3.15 does not yield the value of the probability directly. The area bounded by the x-axis, $f(v)$, vertical lines $x = v_1$ and $x = v_2$ is the probability $P(\lambda)$ that a given velocity v will fall within the range v_1 to v_2 ($v_2 > v_1$). The numerical value for this area may be obtained by integrating Equation 3.15 between the limits 0 and λ_1, and then again between 0 and λ_2, and subtracting the first value from the second. $\lambda_1 = \frac{v_1}{\alpha}$, $\lambda_2 = \frac{v_2}{\alpha}$:

$$P(\lambda) = \theta(\lambda_2) - \theta(\lambda_1), \tag{3.16}$$

where

$$\theta(\lambda) = \frac{2}{\sqrt{\pi}} \int_0^\lambda e^{-\lambda^2} d\lambda. \tag{3.17}$$

Equation 3.17, the Normal distribution, is usually evaluated by looking up in tables, which are available in most probability text books (e.g. [60]). The tables are presented in various versions, so care must be taken that the limits of integration, and the multiplier in front of the integral, correspond exactly with those in Equations 3.15, 3.16 and 3.17.

On p. 35 of [46], Xenakis presents a sample calculation of *glissando* speeds for the same sample Cell III ιz. He states that in this case $\alpha = 3.88$, where α is the "quadratic mean of the speeds." He further states on p. 32 that α (with a different value of 3.38, presumably a typographical error) is proportional to the sonic density δ, although he does not provide any expression.

To provide a sample calculation it is necessary to compare Equation 3.15 to the more common expression used in probability text books [60], p. 196:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(x-\mu)^2}{2\sigma^2}} \tag{3.18}$$

To use Table 3 in [60] it is necessary to choose values of σ, x and μ such that Equation 3.18 is equivalent to Equation 3.15. To do this, set:

$$\begin{aligned} x &= \frac{v}{2}, \\ \sigma &= \frac{a}{2\sqrt{2}}, \\ \mu &= 0. \end{aligned} \tag{3.19}$$

Then, evaluate

$$z = \frac{x - \mu}{\sigma} \tag{3.20}$$

where z is the equivalent of λ in Equation 3.17. For example, to calculate the value of $\theta(\lambda)$ corresponding to $v = 1$ in Xenakis *Table of Speeds*, [46], p. 35, using Equations 3.20, set $x = \frac{1}{2} = 0.5$, $\sigma = \frac{3.88}{2\sqrt{2}} = 1.3718$ to finally obtain, per Equation 3.20, $z = 0.3644$. Consulting Table 3 of [60], p. 654, and using linear interpolation between $z = 0.36$ and $z = 0.37$, one obtains a value of 0.1422. Xenakis chooses to work with the area under the entire normal curve, whereas Mendenhall, et al. [60] base their values on the one half of the area under the curve. To compare the Mendenhall, et al. value with Xenakis, therefore, we need to multiply the table value by 2 to finally obtain $0.1422 \times 2 = 0.2844$, which is quite close to the value 0.2869 in Xenakis for $\theta(\lambda)$.

Actual Results

The results of Xenakis' calculations will be compared in detail to the score, first in Cell III ιz, and then in Cell V α.

Cell III ιz The 6.5 measure cell starting at measure 105 is shown as Figures 9 and 10. The instruments participating in Cell III, the *glissando* strings, are Violin 3, Cello 2 and Contrabass 2. The hand-written, circled numbers are the *glissando* entrances in the order in which they occur. It will immediately be noted that events 1 and 2 occur before the barline of the measure 105, and that events 23, 28 and 30 persist past the first half of measure 111. This is typical throughout *Achorripsis* and is mentioned previously in this paper (see Section 3.1.13). The difficulty is to understand how Xenakis intended these "out of range" events to figure in his calculations. The first immediate problem is that there are 30 glissando events rather than the 29 calculated from the value of δ.

Leaving this problem aside, let us first address his calculation for the time between events. To do this, it is helpful to consider the note values he has chosen throughout the piece. In Fig. 8, the basic note values, and their duration expressed as a fraction with 52 as the denominator (to correspond to the tempo of MM = 52) are shown. To realize an exponential distribution of

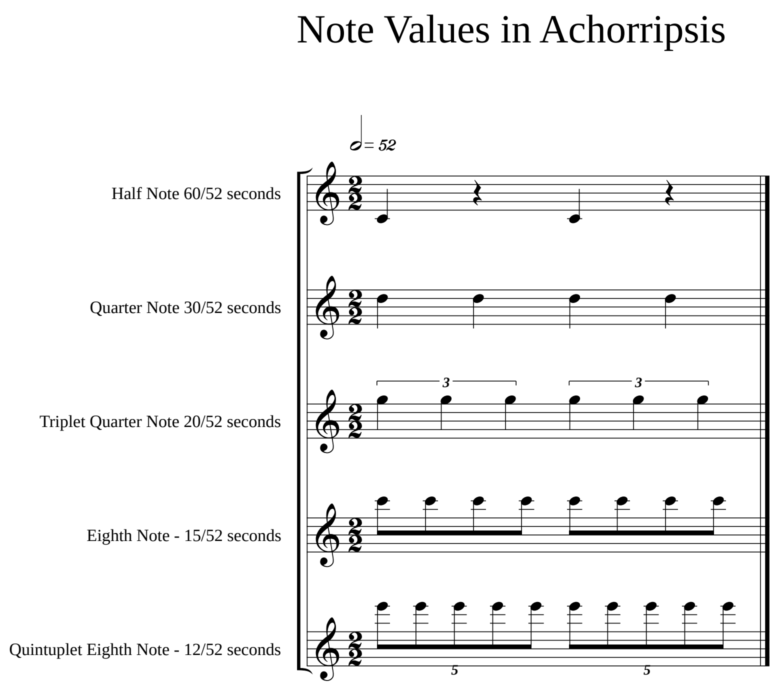

Figure 8: *Note Values in Achorripsis*

Figure 9: *Cell III ιz, Part 1*

Figure 10: *Cell III ιz, Part 2*

time between events, Xenakis draws from a rhythmic "palette" of 2 against 3 against 4 against 5. The smallest time between events is $\frac{3}{52}$ seconds, which occurs between successive entrances of an eight note and a quintuplet eighth note.

All relevant raw data from the score for this cell is shown in Table 1.

Measure No.	Event No.	t	Δt	p_s	Δp_s	p_e	Duration	v
1	1	-45		-29		-7	2.2083	9.96
	2	-36	9	9	38	9	0.8000	0.00
	3	60	96	10	1	12	0.6250	3.20
	4	80	20	16	6	17	0.6333	1.58
2	5	135	55	29	13	31	0.2500	8.00
	6	156	21	14	-15	15	0.9000	1.11
	7	165	9	44	30	43	0.2500	4.00
	8	195	30	9	-35	9	0.2500	0.00
	9	220	25	-24	-33	-23	0.5000	2.00
	10	225	5	37	61	37	0.2500	0.00
3	11	255	30	43	6	39	0.8750	4.57
	12	264	9	-3	-46	0	1.2000	2.50
	13	280	16	-32	-29	-31	1.1666	0.86
4	14	360	80	33	65	33	0.3750	0.00
	15	405	45	48	15	46	0.6250	3.20
	16	408	3	-3	-51	-4	0.8500	1.18
	17	420	12	-26	-23	-21	0.7000	7.14
5	18	480	60	27	53	27	0.1667	0.00
	19	500	20	21	-6	25	1.0833	3.69
	20	504	4	-23	-44	-23	0.4666	0.00
	21	510	6	-24	-1	-24	0.3500	0.00
	22	552	42	-22	2	-24	0.6000	3.33
	23	560	8	-31	-9	-29	2.3333	0.86
6	24	624	64	-5	26	-3	0.3333	6.67
	25	630	6	31	36	32	0.4166	2.40
	26	660	30	-10	-41	-10	0.2000	0.00
	27	680	20	1	11	1	0.3333	0.00
	28	684	4	22	21	18	1.3000	3.08
7	29	720	36	40	18	41	0.3750	2.67
	30	765	45	48	8	48	0.6250	0.00

Table 1: Raw Data from Cell III ιz

In the table heading, t represents the time that the event occurs, in multiples of $\frac{1}{52}$ seconds (in order to work with integer values). Δt is the time between events, in the same units. p_s is the starting pitch of the *glissando* note, in standard numerical pitch notation (each increment of 1 represents a semi-tone, with C4 - middle C - assigned to 0). Δp_s is the interval between successive starting pitches, positive if up, negative if down. p_e is the ending pitch of the *glissando*. In the score, it will be seen that this final pitch is represented by a grace note for absolute clarity. The heading "Duration" is the length of the *glissando* note in fractions of a measure. So for example, the duration of Event 7 is 0.2500 of a measure (ie., a quarter note). v is the speed of the *glissando* note, in units of semitones/measure. So in the case of Event 7, which starts on 44 and ends on 43, the speed $v = \frac{(44-43)}{0.2500} = 4.00$ semitones/measure.

In the following three tables, the actual Δt, Δp_s and v values will be tabulated and compared with the probability distributions, starting Δt in Table 2 The numbers in "Score" are obtained by counting the tabulated Δt, and the

Duration	Δt	Score	Exponential Distribution
0 - 12	9, 9, 5, 9, 3, 4, 6, 8, 6	10	10
12 - 24	20, 21, 16, 12, 20, 20	6	7
24 - 36	30, 25, 39, 36, 30	5	4
36 - 48	45, 42, 45	3	3
48 - 60	55, 60	2	2
60 - 72	64	1	1
72 - 84	80	1	1
84 - 96	96	1	0
	Totals	29	28

Table 2: Comparison of Theory with Data: Δt

numbers in "Exponential Distribution" are the same as in *Table of Durations*, [46], p. 34. Table 2 shows excellent, if not perfect, agreement between the theoretical distribution and the actual distribution in the score. The fact that there are more events in the score than in the theoretical table is related to the fact that throughout the score Xenakis chooses to wander outside the boundaries of his 6.5 measure blocks slightly. There is however another point to be made.

Probability distributions are an attempt to *estimate* the behavior of large numbers of events which occur over time. An individual instance of events (such as the distribution of time events in this 6.5 measure block) can only be expected to approximate to the theory. Xenakis and Squibbs refer to this phenomenon as "The Law of Large Numbers."

The next table shows the comparison of theory and score results for intervals Δp_s between successive starting pitches. Based on Xenakis' values of g and v, (see Equation 3.13), i ranges from $i = 0, \ldots, n$ where $n = 18$. Table 3 may be compared with *Table of Intervals*, [46], p. 36.

i	Pitch Range	Δp_s	Score	Exponential Distribution
0	0 - 4.5	1, 1, 2	3	3
1	4.5 - 9	6, 6, 6, 8	4	3
2	9 - 13.5	13, 9, 11	3	3
3	13.5 - 18	15, 15	2	3
4	18 - 22.5	21, 18	2	2
5	22.5 - 27	23, 26	2	2
6	27 - 31.5	30, 29	2	2
7	31.5 - 36	35, 33	2	2
8	36 - 40.5	38, 36	2	2
9	40.5 - 45	44, 41	2	2
10	45 - 49.5	46	1	1
11	49.5 - 54	51, 53	2	1
12	54 - 58.5		0	1
13	58.5 - 63	61	1	1
14	63 - 67.5	65	1	1
15	67.5 - 72		0	0
16	72 - 76.5		0	0
17	76.5 - 81		0	0
18	81 - 85.5		0	0
		Totals	29	28

Table 3: Comparison of Theory with Data: Δp_s

As with Table 2, agreement between the score data is generally good, with only 3 of the 18 data "slots" off by a count of 1.

The next table shows the comparison of theory and score results for *glissando* speeds, where v are the starting and ending values of "velocity" in semitones/measure for the "bins", and v_m are the mean values of velocity in each bin (i.e., the average of the starting and ending velocities). v_{score} are the values of *glissando* speed (velocity) tabulated in Table 1, in semitones/measure.

v	v_m	v_{score}	Score	Normal Distribution
0				
	0.5	0, 0, 0, 0.86, 0, 0, 0, 0, 0.86, 0, 0, 0	12	9
1				
	1.5	1.58, 1.11, 1.18	3	7
2				
	2.5	2, 2.5, 2.4, 2.67	4	5
3				
	3.5	3.2, 3.2, 3.69, 3.33, 3.08	5	4
4				
	4.5	4, 4.6	2	2
5				
	5.5		0	1
6				
	6.5	6.67	1	1
7				
	7.5	7.142	1	0
8				
	8.5	8	1	0
9				
	9.5	9.96	1	0
10				
		Totals	30	29

Table 4: Comparison of Theory with Data: v

Here there are significant discrepancies between the expected normal distribution (which Xenakis supplies in *Table of Speeds*, [46]) and the velocity values actually calculated from the string *glissandos* and their duration. For example, there are 10 instances of 0 velocity (i.e., no *glissando*, just a held note) and 2 additional instances of velocity between 0 and 1, in the score. for a total of 12 – as compared to the 9 from the Xenakis table. The velocity bins in the range 1 to 3 are depleted, with more instances of higher velocities in the range

7 to 9 than called for. It is curious that these discrepancies exist, especially where these values have been tabulated by Xenakis and he invites the reader to compare these values with the score (as has been done here, with the results shown in Table 4). One speculation is that the constraint of meeting both the exponential distribution for time between events, and the normal distribution for *glissando* speeds may have been too difficult (since Xenakis generated the actual notes without the use of a computer), and that he took care of the exponential distribution first, allowing all of the error to accumulate in the *glissando* speeds. It is impossible to draw a firm conclusion without consulting any notes which may be available in his archives.

Cell V α This cell occurs at the very beginning of the piece, in the *pizzicato* strings. There are no *glissandi*, hence only the time between events and the intervals between successive pitches can be compared to statistical distributions. The raw data for this cell is shown in Table 5.

In the next two tables, these raw score values of Δt and Δp_s will be compared with the exponential and linear probability distributions, respectively. Cell V α is a sparser section of the score with a density $\delta = 3.5$ sounds per measure so that there should be $3.5 \times 6.5 \approx 22$ events (rounding down) or 21 times between events Δt. Assuming that v, the size of the time range, is the same as in Cell III ιz ($v = 0.1$ measure, $\frac{12}{52}$ seconds) then we can calculate probablities for each time range just as in Equations 3.10 and 3.11. The comparison of theory and data is shown as Table 6, where it can be seen that the agreement is very good (with a discrepancy of only one in two time slots).

In comparing the pitch differences between entrances in Cell V α, one must assume that Xenakis intended to use the linear distribution, Equation 3.13 to govern these differences in other timbral classes besides *glissando* strings. Then,

Measure No.	Event No.	t	Δt	p_s	Δp_s
1	1	0		-28	
	2	24	24	12	40
	3	36	12	-21	-33
	4	40	4	1	22
	5	60	20	21	20
	6	80	20	-6	-27
2	7	160	80	-22	16
	8	180	20	-32	10
3	9	240	60	2	34
	10	288	48	23	21
	11	300	12	-2	-25
4	12	372	72	-18	-16
	13	420	48	16	34
	14	450	30	15	-1
5	15	492	42	8	-7
	16	504	12	-27	-35
6	17	600	96	5	32
	18	648	48	4	-1
	19	660	12	7	3
7	20	720	60	-21	-28
	21	750	30	-16	5
	22	760	10	10	26

Table 5: Raw Data from Cell V α

it is necessary to speculate as to the appropriate values of g, n and v. In Cell V α, the maximum pitch range is E1 to B5 (-32 to 23, ie. 55). So if we stick with $n = 18$, then $v = \frac{55}{18} \approx 3$. Based on this assumption, the values of the probability for each pitch difference range will remain the same, but the size of the ranges will be smaller. The total number of pitch difference events in this cell is 21. Comparing data with the linear distribution then yields Table 7.

This table seems to be at serious variance with the linear statistical distribution, calling into question that Xenakis intended that this distribution be used in other timbral classes besides the *glissando* strings.

3.1.16 Analysis Conclusions

A thorough analysis and comparison of theory with the score was carried out in two cells (Cell III ιz and Cell V α). The general conclusions are:

Duration	Δt	Score	Exponential Distribution
0 - 12	4, 10, 12, 12, 12, 12	6	6
12 - 24	20, 20, 20, 24	4	4
24 - 36	30, 30	2	3
36 - 48	42, 48, 48	3	2
48 - 60	60, 60	2	2
60 - 72	72	1	1
72 - 84	80	1	1
84 - 96	96	1	1
	Totals	21	21

Table 6: Comparison of Theory with Data: Δt

1. Xenakis was most rigorous is applying the exponential distribution to the time between events, and less so in applying the linear distribution to the intervals between pitches and the the normal distribution to *glissando* speeds.

2. Clearly his intent was to compose music, so some artistic adjustments were made to his distribution results. However, he followed them closely enough so that they can be usefully examined from the standpoint of sonification.

3. The preparation of the score was a remarkable feat considering that he worked without the help of a computer, but calculated all distributions, and their musical implementation, by hand.

Detailed examination of Xenakis' procedures is challenging and time-consuming, due in part to typographical errors in his text (which can eventually be figured out), and to the obscurity of some of his equations and accompanying mathematical language. To figure out his procedures, it is necessary at times to make a guess about what he was doing, follow through the guess, then if it doesn't work out, guess again. The obscurity may be in part due to the

i	Pitch Range	Δp_s	Score	Linear Distribution
0	0 - 3	1, 1	2	2
1	3 - 6	3, 5	2	2
2	6 - 9	7	1	2
3	9 - 12	10	1	2
4	12 - 15		0	2
5	15 - 18	16, 16	2	2
6	18 - 21	20, 21	2	2
7	21 - 24	22	1	1
8	24 - 27	27, 25, 26	3	1
9	27 - 30	28	1	1
10	30 - 33	33, 32	2	1
11	33 - 36	34, 34, 35	3	1
12	36 - 39		0	1
13	39 - 42	40	1	1
14	42 - 45		0	0
15	45 - 48		0	0
16	48 - 51		0	0
17	51 - 54		0	0
18	54 - 57		0	0
		Totals	21	21

Table 7: Comparison of Theory with Data: Δp_s

translator not knowing mathematical terminology, or Xenakis' overall impatience in documenting his own work, or lack of a technical editor who would have demanded clarifications. His biographer, Nouritza Matossian [58], had extensive interviews with the composer and eventually had complete access to his studio, in which he worked throughout his composing career in Paris. She described it as unheated, never cleaned, full of disarray with scribbled notes in shoeboxes, etc. She archived a lot of material to paint a picture of the man and his life. Someone else should go through his compositional notes to provide a user-friendly treatise on all of Xenakis' mathematical procedures.

It is obvious from his writings in *Formalized Music* that Xenakis was constantly expanding his knowledge of mathematics, and later, computer programming, far beyond what he would have learned as an engineer at Athens Polytechnic. His inventiveness in the application of his new-found knowledge to the composition was unique and profound.

3.1.17 Perception

The overall impression on the listener will be unique to the individual. However, the sound of *Achorripsis* is quite transparent. It is a sparse piece, and the probabilistic events can be clearly heard (especially if one listens with the matrix (Fig. 6)). *Achorripsis* could be perceived as a sonification of events occurring in real time. Each of the timbral classes could be considered as a separate data stream (for a total of seven). The data could represent any phenomenon governed by the probablility distributions similar to those chosen by Xenakis. In this scenario, every musical event would represent one data point or event. Each 15 second cell would then represent a group of data points which conform to a specific statistical distribution (for timing, pitch change, and (when specified) *glissando* speed in the strings.

It is probably easiest to first focus on the timing of the events, which is governed in *Achorripsis* by the exponential distribution, see Equation 3.10, and the associated explanation. It is critical to realize that while Equation 3.10 provides for the timing of events to fall within a continuum, the *Achorripsis* uses only discrete values, based on the durations of notes which he chooses to use in the score, see Fig. (8). The basic note values, and their durations, are expressed as the numerator of a fraction with 52 as the denominator (to correspond to the tempo of MM = 52). To realize an exponential distribution of time between events, Xenakis draws from a rhythmic "palette" of 2 against 3 against 4 against 5. The smallest time between events is $\frac{3}{52}$ seconds, which occurs between successive entrances of an eighth note and a quintuplet eighth note.

The theoretical and actual event timings for Cell Vα are shown in Table 6. The corresponding section of the score is shown as Fig. 11.

Achorripsis

Iannis Xenakis

Figure 11: *Cell Vα*

In the column of Table 6 labeled Δt, is a list of all the "times between events" which occur in Cell Vα (21 in total), tabulated by duration range. So, e.g., in the first row, the exponential distribution calls for 6 Δt's in the range 0 - 12. Examining the score, and calculating the time between entrances of the 22 notes, we find six discrete values of Δt in the range 0 - 12: 4, 10, 12, 12, 12, 12, where, e.g., 4 implies $\frac{4}{52} = 0.0769$ secs. This value is the time between the entrance of the $D\sharp$ in the Contrabass, and the $C\sharp$ in the Violoncello, in the first measure (third and fourth entrances in that measure).

If the musical constraints were absent, any 6 duration values in the (continuous) range 0 - 12 would still satisfy the exponential distribution. The use of discrete values, drawn from the "palette" of note values (Fig. 8), results in a discernible rhythmic signature, which corresponds to the density $\delta = 3.5$ (average number of events per measure) chosen for this cell (see Fig. 6).

Since all 89 "active" cells each have a distinct value for the density δ (the circled numbers in Fig. 6) it is tempting to speculate that the use of discrete note values makes it easier (or at least more enjoyable) for the listener to recognize, from the rhythm of the cell, what its density is. To test this speculation, a "game" has been set up, for the pizzicato sections, as a preliminary listening test, at the author's web site [61] in which the player will attempt to associate soundfiles from individual cells in *Achorripsis* with the displayed matrix cells. Soundfiles which satisfy the exponential distribution but use a continuous range of Δt, will be available for comparison.

A follow up to these tests would be to determine if it is easier for a listener to track two or more streams more easily with the discrete rhythmic configuration (Xenakis has, at most, six going simultaneously). How many different streams could the listener be expected to track? From a musical standpoint, *Achorripsis* can be performed under a good conductor, who is able to hear all of these events and determine whether or not a mistake has been made.

Xenakis' use of the linear distribution for the intervals between successive pitches is more difficult to correlate with the actual notes in the score. Furthermore, these distributions do not appear to be affected by parameters which change from one cell to the other. That is, other than constraints imposed by the ranges of the instruments, this linear distribution appears to be the same in all cells, and thus has a "neutral" influence, providing more of a vehicle for experiencing the rhythmic events.

Xenakis makes a useful contribution to the "mapping problem" by choosing to map the concept of velocity to *glissando*. In this mapping, the rate of change of the pitch (in, e.g., semitones per measure) is used to represent some sort of velocity (in, e.g., meters per second). This choice leaves open the possibility of representing a vector with magnitude and direction. The *glissando* speed represents the magnitude of the vector. A direction to the "right" could be represented by an upward *glissando*, to the "left," downward. "Up" and "down" could be represented by exponential shaping of the *glissando*. Once this has been done, the actual starting pitch of the glissando, the timbre and the register are still "free" to convey additional information.

The score of *Achorripsis* adheres closely enough to the composer's stated statistical distributions, especially in the time domain to be considered a useful contribution to the "mapping problem" in sonification. The work is challenging in that a lot of information is being conveyed in each cell, and so probably pushes the loading of sound to the limit. His musical rendering of these distributions may make them easier to grasp. Listening tests are needed to test this thesis. The mapping of *glissando* speed to velocity has useful properties for conveying vector quantities with both magnitude and direction, for example in the sonification of computational fluid dynamics [2].

3.2 Earth's Magnetic Field

Earth's Magnetic Field [62] [63] by Charles Dodge is probably the first sonification of scientific data realized with the use of a computer synthesis language (in this case, Music 4BF). Dodge was invited by geophysicists at the NASA Goddard Space Center to sonify data known as the K_p index, which measures changes in the earth's magnetic field due to its interaction with the solar wind. The data lends itself well to sonification because it only assumes one of 28 distinct values. Data for the years 1932 to 1961 have been compiled by Julius Bartels in the form of diagrams that resemble musical notation [64].

The K_p index is recorded at 3 hour intervals at 12 monitoring stations and then averaged. There are 8 values per day and 2920 values per year. Dodge chose the year 1961 and mapped the values of the index to 4 octaves of the diatonic C scale in the meantone temperament [65]. Thus, the pitches are wholly determined by the values of the index. Another characteristic of the data, known as "sudden commencements," refers to sudden increases in the index. Dodge used these commencement points as sectional markers in the piece. He plotted the highest value of the index occuring in each section (20 in all) vs. the length of the section. This graphical pattern was used, for example, to control tempo changes in the first part of the work.

A compelling part of *Earth's Magnetic Field*, when considered as a sonification, is interpretation. The physical significance of the magnetic field is that it provides protection against the intermingling of the solar wind with the planetary atmosphere, as well as shielding the earth from other harmful radiation. The mapping of K_p index to a diatonic scale gives the music a "benign" character. The lack of harmonic modulation provides a sense of stability. The timbral variation, spatialization and reverberation give an impression of "swirling radiance." Thus Dodge's mapping strategy, as well as his artistic shaping of the piece is appropriate and suggestive of the significance of the data. Despite the determination of all pitches from the Bartel's diagram, *Earth's Magnetic Field* reflects the personality of the composer and his compositional style.

3.3 The Red and White Cows

The Red and White Cows by Daniel Goode was composed in July 1979. The piece consists of two pages of typewritten instructions for the sonification of a mathematical problem, by the same name, proposed by J. V. Pennington [66].

> A rancher bought a white cow, and in the following year a red one. Each succeeding year he duplicated his purchases of the preceding two years, buying the same number of cows, of the same color and in the same order. Thus, in the third year, he bought a white and then a red cow; and so on. What was the color of the nth cow?

The first few years look like:

```
1, 2, 1-2, 2-1-2, 1-2-2-1-2, 2-1-2-1-2-2-1-2, etc.
```

where `1` stands for a white cow and `2` stands for a red cow. The commas delimit the years.

There does not appear to be much to sonify in this piece. The composer specifies that `2` is always higher in pitch than `1`, and that `2` is always two beats, and `1` is always one beat, at whatever tempo. Within a given year (i.e., between commas), there are to be no sudden changes in sounds or tempo. After a comma, a discrete change in one or more qualities (pitch, timbre, tempo, dynamics, articulation, density, harmony, emotional import) should occur. The sounds chosen for 1 and for 2 may be single or complex, pitch-like or noise-like.

The composer did an instrumentation for violin, viola and piano which was recorded and released on an LP entitled *Sleepers*. Other realizations have been done by other groups, including a recent performance on pitched percussion at Dartmouth College in 2003, with the composer present. *The Red and White Cows* is a fascinating instance of a sonification in which the performers, as well as the composer, share a role in defining the mapping of the numerical sequence to sounds.

3.4 Lottery Piece

Lottery by Nick Didkovsky, mentioned by Polansky in `http://music.dartmouth.edu/~larry/misc_writings/out_of_print/skin.intro.html`, and described in detail in `http://www.doctornerve.org/nerve/pages/lottery/lottery.shtml`, is unique in that it is a real-time sonification of a social experiment based on a class of decision-making problems called "prisoner's dilemmas." In such situations, the individual prisoner must make a choice in isolation. However, the collective outcome of these isolated choices has a profound effect on the group of prisoners as a whole. If the choices of the group as a whole are "cooperative," the whole group will fare well. However, if some members of the group elect to "defect" (i.e., make choices that would give them the upper hand), the "defectors" generally fare better than the rest of the group.

A special case of a prisoner's dilemma is the Luring Lottery, in which there is a maximum amount of prize money P, and participants may enter as many ballots as desired at no cost and with minimal effort. The twist is that the actual amount of prize money available is P/N, where N is the total number of ballots that have been entered. Thus, if one or more participants attempt to win by entering large numbers of ballots, P/N will be so small that for all intents none of the participants will receive any money at all. There are intermediate cases in which a participant might submit a modestly large number of ballots, say 100, and gain a reasonable sum from doing so provided that none of the

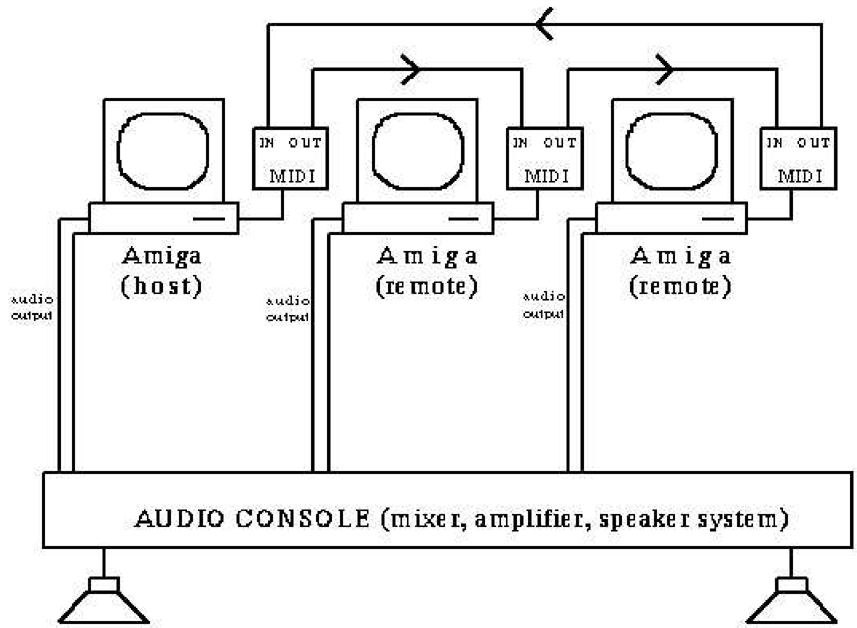

Figure 12: *Lottery Piece Computer Connection*

other participants did the same thing. A Luring Lottery was actually conducted in 1983 by Douglas R. Hofstadter, with a maximum prize of $1,000,000. The only requirement to participate was to write the number of ballots desired on a postcard and mail it to the specified address. The results were disappointing, in that many wrote extremely large numbers on their postcards, such that P/N vanished and no money was distributed.

Didkovsky decided to "introduce a Luring Lottery as a coercive device in a musical performance," to explore the possibility of behavior adjustment to a more cooperative state. *Lottery* is a live piece which was performed once on April 1, 1990 by Didkovsky and three other performers, on four Amiga computers connected together via a MIDI protocol (see Fig. 12). Each participant is allowed to edit the spectral characteristics of four independent waveforms (see Fig. 13). Participants are also allowed to "vote" by submitting some number of ballots for some collection of the available 16 waveforms, which constitute a

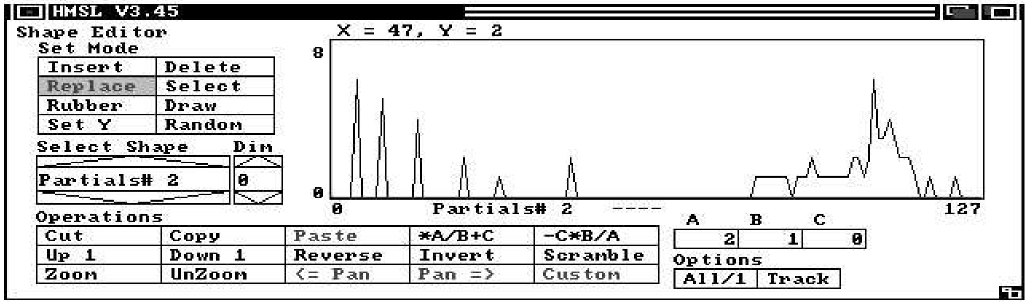

Figure 13: *Lottery Piece Waveform Shaping*

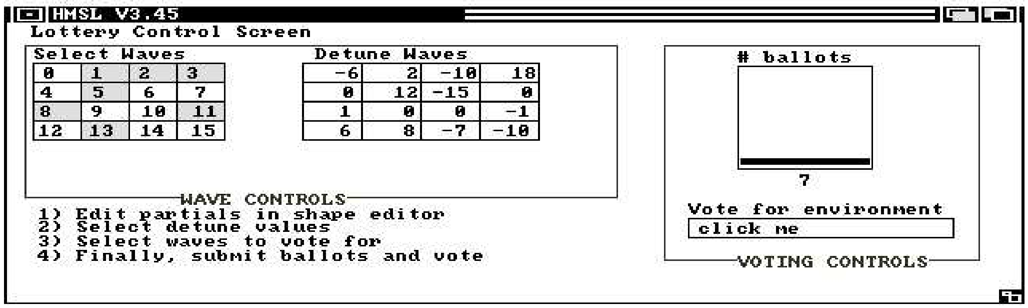

Figure 14: *Lottery Piece Control*

Sound State (see Fig. 14). A key aspect of *Lottery* is that the Luring Lottery repeats, or turns over, every 30 seconds. Each participant has access to a voting history screen, and is able to see how many votes were cast, and by whom, in the immediately previous lottery (see Figure 15). The person who "wins" each lottery is able to control the Sound State for the ensuing 30 seconds; however,

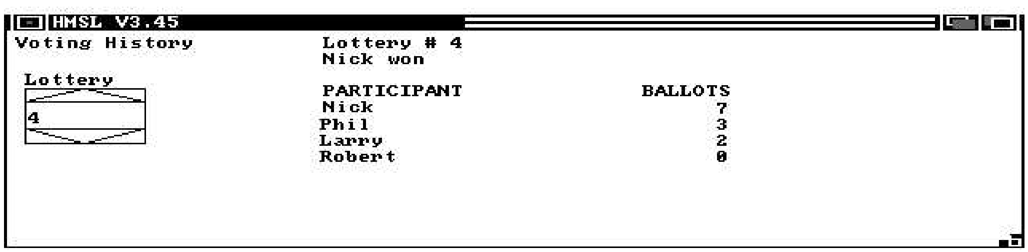

Figure 15: *Lottery Piece Voting History*

the amount of alteration that the Sound State actually undergoes is inversely proportional to the total number of ballots cast. If a large number is cast in a given lottery, the Sound State will be minimally changed. Thus, the sonic equivalent of vanishing prize money is monotony. Didkovsky points out one flaw in the sonfication scheme: the participants may vote with a reasonable number of ballots for the *same* Sound State, and this choice would have a similar effect on the sound.

4 Sonification Case Studies

In this chapter, some of the practical and technical aspects of sonification will be discussed, followed by a presentation of some case studies. Each case study illustrates a different application of sonification. The first example uses sonification for the purpose of monitoring complex, numerical processes, in this case, the solution of computational fluid dynamics (CFD) simulations. The second example is a sonification of historical data, in which the "playback" speed is increased so the data which was originally collected in the course of one year could be "heard" in approximately 20 minutes. The third example uses sonification in a real-time situation, in this case, financial trading, where there is too much data to be absorbed through visual displays alone.

4.1 Technical Aspects

In the ten years since the first ICAD conference in 1992, significant improvements in the audio capabilities of personal computers have been realized. The main result of these improvements is that many sound synthesis techniques can be accomplished "in real time," which is particularly relevant to the sonification of real-time data streams.

An all-purpose platform for designing and realizing sonifications does not currently exist. A specification for such a platform might include:

1. Data Inputs
    (a) Real-Time data from any source (sockets, data busses, XML, etc.)
    (b) Flexible file-handling, importing.
    (c) Interchange with other applications (spreadsheets, DDE, etc.)
2. Data Processing

    (a) Statistical analysis, range determination, normalization, etc.

3. Mapping Tools

    (a) Definition of Data Scenarios

        i. Movement monitoring of one or more variables.

        ii. Approach of one or more variables to one or several targets.

        iii. Interactive activity (online transactions or negotiations).

        iv. Feature search in complex data.

        v. Pattern search in historical data.

        vi. Process monitoring.

        vii. Global expectation vs. reality.

    (b) Wizard-style interface to guide the user through choices of:

        i. Musical style

        ii. Instruments

        iii. Choice of sonification schemes

4. Sound Generation Methods

    (a) Sample manipulation.

    (b) Synthesis

        i. Additive

        ii. Subtractive

        iii. FM

        iv. Physical Modeling

        v. Granular

        vi. Distortion

vii. Sinusoidal/Noise

(c) Spectral Manipulation

i. Phase Vocoding and other FFT techniques

ii. Cross Synthesis

5. Sound Output Methods

(a) Speakers, Headphones

(b) Spatialization (HRTF, VBAP, etc.)

(c) Direct-to-disk recording

4.1.1 Programming Environment

Several proposals for the development of a sonification design environment have been made. Kaltenbrunner [67] has drafted a framework for a shared auditory user interface (AUI) which would incorporate auditory widget libraries that mimic the X-Windows GUI framework in the Unix system. The basic idea would be to allow multiple sound widgets or sound APIs to have access to a *sound server*. The framework could be implemented in the Java programming language so as to ensure portable code for the high-level interfaces. Kaltenbrunner coins the phrase "hear and feel" to suggest an analogy with the expression "look and feel," which is often used to refer to GUIs. Kaltenbrunner's recommendations are embryonic, intended mainly to stimulate discussion.

Ben-Tal, et al. [68] have proposed *SonART* as a platform-independent set of methods to map data to sonification parameters. *SonART* proposes to use STK as its underlying audio engine [69]. The Synthesis Tool Kit (STK) is a collection of classes in C++ intended for the rapid creation of sound synthesis and audio processing systems. *SonART* is still in the design/conception stage, and is not currently available. Other proposed sonification tools are *MUSART*

[70], *Listen* [71] and *Personify* [72]

In the absence of high-level sonification toolkits, in which the user could focus mainly on mapping, the next best strategy is to use existing API's or synthesis software. Established languages such as Max/MSP (`www.cycling74.com`), CSound [73] and KYMA (`www.symbolicsound.com`) are primarily designed for the composition and performance of electro-acoustic music. These programs provide a variety of synthesis techniques to the composer and sound designer. KYMA was originally put forward as a tool for sonification design [5] and is currently used for some data sonification projects in which CPU-intensive synthesis or other processing is required. Early versions of KYMA had, for example "1st order - data to frequency and amp" and "data to distortion" objects. Kramer used an early version of Max to realize an early Sonification Toolkit [74]. Csound has been used as a sonification tool in conjunction with "R," one of the most powerful and developer friendly statistics packages [75].

The main limitations of these music languages for general-purpose sonification are the difficulty of implementing sonification-specific interfaces, lack of built-in tools for dealing with data and file streams, and (in the case of Csound) the processing time for building the soundfile (which is a problem for sonifying real-time data). APIs such as JSyn (`www.softsynth.com`), JMSL (`www.algomusic.com`), SMS (`http://www.iua.upf.es/~sms/`) and STK [69] have the advantage that the rich array of data-processing classes (in Java or C++) are available to the programmer. Implementation of the sound objects is considerably more work in these APIs, however, Burk has implemented a graphical interface in JSyn called "Wire" which enables the interconnection of objects in a manner which resembles Max/MSP. The resultant instrument may then be output to Java source code as an extension of a `SynthNote` or `SynthCircuit` object, and incorporated as a Class into the main program.

All of the sonifications in this thesis project have been implemented in JSyn. In my opinion, the extra effort required to set up the sound synthesis objects is more than offset by the ease of piping data from a wide variety of data sources to the sound object input parameters. The Java object-oriented design allows easy encapsulation of existing objects into more complicated objects, or the extension of existing objects to incorporate new features [76]. It is also easy to design user-interfaces using “Visual Studio” style rapid application development (RAD) tools in integrated development environments (IDEs) such as Metrowerks Code Warrior (`www.metrowerks.com`).

One drawback of using Java for sound applications is its slowness at sample-level manipulations. This limitation is being reduced as new versions of the Java Sound API are released (`http://java.sun.com/products/java-media/sound/`. The basic problem is that Java is compiled and linked into machine-independent “bytecode,” which must then be “interpreted” by the Java Run Time system specific to each computer platform. The issue can be circumvented by compiling computer-intensive sample level routines into a native library, which are then called from the main Java program using the Java Native Interface (JNI, `http://java.sun.com/docs/books/tutorial/native1.1`). Burk has implemented this approach with JSyn, using the C-based code CSyn.

JSyn provides a set of unit generator, sound in, sound out, filter, reverberation and delay objects, as well as sample readers which can handle both `.wav` and `.aiff` formats. Several higher-level classes, such as FM instrument design and granular synthesis, are provided as examples. Synthesis techniques involving spectral manipulation or physical modeling are not available. Such techniques, at the present time, may have some limitations in dealing with real-time data, although recent advances have been made in this area, for example ATS [77].

4.1.2 Spatialization

The author is indebted to his colleague, Dr. Ville Pulkki at the Helsinki University of Technology, who provided the information, the references and some of the text in this section 4.1.2. The material is herein incorporated with the permission of Dr. Pulkki. This information, in a condensed form, has been accepted as a paper at the International Conference on Auditory Display, July 2003 in Boston [3].

Spatialization can be used in sonification as an additional mapping dimension. In some cases, the perceived direction of a sound source corresponds directly with the location of the data. If not, it is nonetheless true that when data is sonified immersively around the listener, more auditory information may be decoded, compared to monophonic or stereophonic methods. The physical location of data to be sonified is frequently as important as the value of the data itself. A significant aspect of the "mapping" problem then becomes how to represent different spatial locations in auditory space. Mapping "space to space" by applying the techniques of 3D sound spatialization to data sonification seems a natural solution to this problem. The direction of the sound source itself thus contains information in a straightforward manner, e.g., the geographic location of a hail storm may be mapped to the direction of a virtual source.

Unlike visual displays, which must be presented directly in front of the user, sounds may placed around the listener in spatial auditory displays. When the listener hears sounds from different directions, it is easier to parse an auditory scene consisting of several simultaneous sounds (the cocktail party effect). When multiple sounds are presented only monophonically or monaurally they mask each other, and the listener cannot distiguish between them. So, more information may be transmitted in spatialized displays.

There are many technical approaches to audio spatialization, based either on the use of headphones or loudspeakers. Loudspeakers can be used with two or more channels. Some properties of existing spatialization systems will be briefly reviewed. The most desirable system for sonification would possess a relatively stable spatialization quality with relatively low numbers of loudspeakers in a large listening area, so that multiple listeners would perceive as similar a 3-D soundscape as possible. The VBAP (Vector Base Amplitude Panning) method was chosen, since it fulfills these requirements quite well, and because of its applicability to arbitrary numbers and locations of speakers.

Directional Hearing

Spatial and directional hearing have been studied intensively, for overviews, see for example [78]. The duplex theory of sound localization states that the two main cues of sound source localization are the interaural time difference (ITD) and the interaural level difference (ILD) which are caused, respectively by the wave propagation time difference (primarily below 1.5 kHz) and the shadowing effect by the head (primarily above 1.5 kHz). The auditory system decodes the cues in a frequency-dependent manner. The monaural spectrum is also used to decode the elevation of sound source, and the effects of head rotation to ITD and ILD also help in localization.

The resolution of directional hearing is in optimal conditions with only one sound source about 1° at its best, and about 20° at its worst, depending on the sound source direction. However, when there are multiple sound sources presented at the same time, the resolution drops to at best 20°. An important phenomenon in directional hearing is the precedence effect. When a coherent sound arrives from multiple directions within a short (< 5-30 ms) time window, the direction of the first arrival dominates in direction decoding. This is a complicated phenomenon that is summarized in [79].

Spatialization Methods

The intent of spatial sonification is to synthesize spatial impressions; not to recreate spatial sound that existed and was recorded on some occasion. Different methods for the positioning of virtual sources are briefly reviewed. We are primarily concerned with the production of immersive 3-D sound stages, that would be audible by several people in a large listening area with about 10 loudspeakers. Therefore, headphone methods such as HRTF processing were not considered. Wave field synthesis requires a much larger number of loudspeakers and was not used.

HRTF Processing

Headphone Reproduction In headphone listening a monophonic sound signal can be positioned to sound as though it is coming from virtually any direction, if the HRTFs (head-related transfer functions) for both ears are available [80, 81]. The audio signal is processed by a digital filter which models the measured HRTF. The method simulates the ear canal signals that would have been produced if a sound source had originated from a particular direction. If a listener moves her head during listening, then the movements should also be taken into account in processing, otherwise the sound field will seem to move along with the listener.

Cross-Talk Cancelled Loudspeaker Listening If, in a stereophonic setup, the cross-talk between a loudspeaker and the contralateral ear is cancelled by some method, it is possible to control the sound signal arriving at a listener's ear canals [80]in a precise manner. However, the best listening area (sweet spot) is very small. Cross-talk cancelled reproduction is otherwise similar to HRTF techniques over headphones, with the addition of room transfer functions.

Amplitude Panning

Amplitude panning is the most frequently used panning technique. Panning is achieved by applying the same audio signal to one or more speakers, but with different amplitudes for each speaker. The amplitudes are controlled with gain factors denoted by g.

$$x_i(t) = g_i x(t), \quad i = 1, \ldots, N, \tag{4.21}$$

where $x_i(t)$ is the signal to be applied to loudspeaker i, g_i is the gain factor of the corresponding channel, N is the number of loudspeakers, and t is the time parameter. As a result, the listener perceives a virtual source from a direction which is a function of the gain factors. When all the loudspeakers are on a horizontal plane, pair-wise panning is often used, in which sound is applied to two adjacent loudspeakers [82]. If elevated or descended loudspeakers also exist, triplet-wise panning can be utilized, in which an audio signal is applied to three loudspeakers which form a triangle from listener's view point.

Ambisonics

Ambisonics is basically a microphoning technique; however, it can also be used to synthesize spatial audio as an amplitude panning method [83]. In this case it is an amplitude panning method in which the audio signal is applied to *all* loudspeakers placed evenly around the listener with gain factors

$$g_i = \frac{1}{N}(1 + 2\cos\alpha_i), \tag{4.22}$$

where g_i is the gain of ith speaker, N is the number of loudspeakers, and α is the angle between loudspeaker and panning direction. The audio signal therefore emanates from all loudspeakers. Second-order Ambisonics applies the audio to a similar loudspeaker system [84] with gain factors calculated from

$$g_i = \frac{1}{N}(1 + 2\cos\alpha_i + 2\cos 2\alpha_i). \tag{4.23}$$

The sound is still applied to all of the loudspeakers, but the gains have significantly lower magnitudes on the side which is opposite the panning direction. This creates fewer artifacts. To get an optimal result, the loudspeakers must be in a symmetric layout. Increasing the number of loudspeakers beyond a certain number, however, no longer enhances the directional quality. The main problem with ambisonics is that the same audio signal appears on all of the loudspeakers. The directional quality degrades quickly outside the sweet spot, since the virtual source appears to emanate from the loudspeaker nearest the listener due to the precedence effect [79]. Also, even in the best listening position, the directional quality is not optimal, since the directional cues produced with ambisonics may differ significantly from the real source cues [85].

Vector Base Amplitude Panning

Vector base amplitude panning (VBAP) is a method for the calculation of gain factors for pair-wise or triplet-wise amplitude panning [87]. In pair-wise panning, VBAP is a vector reformulation of the tangent law [88]. Unlike the tangent law, however, it can be generalized easily for triplet-wise panning. In VBAP the listening configuration is formulated with vectors; a Cartesian unit vector $\mathbf{l}_n$ points to the direction of loudspeaker n, relative to the position of the listener. In triplet-wise panning unit vectors $\mathbf{l}_n$, $\mathbf{l}_m$, and $\mathbf{l}_k$ then define the directions of loudspeakers n, m, and k, respectively. The panning direction of a virtual source is defined as a 3-D unit vector $\mathbf{p} = [p_n \ p_m \ p_k]^T$. A sample configuration is presented in Fig. 16.

The panning direction vector $\mathbf{p}$ is expressed as a linear combination of three loudspeaker vectors $\mathbf{l}_n, \mathbf{l}_m$, and $\mathbf{l}_k$, in matrix form:

$$\mathbf{p}^T = \mathbf{g}\mathbf{L}_{nmk}. \tag{4.24}$$

Here g_n, g_m and g_k are gain factors, $\mathbf{g} = [g_n \ g_m \ g_k]$ and $\mathbf{L}_{nmk} = [\mathbf{l}_n \ \mathbf{l}_m \ \mathbf{l}_k]^T$.

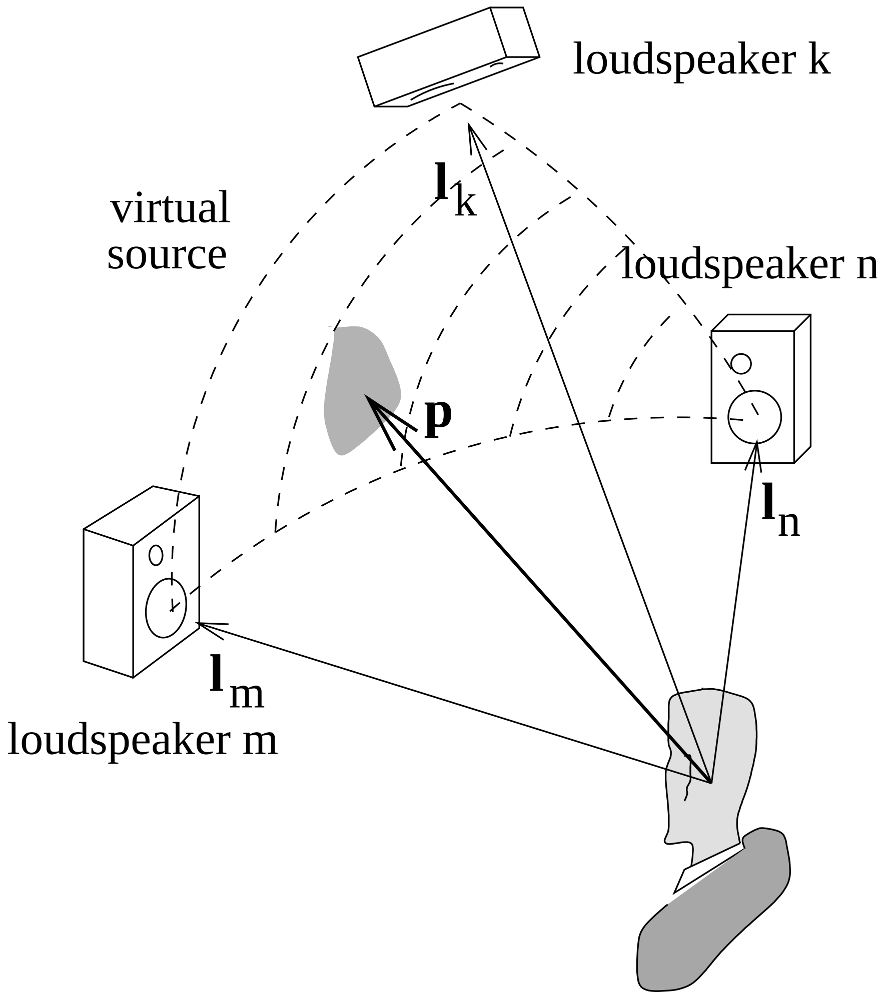

Figure 16: *A loudspeaker triplet forming a triangle formulated for three-dimensional vector base amplitude panning (VBAP) [86].*

Vector $\mathbf{g}$ can be solved

$$\mathbf{g} = \mathbf{p}^T \mathbf{L}_{nmk}^{-1} \tag{4.25}$$

if $\mathbf{L}_{nmk}^{-1}$ exists, which is true if the vector base defined by $\mathbf{L}_{nmk}$ spans a 3-D space. Eq. 4.25 calculates barycentric coordinates of vector $\mathbf{p}$ in a vector base defined by $\mathbf{L}_{nmk}$. The components of vector $\mathbf{g}$ can be used as gain factors; a scaling of them may be desired. When more than three loudspeakers are present, the loudspeaker system is triangulated, and one triplet is used at one time for panning.

The perceptual quality of virtual sources produced with VBAP, or more generally with amplitude panning has been studied in [86]. It was found that VBAP accurately predicts the azimuthal angle between the median plane and virtual source direction. Perception of the elevation angle was found to be more dependent on the individual. However, with triplet-wise panning, the perceived direction of virtual source is generally inside the loudspeaker triangle. This holds even if the listener is not located in the sweet spot.

Wave Field Synthesis

When the number of loudspeakers is very large, Wave Field Synthesis [89] can be used. It reconstructs a whole sound field in the listening room. Theoretically it is superior as a technique, but unfortunately it is impractical in most situations. The most restricting boundary condition is that the system produces the sound field accurately only if the loudspeakers are at a distance of maximally a half wavelength from each other. The centroids of loudspeakers should thus be a few centimeters from each other to be able to produce high frequencies correctly also, which cannot be achieved without a very large number of loudspeakers.

Java VBAP Implementation

Wave-field synthesis has proven to be the most accurate spatialization method. However, for 3-D rendering, tens of thousands of loudspeakers are typically required. In this project, two listening rooms with 8 and 12 loudspeakers with different loudspeaker layouts were available. VBAP was therefore chosen to be used as the spatialization scheme, since it performs relatively well with loudspeaker arrays of modest size even outside the best listening position, and can be easily configured for different loudspeaker setups.

JSyn is a Java API which provides an extensive toolkit of unit generators, filters, effects, mixers, etc. which may be encapsulated and customized by the user [90]. The low-level synthesis is carried out by native-compiled C subroutines. It is convenient to design sonifications in this environment because the necessary input/output, GUI design, and analytical calculations and sound generation may all be implemented in the same project. The existing VBAP C subroutines were accessed from Java via the Java Native Interface (JNI).

In the current VBAP implementation, an automatic method is used to divide the loudspeaker setup into pairs or triplets based on loudspeaker directions. An automatic triangulation was also adapted based on existing code [91]. It selects the triangles based on the greedy triangulation method. The method outputs the inverse matrices and other data needed in VBAP gain calculation. The system can be thus adapted to a new loudspeaker setup simply by defining the number and directions of loudspeakers in it.

In the Java implementation, a single VBAP object is instantiated using the loudspeaker location data (azimuth, elevation). The desired azimuth and elevation of a single virtual source is calculated from the relevant data and passed to the VBAP object, which then returns the gain factors for each loudspeaker. These are then used, via a standard JSyn `SynthMixer` object, to control the gains on the output channels.

Realization

HUT Acoustics Lab Listening Room The spatial sonification tests, using the meteorological data (see Section 4.2.2) were first realized in the Helsinki University of Technology Acoustics Lab Listening Room (Fig. 18) during April-May of 2002, using 5 Genelec speakers in the horizontal plane, at azimuthal positions of $30°$, $-30°$, $90°$, $-90°$ and $180°$, and 3 Genelec speakers at an elevation of $45°$, at azimuthal positions of $40°$, $-40°$ and $180°$ and a Genelec subwoofer. As may be seen in the figure, the elevated speakers were positioned through the use of

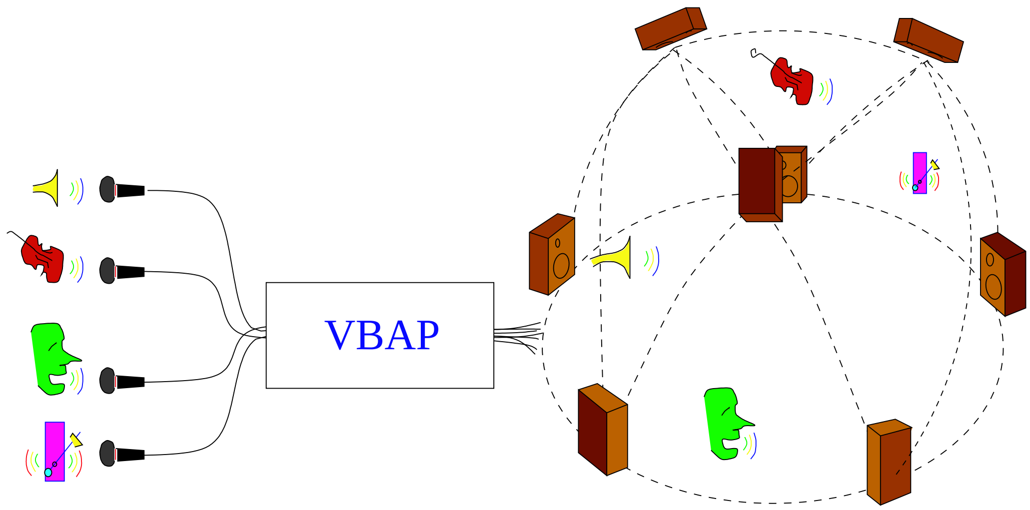

Figure 17: *Loudspeaker Array [86]*

Figure 18: *HUT Acoustics Lab Listening Room*

microphone stands.

Bregman Electronic Music Studio Listening Room After the mapping scheme for the meteorological data (see Section 4.2.2) was modified, and a user interface was added, the sonification was realized a second time, during January 2003, in Bregman Electronic Music Studio listening room. The system consists of a 12 speaker system, using a Macintosh G4 computer, 2 Motu 896 firewire interfaces, 8 Hafler Trans-nova reference monitors on a horizontal 8 loudspeakers on the horizontal plane at azimuthal positions of $15°$, $-15°$, $45°$, $-45°$, $90°$, $-90°$, and $150°$, $-150°$, and 4 Yamaha loudspeakers an elevation of $45°$ at azimuthal positions of $0°$, $90°$, $-90°$ and $180°$.

4.1.3 Timbral Control

Choosing “natural” and “pleasant” sounds is an important consideration in designing sonifications for use in practical situations. One aspect of “naturalness” is pleasing timbre or tone quality. Chowning and others have observed that the time evolution of the spectral components in a sound determine, to a large degree, its timbre [92] and “liveliness.” It is also easier to detect the location of spectrally rich sounds in a spatial sonification [79].

Systematic control of timbre has been cited as a desirable feature of any sonification system [19]. In this study, the implementation of an FM formant instrument was performed using the JSyn API, allowing data-driven manipulation of timbre in addition to other musical parameters such as pitch, envelope, duration, loudness, tempo, etc.

A three oscillator instrument consisting of one modulator and two carriers was chosen for the sonification, (see Fig. 19). The inputs available for manipulation are labelled at the top of each unit. The timbre of the instrument may be adjusted via the manipulation of five parameters:

1. The index of modulation I.
2. The ratio of the frequency of Carrier 1 c_1 to the modulating frequency m (c_1/m).
3. The ratio of the frequency of Carrier 2 c_2 to that of Carrier 1 (c_2/c_1).
4. The ratio of the "depth" of Carrier 2 d_2 to that of Carrier 1 d_1 (d_2/d_1). If, for example, this ratio is less than one, the frequency deviation of Carrier 2 will be less than that of Carrier 1.
5. The ratio of the amplitude a_2 of Carrier 2 to that of Carrier 1 (a_2/a_1).

This instrument was originally proposed by Chowning [92] as a means of introducing a formant peak into the spectrum. It was implemented in JSyn by encapsulating three `FMOperator` objects into a `SynthNote` object. Each `FMOperator` is controlled by an separate envelope, so that the time evolution of the modulation index, the formant peak and the fundamental tone (Carrier 1) can be independently controlled.

4.2 Case Studies

The following three sonification projects were carried out at the Bregman Electronic Music Studio from January 2001 to May 2003.

4.2.1 Computational Fluid Dynamics

History

The non-linear partial differential equations governing the conservation of mass, momentum and energy in fluids were derived from first principles in the first half of the nineteenth century by J. Navier [93] and G. Stokes [94]. Except for a few very restricted cases (e.g., fully-developed laminar flow in a duct), these equations could not be solved by analytical methods, and thus remained

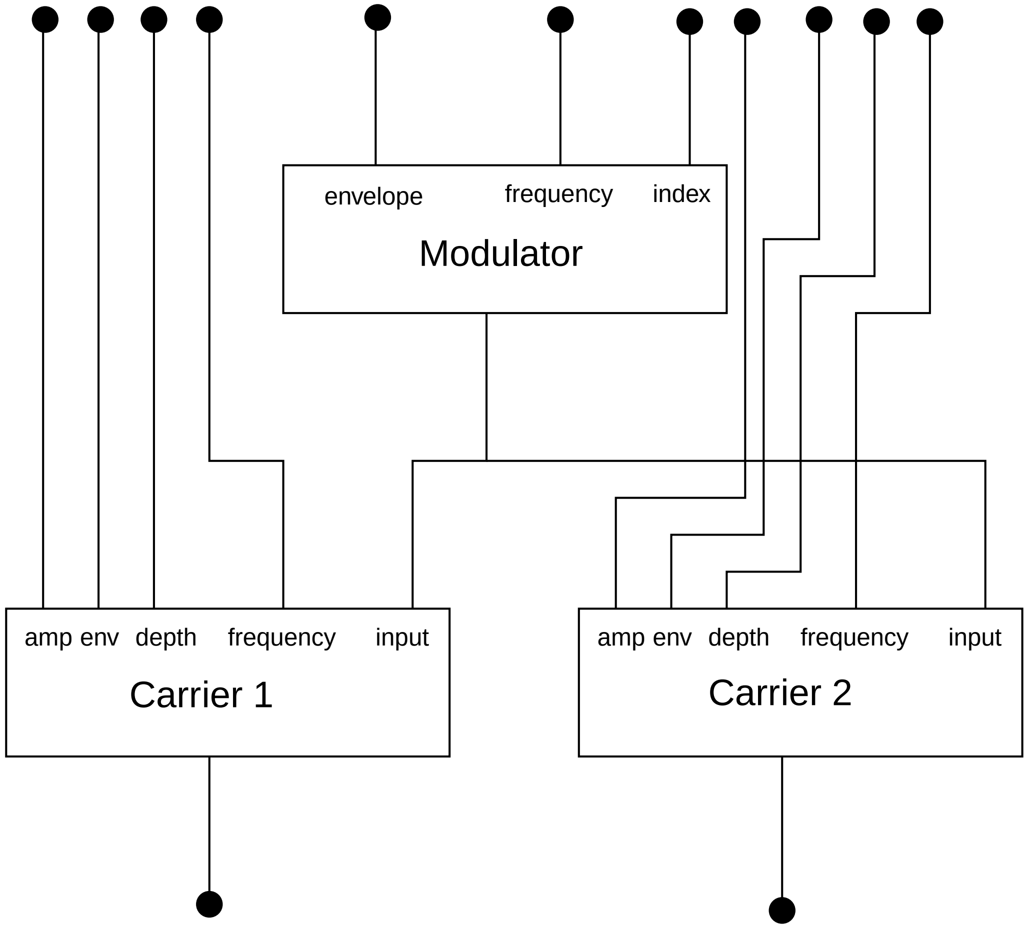

Figure 19: *FM Formant Instrument*

a mathematical curiosity until numerical methods on high speed computers became available in the second half of the twentieth century [95]. During the past twenty years, general-purpose CFD software has emerged as a practical tool for applying the Navier-Stokes equations to the solution of realistic fluid flow problems in engineering and physics. Today, several commercial CFD packages are routinely used by engineers and scientists in such diverse fields as hyrdrodynamics, aerodynamics, biomedical engineering, process industry, heating, ventilating and air conditioning, environmental engineering, etc.

The CFD Process

A typical CFD analysis is carried out in six stages:

1. The complex, real-world situation to be analyzed is reduced to a practical CFD project based on:
    - The limitations of the CFD model being used.
    - The available time and computational resources.
    - Engineering judgment as to what details of the flow field are essential to the analysis.
2. The *geometry* of the region of interest is either imported from a CAD package or constructed from scratch in the CFD package. Imported geometries often contain details which are extraneous to the CFD analysis and must either be modified or removed.
3. A *computational grid* is generated which must generally satisfy the following constraints:
    - The total number of grid points must not be so large as to overwhelm the limitations of CPU storage and speed.
    - The grid must be fine enough to resolve the details of the flow field which are of interest to the user.
    - The characteristics of the grid must be compatible with the solver: there can be no sudden changes in cell size and cells may not have too high an aspect ratio or be extremely skewed.
4. *Boundary conditions* must be applied to all regions of the computational domain and the *physical properties* of the fluid(s) must be specified.
5. Various solution control parameters and solver options must be set. The *solver* is then started and must be monitored until a converged solution is achieved.

6. The results are then *post-processed*, sometimes within the CFD package itself, or else exported to a data visualization package.

This cycle is often repeated several times before a final, satisfactory result is obtained.

CFD and Sonification

To the author's knowledge the above-described process is currently carried out entirely in the visual domain. There are no commercial CFD packages which make use of sound to enhance the interaction between the engineer and the data.

Furthermore, very little research into sonification and CFD has been published. McCabe and Rangwalla [96] presented two examples of auditory display: the simulation of an artificial heart pump and rotor-stator interaction in turbomachinery. In the first, MIDI sound was used to enhance the post-processing of the artificial heart simulation, in particular to signal global changes in the system such as the opening and closing of a heart valve. In the second, the time-varying pressure field predicted by the model was rendered directly into sound. The simulated sound was then compared with known characteristics of the actual sound, and conclusions then drawn about the validity and accuracy of the CFD model.

The potential, however, for the use of sound in almost every aspect of the CFD process, seems considerable, and could be considered at many of the stages:

1. The model *geometry* could be sonified so that glitches and discontinuities in curves and surfaces, which often occur upon transfer from a CAD package to a CFD package, could be quickly identified. In this mode, a smooth geometry would have a pleasing, harmonious sound, and discontinuities could be represented by bursts of noise or discordant pitches.

2. The *grid* could be sonified in an analogous manner, through the use of unpleasant sounds to highlight badly skewed or high aspect ratio cells.

3. Real-time sonification of the *solver* would be analogous to the frequently cited example of the auto mechanic listening to the car engine. CFD solvers work iteratively on non-linear mathematical systems. They frequently diverge or "hunt" without reaching a converged solution. The CFD analyst is required to adjust many solution control parameters, and even choose between alternative solution strategies, in the hope that the solver engine will "behave" and provide a converged solution. Monitoring the solution process via the display alone is monotonous and unproductive; adding sound would allow the analyst to:

   - Listen to the solution in the background while pursuing other tasks.
   - Recognize favorable sound patterns which indicate good choices of solution control parameters and reflect a smoothly running computational "engine."

4. There are many aspects of adding sound to the *post-processing* of CFD data:

   - The recognition of significant patterns in the (sonified) data which are not apparent from visual displays.
   - Increased productivity in examining large amounts of data by concurrent visual and aural displays.
   - The sonic codification of global events in the flow field which are difficult to perceive from the visual diplay of local details [96].
   - The comparison of sound rendered from simulations of the pressure field with recorded sounds from the corresponding experiment [96].

CFD and Music

As a corollary to sonification work, there are various ways to consider CFD as a source of new music. First, CFD provides a numerical mirror into the natural world, which has long been an inspiration to composers. For example, a classic CFD result is to predict von Karman vortex streets [97] in the wake of a circular cylinder in a cross wind. This periodic vortex shedding phenomenon sometimes gives rise to audible frequencies, as when the telephone wires "sing" in a breeze. The scope of CFD also includes the modeling of the compression waves which produce sounds, even in brass or wind instruments. So CFD could be thought of as a tool to enhance the relationship of the composer with sound phenomena in the natural world.

Second, the graphical representations of the results of CFD models are often aesthetically interesting. A common method of visual rendering is to trace the paths taken by fluid particles, (see for example Fig. 20). In flowfields with obstructions and/or complex geometry, especially if natural convection is present, the particles often get caught up in a complex structure of vortices, such as the horseshoe vortex shown in Fig. 20. Under the appropriate conditions, smaller and smaller vortices are spawned from their parents, until a turbulent or chaotic flow structure results. Other graphical renderings include the representation of local fluid velocities by vectors whose size, direction, and color depend on velocity magnitudes and other parameters, (see Fig. 21). Both styles of visual renderings (and others not mentioned here) could be taken as a point of departure for sound exploration.

Thirdly, CFD could be used as a tool for algorithmic composition, following a tradition started by Iannis Xenakis [46] [48] who drew extensively on mathematical formulations from science and engineering. CFD is particularly attractive since the algorithms it uses are iterative by nature, and generate numbers that

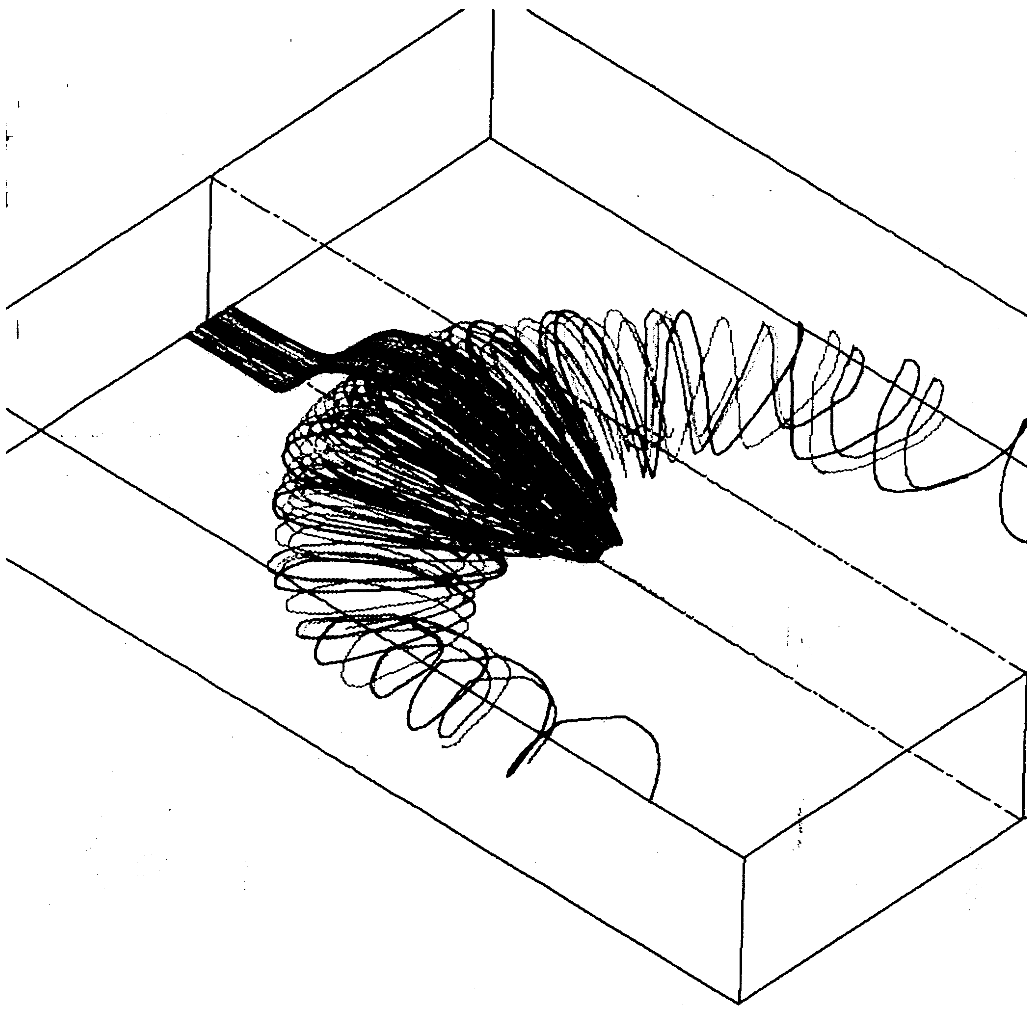

Figure 20: *Particle Paths*

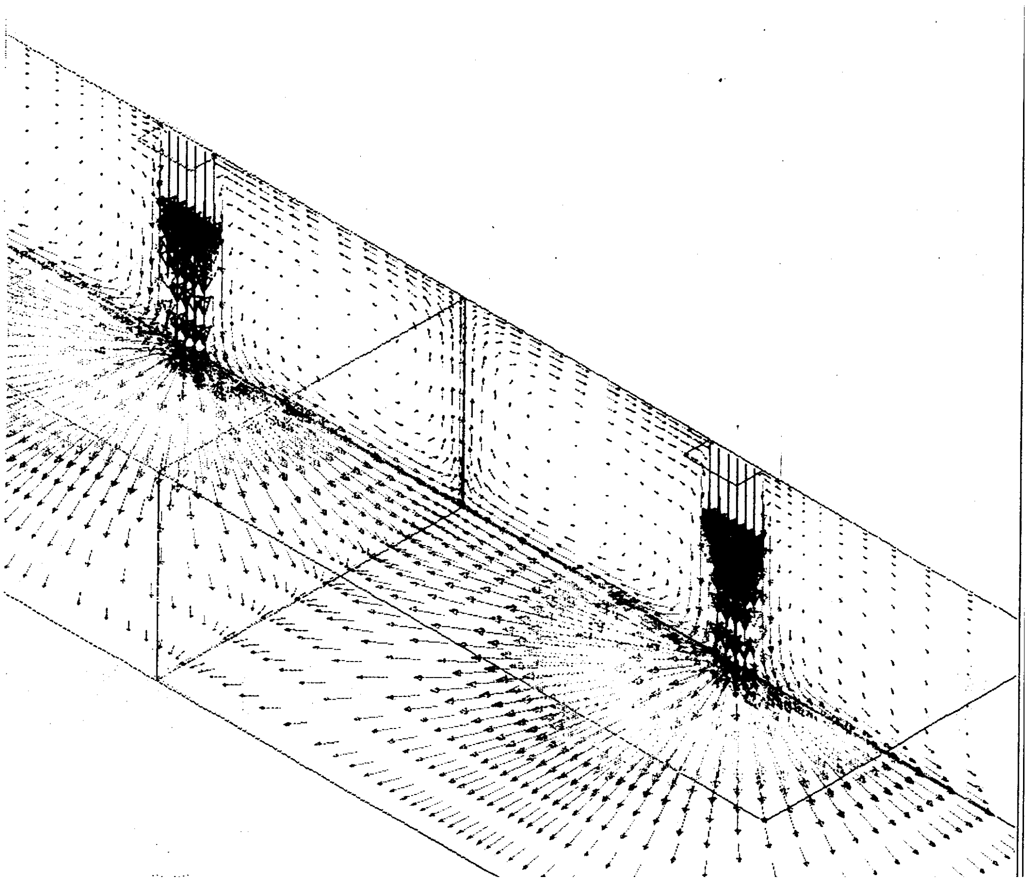

Figure 21: *Vectors*

evolve over time. There is therefore a potential for CFD to generate live musical compositions.

Scope of Present Work

This section will focus on the real-time sonification of Stage 5 (solution) of the CFD process.

An academic CFD research code TEACH-L [98], originally written in FORTRAN, was ported by the author to Java in order to make use of JSyn [90], a digital audio synthesis package. JSyn (**J**ava **Syn**thesis) is a Java API (**A**pplication **P**rogramming **I**nterface), which provides several classes of objects that can create and modify sound. A fast DSP synthesis package written in C lies beneath JSyn's hood. All Java synthesis calls are passed transparently to the C engine.

TEACH-L solves the steady, laminar equations of the conservation of mass, momentum (Navier-Stokes) and energy on a two-dimensional, cartesian grid, using a hybrid differencing scheme and the SIMPLE (**S**emi-**I**mplicit **M**ethod for **P**ressure-**L**inked **E**quations) algorithm [99] to correct the pressure field. The algebraic equations are solved line-by-line (LBL), using the tri-diagonal matrix algorithm (TDMA). The TEACH-L code is extremely compact. There is no user-interface. To set up the geometry, grid, boundary conditions, and physical properties, the user must write her own subroutines using the templates provided. This structure was retained in the Java port, so that TEACH-L consists of a CFD API.

The compactness of TEACH-L, together with the flexibility of the JSyn API, afforded a reasonable implementation on a Macintosh G3 Powerbook, providing a tool for the in-depth exploration of various sonification strategies for a CFD solver.

In the following sections, the physical and mathematical basis of TEACH-L will be presented. One sonification strategy will be explored, followed by results and conclusions.

Physical and Mathematical Basis of the CFD Solver

The Governing Equations In the case of steady, two-dimensional flow, the continuity (conservation of mass) equation is:

$$\frac{\partial}{\partial x}(\rho u) + \frac{\partial}{\partial y}(\rho v) = 0. \tag{4.26}$$

where ρ is the local fluid density (kg/m^3), u and v are the local fluid velocities (m/s) and x and y (m), correspond to the cartesian coordinate system.

For incompressible flow, the momentum equations are for the x direction:

$$\rho u\frac{\partial u}{\partial x} + \rho v\frac{\partial u}{\partial y} = -\frac{\partial p}{\partial x} + \frac{\partial}{\partial x}(2\mu\frac{\partial u}{\partial x}) + \frac{\partial}{\partial y}(\mu[\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x}]), \tag{4.27}$$

and for the y direction:

$$\rho u\frac{\partial v}{\partial x} + \rho v\frac{\partial v}{\partial y} = -\frac{\partial p}{\partial y} - \rho g + \frac{\partial}{\partial x}(\mu[\frac{\partial v}{\partial x} + \frac{\partial u}{\partial y}]) + \frac{\partial}{\partial y}(2\mu\frac{\partial v}{\partial y}), \tag{4.28}$$

where g is the acceleration due to gravity (m/s^2), p is the fluid static pressure (Pa) and μ is the fluid dynamic viscosity (kg/ms).

The energy conservation equation for the fluid, neglecting viscous dissipation and compression heating, is:

$$\rho c_p(u\frac{\partial t}{\partial x} + v\frac{\partial t}{\partial y}) = \frac{\partial}{\partial x}(k\frac{\partial t}{\partial x}) + \frac{\partial}{\partial y}(k\frac{\partial t}{\partial y}), \tag{4.29}$$

where c_p is the fluid specific heat at constant pressure (J/kg K), k is the fluid thermal conductivity (W/m K), and t is the fluid static pressure (K).

Discretization To solve the non-linear partial differential equations from the previous section, it is necessary to impose a grid on the flow domain of interest, (see Fig. 22). In TEACH-L, discrete values of fluid velocities, properties,

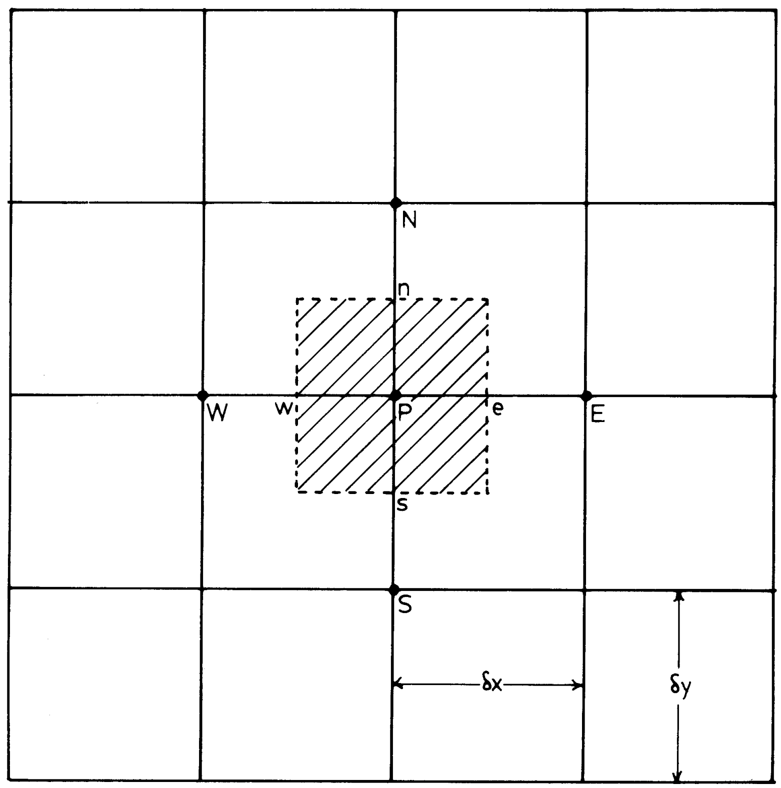

Figure 22: *Control Volume*

pressure and temperature, are stored at each grid point (the intersection of two grid lines). To obtain a matrix of algebraic equations, a control volume is constructed (shaded area in the figure) whose boundaries (shown by dashed lines) lie midway between grid points P and its neighbors N, S, E, W. A complex process of formal integration of the differential equations over the control volume, followed by interpolation schemes to determine flow quantities at the control volume boundaries (n, s, e, w) in Fig. 22, finally yield a set of algebraic equations for each grid point P [100]:

$$(A_P - B)\phi_P - \sum_c A_c \phi_c = C, \tag{4.30}$$

where the subscript c on $\sum$, A and ϕ refers to a summation over neighbor nodes N, S, E and W, ϕ is a general symbol for the quantity being solved for (u, v or t), A_P, etc. are the combined convection-diffusion coefficients (obtained from integration and interpolation), and B and C are, respectively, the implicit and explicit source terms (and generally represent the force(s) which drive the flow, e.g., a pressure difference).

Solution Equation (4.30) must be written at each node where the value of ϕ_P is required; doing so will generate an $N \times N$ matrix of simultaneous equations, where N is the number of nodes in the solution domain. It is usually impractical to invert this matrix directly, so instead a line-by-line (LBL) scheme is used, (see Fig. 23). In the LBL scheme, only one line at a time is solved, quantities on other lines are considered to be "known." Thus, a smaller $n \times n$ tri-diagonal matrix results, where n is the number of nodes in the j direction. The solver starts at the first i-index line in the solution domain, and "sweeps" from left to right.

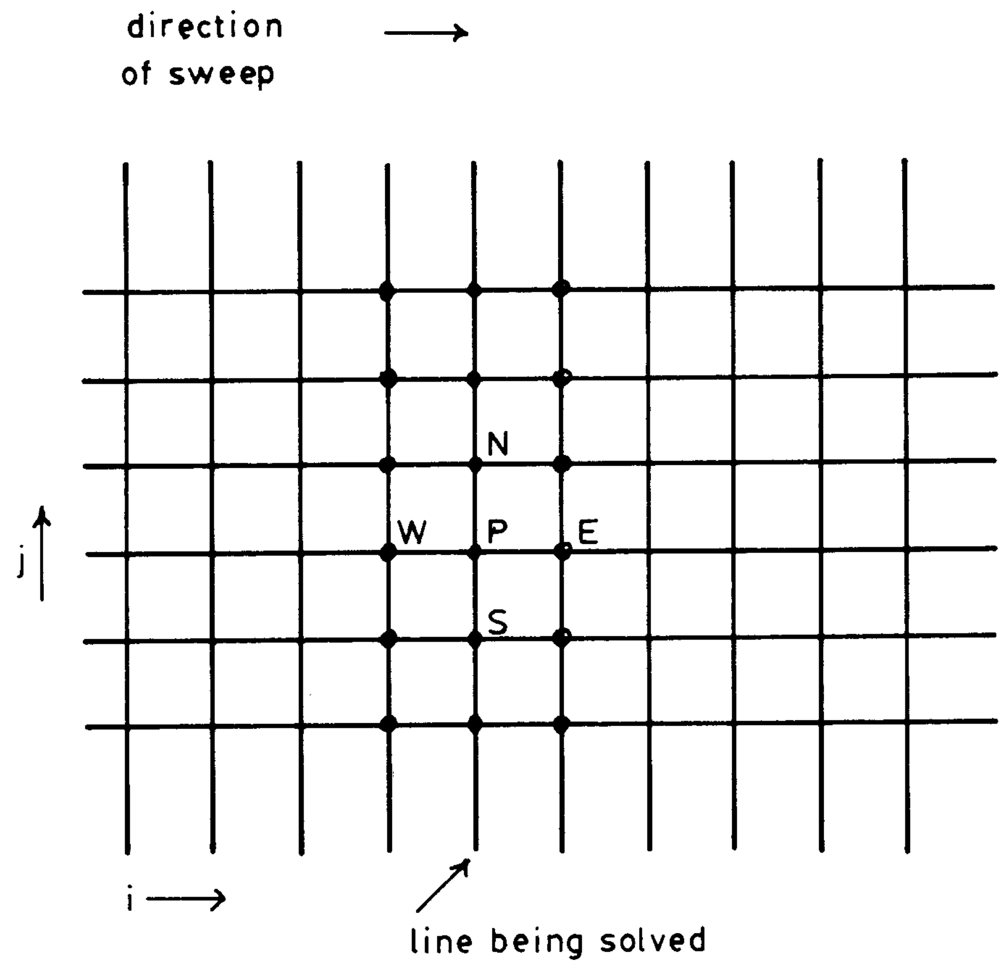

Figure 23: *Line-by-Line Scheme*

Because the partial differential equations are non-linear, the resulting algebraic matrix equations will be also, that is, the convection-diffusion coefficients A_P, etc. will themselves be functions of the ϕ's. An iterative solution is thus required, as follows:

1. The coefficients for Eq.(4.30) in which $\phi = u$ are formed, and the current global "error" or "residual" is calculated. The matrix of coefficients is solved by LBL and new values of $\phi = u$ are obtained.

2. The same is done for $\phi = v$ and $\phi = t$.

3. A "pressure correction" equation, derived from the mass conservation equation is solved in a similar manner, the values of pressure at each note are updated [99].

4. The global errors calculated in each of the above steps are then all compared to a set of target values. If these errors fall below the target, the solution has converged and the calculation stops. Otherwise, the calculation resumes at step 1.

Solution Control Parameters There are two sets of parameters which are set by the user to control the progress of the solution, the underrelaxation factors and the sweep controls. In an iterative, numerical solution, underrelaxation slows the rate of change of a variable from one iteration to the next. Underrelaxation is necessary for numerical stability and to avoid divergence of the solution. Sweep controls for each variable are set to control the number of times, per iteration, the LBL procedure is applied to the coefficient matrix. The greater the number of sweeps, the better the matrix inversion at a particular iteration (but the greater the amount of CPU time required.)

Duct Flow Example

As an example of a solver sonification, the simple problem of steady, laminar, two-dimensional developing flow in a planar duct will be considered (see Fig. 24). In this flow situation, fluid enters at the left with a uniform velocity profile, which develops, as the fluid reaches the end of the duct on the right, into a parabolic profile which is characteristic of "fully-developed" flow, Eq. 4.31:

$$u = \frac{3}{2}u_m(1 - \frac{4y^2}{c^2}), \tag{4.31}$$

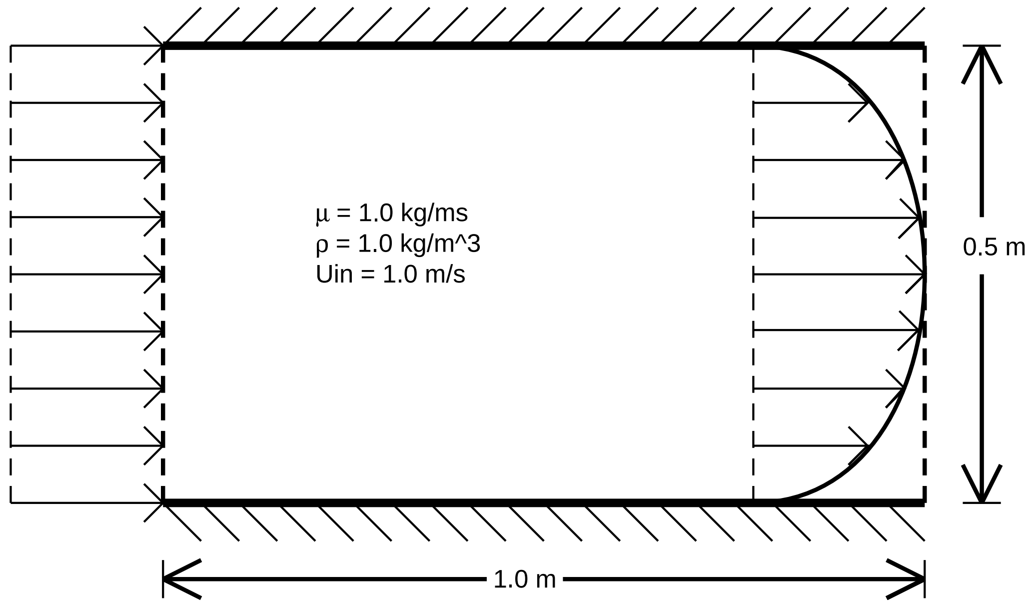

Figure 24: *Developing Flow in a 2D Duct*

where u_m is the average velocity (in this case, 1 m/s), c is the height (in this case 0.5 m) and $y = 0$ at the centerline of the duct. When the flow is fully developed, the transverse velocity component v vanishes, and the pressure p changes only linearly with x:

$$\Delta p = \frac{12 u_m \mu \Delta x}{c^2} \tag{4.32}$$

which gives the pressure drop Δp over some x-direction length Δx, and where the viscosity μ is in this case 1 kg/ms.

A coarse 7×7 evenly spaced staggered grid was used, see Fig. 25 yielding a total of $5 \times 5 = 25$ internal or "live" cells at which the values of u, v and p are updated at each iteration by the solver. With this very simple configuration, the solver converges in about 20 iterations.

Sonification Strategy

The purpose of this sonification was to gain insight into the solver by listening to its progress in real time. To accomplish this, 5 sine oscillators were set

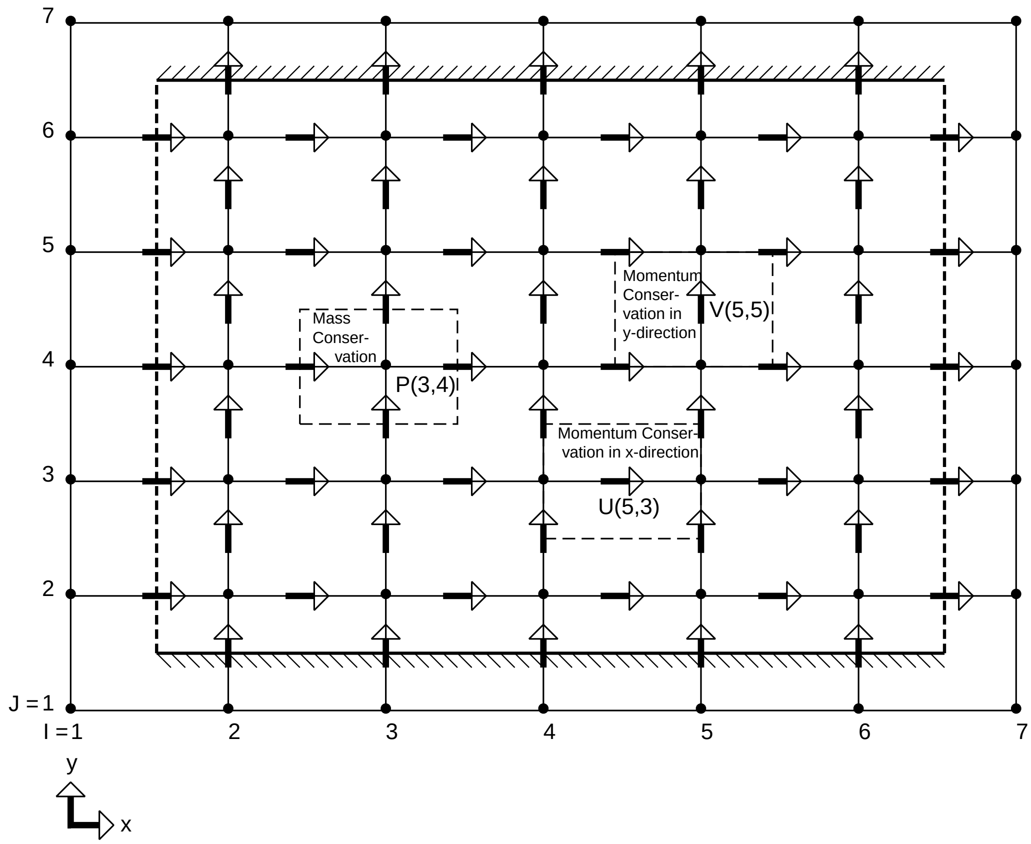

Figure 25: *Storage Locations of Velocities and Pressure*

up to correspond to each column of "live" grid points. As the solver (for u, v then p) sweeps through the domain from left to right, first the column at $i = 2$ sounds, from $j = 2, 6$, with slight arpeggiation, and so on, with a slight pause for each column, through $i = 6$. Thus, each iteration produces 20 notes each for u and v, and 25 notes for p, 65 total. See Fig. 25 for the storage locations of u, v and p. The fact that there are fewer notes for the velocities is by virtue of the staggered grid i.e., the pressure and the two velocity components are not stored at the same locations. This arrangement facilitates the integration of the partial differential equations over the control volumes for the the conservation of mass, momentum in the x-direction and momentum in the y-direction.

In most CFD simulations, general trends and flow behavior are known to the engineer in advance of the calculation. In this sonification, the pitch strategy for each variable was selected so that the anticipated behavior in the flow direction could be "heard":

1. Development of u from a flat to a parabolic profile.

2. Vanishing of v.

3. Linear decrease of p in the flow direction.

Pitch Mapping The values of u range from uniformly 1.0 at the inlet (near $i = 1$) to values in the range $0.0 \leq u \leq 1.5$ at the flow exit (near $i = 5$). The initial guess for u throughout the domain is 0.0. As the solution iterates, values of u well in excess of 1.5 may result. The values of u were mapped to pitch and scaled so that a value of $u = 1$ would yield 440 Hz, with 55 Hz set as the minimum value for values of $u < 0.125$.

The values of v range from approximately 0.1 near the inlet to very small numbers near the exit (the initial guess throughout the flow field, as for u, is $v = 0.0$). The values of v were also mapped to pitch, but scaled up so that a value of $v = 1$ would yield 8400 Hz. All nodes with values of $v <$ about 0.05 were mapped to a major triad with pitches (150, 300, 375, 450) Hz, for nodes $j = 2, \ldots 6$.

The values of p in incompressible fluid flow are generally calculated relative to some numerically convenient reference value, in this case $p_{ref} = 1.0$ Pa at node $(i, j) = (2, 2)$. Values of p are calculated *relative* to $p(2, 2)$, where the significant pressure information is the Δp between nodes which drives the flow. Using this scheme, the values of p range from $-40. \leq p \leq 1.0$. The local value of p was added to a different major triad (100, 200, 300, 400, 500) Hz for $j = 2, \ldots 6$. Because the values of the pressure are fairly uniform at each i

location, one hears, for this scheme, a succession of major triads whose pitch either increases, decreases or stays the same.

Envelope The envelope (attack, sustain, decay) characteristics were derived from the matrix coefficients for each variable at each node:

$$\begin{aligned} t_1 &= A_N, \\ a_1 &= A_N/A_P, \\ t_2 &= A_S, \\ a_2 &= A_S/A_P, \\ t_3 &= A_E, \\ a_3 &= A_E/A_P, \\ t_4 &= A_W, \\ a_4 &= A_W/A_P, \\ t_5 &= 1.0, \\ a_5 &= 0.0. \end{aligned} \tag{4.33}$$

Equations 4.33 represent a 5 frame envelope where a_1 is the amplitude of the first frame, t_1 is its rise time, etc. The values for the fifth frame are set to constant values of 1.0 and 0.0, respectively, to ensure that the oscillator will turn off. The values of t_1 were constrained to be at least 0.05 secs, to avoid clicks due to zero-length frames. Division of all neighbor coefficients A_N, etc. by the center coefficient A_P ensures that $0.0 \leq a_n \leq 1.0$, since the center coefficient is always greater in magnitude than its corresponding neighbors.

Duration In general, as the solver proceeds and makes available the latest values of variable and coefficient at each node, the mapped pitches and envelopes were queued to the oscillators. However, to make the sonic result more intelligible, some delays were added. Firstly, very slight delays were added between nodes in each column, to produce something like an arpeggiated chord. Secondly, a longer delay was added at the conclusion of each column, before proceeding to the next column. Thirdly, at the conclusion of the sonification of a variable at all 25 nodes, a longer delay was added proportional to the global error calculated for that variable. Thus, as the calculation proceeds, this delay decreases as the error is reduced. Finally, at the conclusion of a single iteration, a longer delay was added. It is thus possible for the listener to distinguish, based on the delay between notes, the different stages of the calculation.

Timbre No specific timbral mapping was attempted, since a sine oscillator was used for all notes. However, because each note was given a unique envelope, some striking timbral differences resulted, mainly from different attack times.

Results

The parameter/variable mapping choices described in the previous section had the following effects:

- The u velocity sonification captured the oscillations of the solver well. The solver initial guess of $u = 0$ was followed, at the second iteration, by values considerably overshooting the final result. This behavior sounded like a sequence of low, followed by high frequency chord progressions, finally settling out to a noticeably repeating sequence during the final iteration. Owing to large values of the neighbor coefficients A_N, etc. and the large values of center coefficient A_P at most nodes relative to the neighbors, the attack times were long, and the amplitudes low. Thus the u velocity sonification sounded ethereal and slowly evolving and was easily distinguished from the v velocity and p data.

- From the v velocity sonification, it was easy to hear an initial guess of $v = 0$, via the major chord, followed, after oscillations, by non-zero (higher frequency) sounds at the inlet, and vanishing values near the flow exit. The attack times for this variable were shorter than those for the u velocity, and the amplitudes higher. It was thus easy to hear the transition from the u to v sonification.

- The p sonification was very different in timbre to those for u and v, owing to much shorter attack times and higher amplitudes. It was thus easy to perceive the onset of p data, and to hear, at each iteration, whether the pressure was increasing, staying the same, or finally, as in the converged solution, decreasing in the direction of flow.

Conclusions

The sonification of this simple duct example afforded enhanced interaction with the data in two important ways:

1. The ability to monitor the progress of convergence of the data throughout the flow field. In most CFD packages, visual monitoring of solution progress is only practical at one or two locations. The addition of sound allows the engineer to hear that, on a global basis, the solution is or isn't evolving as expected, and roughly where in the domain a problem might exist if there is one.

2. The ability to notice differences in data (the neighbor and center coefficients), via the envelope characteristics. The different timbres in the three variables aroused curiosity, and triggered further investigation of these coefficients, in order to determine if their values were correct, and to question why they were different for each variable.

In general, it is clear that the addition of sound has enormous potential for the critical examination of CFD simulations and merits further investigation.

Further Work

The current sonification was extended to add spatialization of the sound, so as to hear a sweeping progression from left to right, corresponding to the entrance and exit of the duct, respectively. More complex "instruments" were constructed (see Section 4.1.3), to enhance the effect of local coefficients, and to add the effect of additional coefficients not currently mapped, such as source terms, and local conservation errors. Using the FM formant instrument (see Fig. 19), the u velocity, v velocity and pressure were sonified as one *klang*, instead of being sonified in succession, as with the sine oscillator instrument. The index of modulation was mapped to conservation imbalance. The higher the degree of imbalance, the higher the modulation index. The result was that the cells in which a significant imbalance existed had an "edgy" or "scratchy" sound, when sonified. Cells which were well balanced had a "smooth" sound, when sonified. This sonification drew attention to the fact that significant mass imbalance occurs predominantly in the cells just upstream of the duct exit.

Thus, the sonification scheme resulted in the exposure of a significant area of the solution domain where imbalances were higher than in other regions. In a practical situation, the numerical analyst might not otherwise have been aware of this behavior, and would probably investigate the cause. Such numerical anomalies are typical in large, computer-intensive simulations, but are sometimes difficult to locate using currently available visual and analytical techniques.

The ultimate goal for the sonification of complex numerical processes should be to map every available parameter in the numerical world to some recognizable characteristic in the sound domain, in such a way that it can be distinguished and singled out for inspection and further investigation if warranted. It is also conceivable that such sonifications could be perceived in a global or coherent sense, as distinguished from perception of the component parts.

Musical Compositions

The duct flow example using the sine oscillators was used as an algorithmic composition tool by recording several solver runs with different parameter settings so as to change the tempi or pitch centers. Various soundfiles were created, processed, and mixed using ProTools software. The result is a piece entitled *CFD Sound V*.

The second sonification, using the FM Formant instrument, was organized as a live piece entitled *Duct Tape*, and was performed during the summer of 2002. The flow simulation was run in real time on a Macintosh G3 Powerbook computer; however, the speed (pace) of the solution was manipulated by the author.

4.2.2 Meteorological Data

As a case study, a sonification of historical meteorological data sets was performed, in which the occurrence and location of all hail storms in the continental United States from 1955 to 1995 was provided [101].

The sonification of historical data at a "playback" speed faster than the occurrence of the original data has been used to discover larger scale trends or patterns in financial [102], geological [14], [15] meteorological [103], and other data, which were not apparent from (sometimes more tedious) visual or analytical methods. Rapid scanning of large multi-variable data sets has also been used to "search," e.g., oil well logs [104], for anomalies or significance – another fruitful area in which to use sound spatialization techniques.

The meteorological data files chosen for sonification provide (chronologically, one line for each event) the date, time (to the nearest minute) and location of every hail storm in the continental United States recorded at a National Weather Station (NWS), together with the diameter of the hail, and the number of injuries. The location information is by State, County, NWS office (a three letter code), latitude and longitude. A few lines from a file of 1995 data is shown as an example, only the columns used in the sonification are shown. The format is as originally used in a Data General mainframe computer database. The columns contain the following information:

950150101272330	6	342808752	HSV	0	00059	0	75C
950302801280029	6	320208853	MEI	0	00023	0	75C
950312801280030	6	322408839	MEI	0	00075	0	75C
950322801280100	6	314208906	MEI	0	00067	0	75C
950160101280715	6	312208551	MGM	0	00031	0	75C

1. The year, sequence number, state, date and time.

2. A numerical code (6 or 7) to indicate the occurrence of either hail larger than 0.75 inch in diameter (6) or thunderstorms winds in excess 50 knots. *This number was always 6, indicating that the files actually contained only hail storm information.*

3. The latitude and longitude in degrees and minutes.

4. A three letter code for the name of the NWS at which the data was recorded.

5. The number of fatalities. *This number was always 0, indicating either a failure to record deaths, or that nobody died.*

6. The number of injuries, and the FIPS (Federal Information Processing Standards) code for the county.

7. An economic damage code ranging from 0 (less than 50 US dollars) to 9 (between 500 million and 5 billion USD). *This number was always 0 and not used in the sonification.*

8. The size of the hail in hundredths of an inch ($75 = 0.75$ inches).

As an example, the first entry depicts a storm (hail, diameter 0.75 inches) which took place on 27 January, 1995 at 2330, in the state of Alabama (FIPS code 1), at a latitude of 34 degrees 28 minutes, and a longitude of 87 degrees 52 minutes, recorded at the Huntsville (HSV) weather station, in Franklin county (FIPS code 59), in which no injuries occurred. The next 3 entries track, presumably, the same storm, occurring just past midnight the following day in Mississippi (FIPS code 28) in three counties Clarke (23), Lauderdale (75) and Jones (67), but all reported at the Key Field Airport in Meridian (MEI) (see Fig. 26).

In many cases, the same storm system may be reported at different locations at successive times as it takes its course. Applying a spatial sonification which displays each report at a different location in auditory space provides the listener with a sense of a storm's movement. Furthermore, over the course of the year there is seasonal variation in the location and frequency of hail storms.

Musical Compositions

Two musical compositions were prepared from the hail storm data. The first *Digital Swamp* was prepared as a collaborative piece with Stefan Tomic, a colleague in the graduate program. In *Digital Swamp*, sonifications of the same data by Childs and Tomic were mixed together in a collaborative session, using ProTools. While the two sonification schemes were different and formed

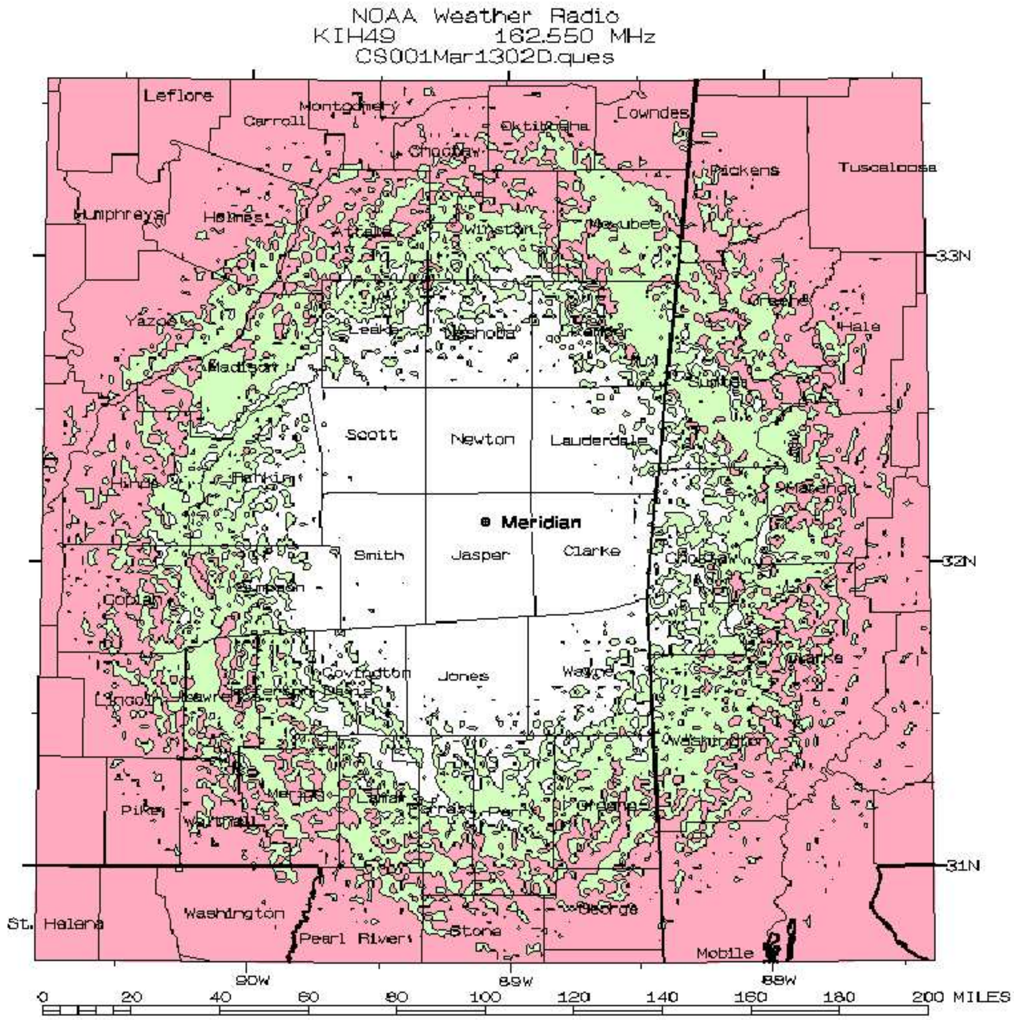

Figure 26: *Meridian, Mississippi Weather Station Area*

contrasting sections, the fact that the same data set was used lended coherence to the piece.

The second composition "Meridian" was prepared as a four track mixdown of an eight track spatialized sonification of the 1995 hail storm data. The piece was played in a Winter 2003 electro-acoustic concert at the Top of the Hopkins Center at Dartmouth College.

4.2.3 Financial Data

The trading of financial securities is influenced by cycles of rising and falling prices which operate over different time scales. Investment strategies vary, according to the particular cycle the financial trader wishes to exploit. Proprietary traders seek to exploit second by second fluctuations and typically execute hundreds of trades in one day. Hedgefund managers seek to capitalize on both upward and downward trends by taking "short" and "long" positions, perhaps altering their positions once or twice in the course of a day. Portfolio managers might evaluate stock or bond performance over time scales of a month or longer. The general public typically invests in IRAs or 401(k) plans in the hope that, eventually, the value of their investments will increase. It is sometimes argued that if the stock market prices are compared over any 11 year period, the trend is always favorable.

Fascination with market cycles, and the desire to understand them for profit, has given rise to a field sometimes referred to as financial technology, or financial engineering. Sophisticated mathematical models are proposed, and then tested on historical financial data. Traders and analysts also seek to predict psychological market behavior, such as panic or optimism. Unsubstantiated rumors often result in substantial and apparently irrational market fluctuations. In the prelude to the Iraq war of 2003, on the day in which the weapons inspectors discovered empty canisters which they presumed had contained chemical agents, bond prices were falling earlier in the day on favorable economic news. Stocks were rising. The news about the canisters was announced in mid-afternoon causing a sudden reversal in these trends. Further news later in the day which mitigated the crisis resulted in a re-establishment of the earlier trends.

Several researchers have experimented with sonifications of historical financial data. Richard Voss, who proposed that the audio power of many musical selections and English speech varies as approximately $1/f$ where f is the frequency [105], found that this correlation also appeared to apply to a melody line generated from IBM stock prices (`http://www.aps.org/apsnews/1299/129905.html`). Over the course of his career at the IBM Thomas J. Watson Research Center, Voss, who specializes in scientific computer graphics, has studied the application of fractals to art, music and the stock market. In 1999, he presented the IBM melody as an "aural representation of Brownian motion."

Kramer [74] performed a series of sonifications of historical Dow Jones, bonds, the US dollar, interest rates and commodities data recorded over a period of 4 1/2 years. He mapped stocks to pitch, bonds to pulsing speed, the dollar to brightness, interest rates to detuning and commodities to attack time. His mapping scheme is interesting in that all 5 variables are presented simultaneously. However some confusion results, at least to my ear, because the pitch seems to be the most predominant variable. It is also difficult (in my opinion) to relate an increase or decrease in a price to variables other than pitch. It would possibly make more sense to use pitch (for the price) of all 5 variables, using timbral and/or spatial variation to distinguish between Dow Jones, bonds, etc.

Jonathan Berger of CCRMA, Stanford, published several sonifications of historical financial data from the year 2000 on the web, in which he used pulses of filtered noise with vocal formants (`http://www-ccrma.stanford.edu/groups/soni/index.html`). The pitch content of the filtered sound tracks prices, and the filter bandwidth tracks trading volume.

Nesbitt and Barrass [106] sonified "bid" and "ask" data, together with other market variables (depth of market), in a real-time scenario, in conjunction with a virtual reality display requiring 3D goggles. The system was designed and tested by people who had no stock market expertise.

The financial data project was to sonify of one day of Nextel stock trading data (Feb. 5, 2002). Nextel is a telecommunications company which trades on the Nasdaq Electronic Exchange. On this particular day, the stock started at \$6.89 a share in the early morning (8 AM), fell to below \$4.38 a share on the expectation of an unfavorable earnings report, but then recovered and finally closed above \$5.00 when an unexpectedly good report was issued prematurely by the company. The data was read from a file, which contained the time of the trade (to the nearest 0.01 sec), the trade price, and the volume of the trade (number of shares). A `TriangleOscillator` was used from the JSyn API. Every note represents one trade or "print." Price was mapped to pitch, trade volume to loudness. The timing was calculated directly from data file, with an acceleration factor added to either speed up or slow down the data. At the height of the activity, when trades occurred more frequently than 0.01 sec, the spacing between notes was set to be approximately 0.1 sec. Three longer `SineOscillator` tones were added to provide a reference, or sonic grid, to the sonification. The first and highest tone has the same pitch as the currently highest price which in this case was the opening price. The second, lower tone has the same pitch as the currently lowest price (which starts close to the opening pitch but gradually drifts lower). The third, even lower tone is the same pitch as the target price (the price someone would intend to sell at). In this case, the target price was set to \$5.00 a share. This pitch remains the same throughout the sonification, and thus provides a reference tone. The register of the mapping is fairly high, ranging from 100 Hz to 1000 Hz, and the sound of the early market, with intermittent trades gradually increasing in frequency, is reminiscent of the sound of a bird sanctuary waking up. A composition was prepared from this sonification, entitled *Market Study*.

5 Current and Future Work

5.1 ICAD 2003 Project

The hail storm sonification described in Section 4.2.2 is currently being developed as a historical data exploration tool. Faculty members from the Meteorology Department at Lyndon State College who specialize in severe weather have agreed to provide feedback on the current system.

The basic concept of the current design is as a tool for exploring temporal and spatial patterns. The data set is selected, and navigated, via the Data Exploration Tool, (see Fig. 27). The data may be "listened to" from any desired position (which is always mapped, spatially, to a point directly overhead). For example, a user could elect to "listen" from some point in New York State, only to storms within a certain radius. Approximately 40 years of hail storm data is available. These may be selected from a drop-down list. The transport controls enable the user to Start, Stop, Pause, Step and Replay the immediately previous sound. These controls have been added, both for the purpose of allowing the listener to investigate an unusual sound, and for the sonification designer to fine-tune the mapping.

Currently, the spatial mapping for the geographical location of the storm is constrained to two angles (the azimuth and the elevation). That is, a particular geographic location is mapped to a location on the (approximately) hemispherical surface outlined by the loudspeaker array. Future developments of this spatialization might include a truly 3D virtual source placement, using perhaps delay, or loudness to move the sound relative to the listener within the speaker array.

The advantage of using loudspeaker arrays for sonification is that a group may collectively listen and compare notes. However, there are also many applications in which headphones are the preferred listening mode. A further project

Figure 27: *Hail Storm Data Exploration Tool*

would involve the incorporation of HRTF algorithms into JSyn.

5.2 Financial Data

The real test of sonification is in practical situations. Thanks to the resources of the Dartmouth Entrepreneurial Network, the author was able to have in depth discussions with four financial traders who trade government bonds, municipal bonds, stocks (for mutual funds) and stocks (for a hedge fund).

The data requirements and working environment of the four individuals were substantially different.

The municipal bond trader monitored major indices and US Treasury instruments on a Bloomberg terminal using the "BTMM" screen. He used other screens to research municipal bond offerings, work with spreadsheets, and answer emails. He provided service to portfolio managers who wished to know, at any given moment, the best deal on fixed income securities. He was not able to monitor the "BTMM" screen at all times so missed significant movements. He required a broad indication of market movements that were significant.

The government bond trader worked on a treasuries desk. He was a proprietary trader, making hundreds of trades in a single day. He required frequent updates of minute movements in two major indices, together with proprietary analytics, including actual trading activity on the electronic exchanges.

The stock trader worked at a large investment bank for mutual fund managers. She worked with large number of securities, choosing to pay particular attention to a smaller subset on a given week or day. She was particularly interested when any of the selected securities deviated from the opening price by some specified percentage.

The stock trader for hedge funds was interested in the movement of a large group of stocks relative to their 30, 50 or 200 day moving averages. Typically, movement of a single stock with respect to its moving averages is viewed graphically on a Williams Chart, (see Fig. 28). In the figure, the movement of a business services stock (jagged line) ChoicePoint, Inc., with ticker symbol CPS, relative to its 200 day (lighter smooth curve) and 50 day (darker smooth curve) moving averages. Significant events often occur when the moving averages cross. In this case, on Feb. 27, 2003, the 50 day average descended and crossed the 200 day average midway through the trading day. This was followed by a precipitous drop in the price of the actual stock.

Since it is not generally practical to observe moving average activity for more than a few stocks, even on several screens, the hedge fund trader required a sonification of the approach to the average, or several (say 10 or 20) stocks, in such a way that the loudness or prominence of the sonification would increase only when the stock was "close" to the average.

Research in these working environments is on-going and will possibly be

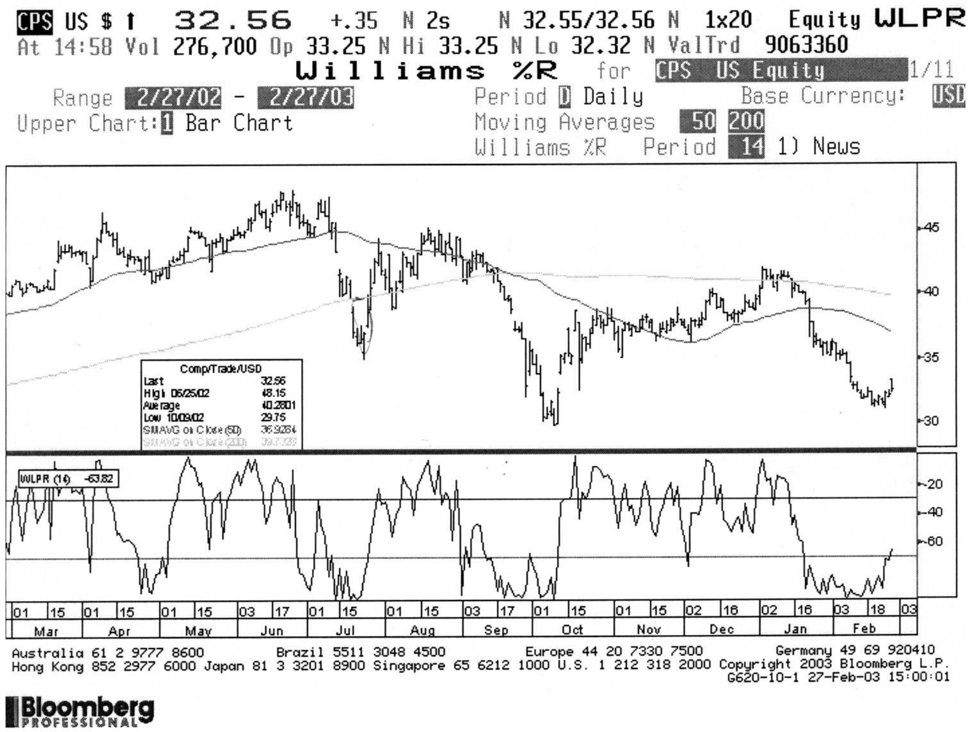

Figure 28: *Williams Chart*

presented at ICAD 2004.

5.3 Future Directions for Sonification

For sonification research to flourish, several challenges lie ahead.

1. Sonification must be recognized by academic and research funding organizations as a legitimate, unique, interdisciplinary field which spans the arts and the sciences. Currently, sonification "lives" within established computer science, acoustics, psychology or computer music departments. Perhaps the "Center of Excellence" model could be used in the establishment of sonification research units. The report for the National Science Foundation written by Kramer, et al., should help in this regard [6].

2. Workers in the disciplines of scientific investigation must learn to respect the process of composers and creative artists, and vice versa, in order for

sonification to thrive. Such mutual respect may require a "paradigm shift" in the current procedures for conference preparation and peer review in both electro-acoustic music and scientific research communities.

3. Development of sonification research environments, such as those which have been suggested by Kaltenbrunner [67] and Ben-Tal [68]. Such environments would facilitate more rapid development of sonification protoypes much in the same way that off-the-shelf GUI development kits do so today.

4. Practical applications for sonification must be found in realistic working environments. Thus far, significant applications exist primarily in medical operating room situations, or in the military.

The author foresees many far reaching practical applications of sonification to meet the needs of knowledge workers who process information rather than material goods for a living, particularly in the fields of financial services, network monitoring, homeland security, manufacturing and healthcare. These workers suffer increasingly from visual data overload and the stress of visual monitoring, which requires them to stare at visual displays.

Beyond the practical, it is hoped that sonification will open an auditory window to patterns not now observed, and give a voice to music previously unheard.

References

[1] E. Childs, "Achorripsis: A sonification of probability distributions," in Nakatsu and Kawahara [107].

[2] E. Childs, "The sonification of numerical fluid flow simulations," in Hiipakka *et al.* [108].

[3] E. Childs, "Using multi-channel spatialization in sonification: A case study with meteorological data," in *Proceedings of the 9th International Conference on Auditory Display* (B. Shinn-Cunningham, ed.), (Boston, MA), 2003.

[4] W. A. Neilson, ed., *Webster's New International Dictionary*. Springfield, MA, USA: G. & C. Merriam, Co, unabridged second edition ed., 1947.

[5] C. Scaletti, "Sound synthesis algorithms for auditory data representations," in Kramer [8], pp. 223–251.

[6] G. Kramer and et al., "Sonification report: Status of the field and research agenda." `http://www.icad.org/websiteV2.0/References/nsf.html`, 1997.

[7] G. Mayer-Kress, R. Bargar, and I. Choi, "Musical structures in data from chaotic attractors," in Kramer [8], pp. 341–368.

[8] G. Kramer, ed., *International Conference on Auditory Display*, no. XVIII in Studies in the Sciences of Complexity, (Jacob Way, Reading, MA 01867), Santa Fe Institute, Addison-Wesley Publishing Company, 1994.

[9] L. Polansky, "Cocks crow, dogs bark: New compositional intentions," *Leonardo Music Journal*, vol. 7, pp. 61–71, 1997.

[10] T. Johnson, *Self-Similar Melodies.* 75 rue de la Roquette, 75011 Paris: Editions 75, 1996.

[11] W. W. Gaver, "Using and creating auditory icons," in Kramer [8], pp. 417–446.

[12] S. Ronkainen, "Earcons in motion – defining language for an intelligent mobile device," in Hiipakka *et al.* [108], pp. 126–131.

[13] C. Hayward, "Sound synthesis algorithms for auditory data representations," in Kramer [8], pp. 369–404.

[14] F. Dombois, "Using audification in planetary seismology," in Hiipakka *et al.* [108].

[15] F. Dombois, "Auditory seismology: On free oscillations, focal mechanisms, explosions and synthetic seismograms," in Nakatsu and Kawahara [107].

[16] K. McCabe and A. Rangwalla, "Auditory display of computational fluid dynamics data," in Kramer [8], pp. 327–340.

[17] E. R. Tufte, *The Visual Display of Quantitative Information.* Box 430, Cheshire, Connecticut 06410: Graphics Press, 1983.

[18] C. D. Wickens, S. E. Gordon, and Y. Liu, *An Introduction to Human Factors Engineering.* New York: Longman, 1998.

[19] C. Scaletti, "Using sound to extract meaning from complex data," in *Extracting Meaning from Complex Data: Processing, Display, Interaction II*, vol. 1459, International Society of Optical Engineering, 1991.

[20] W. T. Fitch and G. Kramer, "Sonifying the body electric: Superiority of an auditory over a visual display in a complex, multivariate system," in Kramer [8], pp. 307–326.

[21] Unknown, 1999. Unpublished lecture at Conductor's Guild Conference.

[22] F. Chopin, "'Etude n. 12 op.25," in *12 studies Op. 25 for piano* (A. Cortot, ed.), pp. 84–91, Salabert Editions, 1914.

[23] C. Scaletti, 2003. Personal communication.

[24] T. L. Read, 1999. Personal communication.

[25] P. Garland, ed., *SOUNDINGS 13: The Music of James Tenney*. Santa Fe: Sounding Press, 1984.

[26] C. Penrose, "Mk ultra," in *American Jingo*, `www.detritus.net/illegalart`: illegal art, 2000.

[27] J. J. Bharucha, B. Tillman, and P. Janata, "Cultural adaptation of the brain to music and speech: an fmri study of listening to indian and western music, hindi and english." Department of Psychological and Brain Sciences, Dartmouth College.

[28] G. Schellenberg, "Name that tune: Identifying popular recordings from brief excerpts," *Psychonomic Bulleting & Review*, 1999.

[29] P. Janata, "Neurophysiological mechanisms underlying auditory image formation in music," in *Elements of Musical Imagery* (R. I. God/oy and H. J/orgensen, eds.), pp. 27–42, Lisse: Swets & Zeitlinger, 2001.

[30] B. Tillman, J. J. Bharucha, and E. Bigand, "Implicit learning of tonality: A self-organizing approach," *Psychological Review*, vol. 107, no. 4, pp. 885–913, 2000.

[31] P. Cariani, "Temporal codes, timing nets and music perception." For Journal of New Music Perception, vol. 30, issue 2 (2001), 2001.

[32] A. D. Patel, E. Gibson, J. Ratner, M. Besson, and P. J. Holcomb, "Processing syntactic relations in language and music: An event-related potential study," *Journal of Cognitive Neuroscience*, vol. 10, no. 6, pp. 717–733, 1998.

[33] F. Faïta and M. Besson, "An event-related potential (erp) study of musical expectancy: Comparison of musicians with nonmusicians," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 21, pp. 1278–1296, 1995.

[34] R. Näätänen, *Attention and Brain Function*. Hillsdale: Lawrence Erlbaum Associates, 1992.

[35] K. Gardner, *Sounding the Inner Landscape*. Rockport, Massachusetts: Element, 1990.

[36] J.-E. Berendt, *The Third Ear*. New York: Henry Holt and Company, 1985.

[37] K. Gann, "Custer and sitting bull." Unpublished Score, 1998.

[38] S. M. Williams, "Perceptual principles in sound grouping," in Kramer [8], pp. 95–126.

[39] A. S. Bregman, *Auditory Scene Analysis*. Cambridge, Massachusetts: The MIT Press, 1994.

[40] D. B. Willingham, *Cognitions: The Thinking Animal*. Upper Saddle River, New Jersey 07458: Prentice Hall, 2001.

[41] P. Janata, B. Tillmann, and J. J. Bharucha, "Listening to polyphonic music recruits domain-general attention and working memory circuits," *Cognitive, Affective and & Behavioral Neuroscience*, vol. 2, no. 2, pp. 121–140, 2002.

[42] S. Adler, *The Study of Orchestration*. New York and London: W. W. Norton and Company, second ed., 1989.

[43] L. Polansky, "Live interactive computer music in hmsl, 1984-1992," *Computer Music Journal*, vol. 18, no. 2, pp. 59 – 77, 1994.

[44] J. Tzelgov, R. Srebro, A. Henik, and A. Kushelevsky, "Radiation detection by ear and by eye," *Human Factors*, vol. 29, no. 1, 1987.

[45] T. M. Madhyastha and D. A. Reed, "A framework for sonification design," in Kramer [8], pp. 267–289.

[46] I. Xenakis, *Formalized Music: Thought and Mathematics in Music*. Hillsdale, NY: Pendragon Press, 1992.

[47] L. Arsenault, "Iannis xenakis's *achorripsis*: The matrix game," *Computer Music Journal*, vol. 26, pp. 58–72, Spring 2002.

[48] I. Xenakis and t. S. Kanach, *Arts/Sciences: Alloys, The Thesis Defense of Iannis Xenakis*. Hillsdale, NY: Pendragon Press, 1985.

[49] J. Harley, "Iannis xenakis bibliography and discography." LEONARDO on-line `http://mitpress2.mit.edu/e-journals/Leonardo/isast/spec.projects/Xenaki%sbib.html`.

[50] "Ccmix - center for composition iannis xenakis." `http://www.ccmix.com`. formerly Les Ateliers UPIC.

[51] B. A. Varga, *Conversations with Iannis.* 3 Queen Square London WC1N 3AU: Faber & Faber, 1996.

[52] I. Xenakis, "The crisis of serial music," *Gravesaner Blätter*, no. 1, 1955.

[53] P. Schaeffer, "Music and computers," in *Music and Technology*, (Paris), UNESCO, La Revue Musical, June 1970. Stockholm Meeting June 8-12, 1970.

[54] J. Appleton, "Machine songs iii: Music in the service of science – science in the service of music," *Computer Music Journal*, vol. 16, Fall 1992.

[55] "Iannis xenakis (1922-2001)." `http://www.bbc.co.uk/music/profiles/print/xenakis.shtml`.

[56] J. Harley, "Formal analysis of the music of iannis xenakis by means of *sonic events*: Recent orchestral works," in *Présences de Iannis Xenakis* (M. Solomos, ed.), 16 place del la Fontaine aux Lions, 75019 Paris, France: Centre de documentation de las musique contemporaine, 2001. `http://www.cdmc.asso.fr`.

[57] "The international conference on auditory display." `http://www.icad.org`.

[58] N. Matossian, *Xenakis.* London and White Plains, New York: Kahn & Averill and Pro/Am Music Resources, Inc, 1991.

[59] R. Squibbs, *An Analytical Approach to the Music of Iannis Xenakis: Studies of Recent Works.* PhD thesis, Yale University, 300 North Zeeb Road, Ann Arbor, MI 48103, November 1996. UMI Microform 9714309.

[60] W. Mendenhall and R. J. Beaver, *Introduction to Probability and Statistics.* Belmont, California: Duxbury Press, 1994.

[61] `http://eamusic.dartmouth.edu/~ed/research.html`.

[62] C. Dodge, *Earth's Magnetic Field.* 15 Columbus Circle, New York, NY 10023: Nonesuch Records, 1970.

[63] C. Dodge, "Earth's magnetic field," in *Columbia Princeton Electronic Music Center, 1961-1973*, New World Records, 1998.

[64] J. Bartels, *Collection of geomagnetic planetary indices K_p and derived daily indices A_p and C_p, for the years 1932 to 1961.* Amsterdam: North-Holland Pub. Co., 1962.

[65] C. Dodge and T. A. Jerse, *Computer Music.* New York: Schirmer Books, 1997.

[66] J. V. Pennington, "The red and white cows," *American Mathematical Monthly*, vol. 2, 1957.

[67] M. Kaltenbrunner, "Y-windows: Proposal for a standard aui environment," in Nakatsu and Kawahara [107].

[68] O. Ben-Tal, J. Berger, B. Cook, M. Daniels, G. Scavone, and P. Cook, "Sonart: The sonification application research toolbox," in Nakatsu and Kawahara [107].

[69] P. Cook and G. Scavone, "The synthesis tool kit," in *Proceedings of the 1999 International Computer Music Conference*, (Beijing, China), pp. 164–166, 1999.

[70] A. Joseph and S. K. Lodha, “Musart: Musical audio transfer function real-time toolkit,” in Nakatsu and Kawahara [107].

[71] C. M. Wilson and S. K. Lodha, “Listen: A data sonification toolkit,” in *Proceedings of the 1996 International Conference on Auditory Display*, (Palo Alto, CA), pp. 1119–1134, 1996.

[72] S. Barrass, “Personify: a toolkit for perceptually meaningful sonification,” in *Proceedings of the Australian Computer Music Conference ACMA ’95*, (Melbourne, Australia), 1995.

[73] R. Boulanger, ed., *The Csound Book*. Cambridge, MA: The MIT Press, 2000.

[74] G. Kramer, “Some organizing principles for representing data with sound,” in *Auditory Display* [8].

[75] E. Neuwirth, “R sings – or using r to sonify data,” in *DSC 2001 Proceedings of the 2nd International Workshop on Distributed Statistical Computing*, (Vienna, Austria), 2001.

[76] B. Eckel, *Thinking in Java*. Upper Saddle River, New Jersey 07458: Prentice Hall, 2000.

[77] J. Pampin, “Ats: A system for sound analysis, transformation and synthesis based on a sinusoidal plus critical band noise model and psychoacoustics.” PAPER/PANEL 1, National Conference of the Society for Electro-Acoustic Music in the United States, 2003.

[78] R. H. Gilkey and T. R. Anderson, eds., *Binaural and Spatial Hearing in Real and Virtual Environments*. Mahwah, NJ: Lawrence Erlbaum Assoc., 1997.

[79] J. Blauert, *Spatial Hearing*. Cambridge, MA: The MIT Press, revised ed., 1997.

[80] H. Møller, M. F. Sørensen, D. Hammershøi, and C. B. Jensen, "Head-related transfer functions of human subjects," *Journal of the Audio Engineering Society*, vol. 43, May 1995.

[81] D. R. Begault, *3-D Sound for Virtual Reality and Multimedia*. Cambridge, MA: AP Professional, 1994.

[82] J. Chowning, "The simulation of moving sound images," *Journal of the Audio Engineering Society*, vol. 19, no. 1, 1971.

[83] D. G. Malham and A. Myatt, "3-d sound spatialization using ambisonic techniques," *Computer Music Journal*, vol. 19, no. 4, 1995.

[84] G. Monro, "In-phase corrections for ambisonics," in *Proceedings of the International Computer Music Conference*, (Berlin, Germany), 2000.

[85] V. Pulkki, "Microphone techniques and directional quality of sound reproduction," in *112th Convention of the Audio Engineering Society*, no. 5500, (Munich, Germany), 2002.

[86] V. Pulkki, *Spatial sound generation and perception by amplitude panning techniques*. PhD thesis, Helsinki University of Technology, 2001.

[87] V. Pulkki, "Virtual source positioning using vector base amplitude panning," *Journal of the Audio Engineering Society*, vol. 45, June 1997.

[88] J. C. Bennett, K. Barker, and F. O. Edeko, "A new approach to the assessment of stereophonic sound system performance," *Journal of the Audio Engineering Society*, vol. 33, May 1985.

[89] A. J. Berkhout, D. de Vries, and P. Vogel, "Acoustic control by wave field synthesis," *Journal of the Acoustical Society of America*, vol. 93, May 1993.

[90] "Jsyn audio synthesis." `http://www.softsynth.com`.

[91] V. Pulkki, "Creating auditory displays to multiple loudspeakers using vbap: A case study with diva project," in *International Conference on Auditory Display*, (Glasgow, Scotland), 1998.

[92] J. M. Chowning, "The synthesis of complex audio spectra by means of frequency modulation," *Journal of the Audio Engineering Society*, vol. 21, no. 7, 1973.

[93] C. L. M. H. Navier, "N/a," *Memoires Acad. R. Sci.*, vol. 6, 1830.

[94] G. G. Stokes, "N/a," *Trans. Camb. Phil. Soc.*, vol. 8, 1845.

[95] J. E. Fromm and F. H. Harlow, "Numerical solutions of the problem of vortex street development," *Physics of Fluids*, vol. 6, 1963.

[96] K. McCabe and A. Rangwalla, "Auditory display of computational fluid dynamics data," in Kramer [8].

[97] H. Schlichting, *Boundary Layer Theory*. New York: McGraw-Hill, 1979.

[98] A. Gosman, B. Launder, F. Lockwood, and G. Reece, "National development progam in computer assisted learning," mechanical engineering dept., Imperial College of Science and Technology, London SW7 BX, 1975.

[99] D. B. Spalding, "N/a," *Int. J. Num. Meth. Engng.*, vol. 4, 1972.

[100] S. V. Patankar, *Numerical Heat Transfer and Fluid Flow*. New York: McGraw Hill, 1980.

[101] "Storm prediction center." `http://www.spc.noaa.gov/archive/hail`.

[102] O. Ben-Tal, J. Berger, B. Cook, M. Daniels, G. Scavone, and P. Cook, "Sonart: The sonification application research toolbox," in Nakatsu and Kawahara [107].

[103] M. Fernström, N. Griffith, and S. Taylor, "Bliain le baisteach - sonifying a year with rain," in Hiipakka *et al.* [108].

[104] S. Barrass and B. Zehner, "Responsive sonfication of well-logs," in *Proceedings of the 2000 International Conference on Auditory Display*, (Atlanta, Georgia), 2000.

[105] R. F. Voss and J. Clarke, "$1/f$ noise in music: Music from $1/f$ noise," *J. Acoust. Soc. Am*, vol. 63, January 1978.

[106] K. V. Nesbitt and S. Barrass, "Evaluation of a multimodal sonification and visualization of depth of stock market data," in Nakatsu and Kawahara [107].

[107] R. Nakatsu and H. Kawahara, eds., *Proceedings of the 8th International Conference on Auditory Display*, (Kyoto, Japan), 2002.

[108] J. Hiipakka, N. Zacharov, and T. Takala, eds., *Proceedings of the 7th International Conference on Auditory Display*, (Helsinki University of Technology, Espoo, Finland), 2001.

A Source Code and Implementation

A.1 Implementation of VBAP in Java JSyn

In order to implement VBAP in JSyn, two steps were necessary:

1. Write an interface from Java to the VBAP code, which was written in C.

2. Modify some of the JSyn Classes so that they would work correctly in a multichannel environment.

A.1.1 Interface from Java to VBAP

Subroutines written in other languages such as C or C++ may be called from Java, however, they must be compiled as native libraries, using the Java Native Interface (JNI), and then called in a specific way from Java. There are two modules which must be accessed in this way. The first is `DLMethod`, which sets up the loudspeaker locations once and for all at the beginning of the calculation (DL stands for Define Loudspeakers). The second is `VBMethod` which takes as inputs the azimuth, elevation and desired spread for the apparent source, and which returns the gains for each output channel needed to create that apparent source in the appropriate loudspeaker triangle. The Java "wrapper" classes for these native method calls are shown below:

```
package com.mw.example;

public class HelloWorld
{
  public native void DLMethod(int dim, double[]
   ls_dirs, double[] ls_sets);
}

package com.mw.example;
```

```
public class VBap
{
  public native void VBMethod(int azi, int elev,
   int spread, double[] ls_sets, double[] gains);
}
```

The original C source code was edited to provide entry points for the native calls:

```
JNIEXPORT void JNICALL
Java_com_mw_example_HelloWorld_DLMethod
(JNIEnv *j, jobject o, jint dim,
 jdoubleArray ls_dirs, jdoubleArray ls_sets)
JNIEXPORT void JNICALL
Java_com_mw_example_VBap_VBMethod
(JNIEnv *j, jobject o, jint azi, jint elev, jint spread,
jdoubleArray ls_sets, jdoubleArray ls_gains)
```

The native library must be built. It is called from inside the Java application as follows:

```
static {
try {
System.loadLibrary("helloworld"); }

catch(Throwable t){t.printStackTrace();}
}
```

A.1.2 Modification of JSyn Classes

In order to work with ASIO-compatible devices, a special native library has been built and is available in the restricted area at `www.softsynth.com`. A folder called exactly "ASIO Drivers" (for the Macintosh) must be present in the working directory of the Java application, and the appropriate driver for the specific device must be placed in the folder.

All JSyn Classes must be instantiated with `synthContext` in the constructor. The author found that the `SynthMixer` and `FMOperator` objects did not have constructors for `synthContext`. This was added by editing the source code and making the change:

```
public SynthMixer2(SynthContext synthContext,
 int   numInputs, int   numOutputs )
throws SynthException
{
super(synthContext);
this.numInputs = numInputs;
this.numOutputs = numOutputs;
nodes = new MultiplyAddUnit[numInputs][numOutputs];

for( int j=0; j<numOutputs; j++ )
{
for( int i=0; i<numInputs; i++ )
{
add( nodes[i][j] = new MultiplyAddUnit(synthContext) );
/* Daisy chain MACs */
if( i>0 )
{
nodes[i-1][j].output.connect( nodes[i][j].inputC );
}
}
}
}
public FMOperator2(SynthContext synthContext) throws SynthException
{
// FIXME - we break up the creation
```

```
// into two stages so we can over ride port assignments
// before calling addPort().
super(synthContext);
osc = new SineOscillator(synthContext);
freqMixer  = new MultiplyAddUnit(synthContext);
ampMixer  = new MultiplyAddUnit(synthContext);
freqScalar = new MultiplyUnit(synthContext);
envPlayer  = new EnvelopePlayer(synthContext);
makeCircuit();
addAllPorts();
}
```

Bibliography

`http://eamusic.dartmouth.edu/~ed/research.html`.

Unknown, 1999. Unpublished lecture at Conductor's Guild Conference.

Stokes, G. G., "N/A," *Trans. Camb. Phil. Soc.*, vol. 8, 1845.

Read, T. L., 1999. Personal communication.

Scaletti, C., 2003. Personal communication.

Spalding, D. B., "N/A," *Int. J. Num. Meth. Engng.*, vol. 4, 1972.

Navier, C. L. M. H., "N/A," *Memoires Acad. R. Sci.*, vol. 6, 1830.

"JSyn Audio Synthesis," `http://www.softsynth.com`.

"Storm Prediction Center," `http://www.spc.noaa.gov/archive/hail`.

"Iannis Xenakis (1922-2001)," `http://www.bbc.co.uk/music/profiles/print/xenakis.shtml`.

Blauert, J., *Spatial Hearing.* Cambridge, MA: The MIT Press, revised ed., 1997.

Matossian, N., *Xenakis.* London and White Plains, New York: Kahn & Averill and Pro/Am Music Resources, Inc, 1991.

Penrose, C., "MK Ultra," in *American Jingo*, `www.detritus.net/illegalart`: illegal art, 2000.

Eckel, B., *Thinking in Java.* Upper Saddle River, New Jersey 07458: Prentice Hall, 2000.

Boulanger, R., ed., *The Csound Book.* Cambridge, MA: The MIT Press, 2000.

Chopin, F., "'Etude N. 12 Op.25," in *12 studies Op. 25 for piano*, (Cortot, A., ed.), pp. 84–91: Salabert Editions, 1914.

Berendt, J.-E., *The Third Ear.* New York: Henry Holt and Company, 1985.

Johnson, T., *Self-Similar Melodies*. 75 rue de la Roquette, 75011 Paris: Editions 75, 1996.

Gann, K., "Custer and Sitting Bull," Unpublished Score, 1998.

Pennington, J. V., "The Red and White Cows," *American Mathematical Monthly*, vol. 2, 1957.

Adler, S., *The Study of Orchestration*. New York and London: W. W. Norton and Company, second ed., 1989.

Dodge, C., *Earth's Magnetic Field*. 15 Columbus Circle, New York, NY 10023: Nonesuch Records, 1970.

Dodge, C., "Earth's Magnetic Field," in *Columbia Princeton Electronic Music Center, 1961-1973*: New World Records, 1998.

Schlichting, H., *Boundary Layer Theory*. New York: McGraw-Hill, 1979.

Schaeffer, P., "Music and Computers," in *Music and Technology*, (Paris), UNESCO, La Revue Musical, June 1970. Stockholm Meeting June 8-12, 1970.

Xenakis, I., "The Crisis of Serial Music," *Gravesaner Blätter*, 1955.

Näätänen, R., *Attention and Brain Function*. Hillsdale: Lawrence Erlbaum Associates, 1992.

Gardner, K., *Sounding the Inner Landscape*. Rockport, Massachusetts: Element, 1990.

Bregman, A. S., *Auditory Scene Analysis*. Cambridge, Massachusetts: The MIT Press, 1994.

Monro, G., "In-phase corrections for ambisonics," in *Proceedings of the International Computer Music Conference*, (Berlin, Germany), 2000.

"The International Conference on Auditory Display," `http://www.icad.org`.

"CCMIX - Center for Composition Iannis Xenakis," http://www.ccmix.com. formerly Les Ateliers UPIC.

Varga, B. A., *Conversations with Iannis.* 3 Queen Square London WC1N 3AU: Faber & Faber, 1996.

Cook, P. and Scavone, G., "The Synthesis Tool Kit," in *Proceedings of the 1999 International Computer Music Conference*, (Beijing, China), pp. 164–166, 1999.

Dodge, C. and Jerse, T. A., *Computer Music.* New York: Schirmer Books, 1997.

Chowning, J., "The simulation of moving sound images," *Journal of the Audio Engineering Society*, vol. 19, 1971.

Gaver, W. W., "Using and Creating Auditory Icons," in *Auditory Display*, (Kramer, G., ed.), (Jacob Way, Reading, MA 01867), pp. 417–446, Santa Fe Institute, Addison-Wesley Publishing Company, 1994.

Willingham, D. B., *Cognitions: The Thinking Animal.* Upper Saddle River, New Jersey 07458: Prentice Hall, 2001.

Neuwirth, E., "R sings – or using R to sonify data," in *DSC 2001 Proceedings of the 2nd International Workshop on Distributed Statistical Computing*, (Vienna, Austria), 2001.

Patankar, S. V., *Numerical Heat Transfer and Fluid Flow.* New York: McGraw Hill, 1980.

Garland, P., ed., *SOUNDINGS 13: The Music of James Tenney.* Santa Fe: Sounding Press, 1984.

Harley, J., "Iannis Xenakis Bibliography and Discography," LEONARDO on-line http://mitpress2.mit.edu/e-journals/Leonardo/isast/spec.projects/Xenaki%sbib.html.

Begault, D. R., *3-D Sound for Virtual Reality and Multimedia*. Cambridge, MA: AP Professional, 1994.

Tufte, E. R., *The Visual Display of Quantitative Information*. Box 430, Cheshire, Connecticut 06410: Graphics Press, 1983.

Williams, S. M., "Perceptual Principles in Sound Grouping," in *Auditory Display*, (Kramer, G., ed.), (Jacob Way, Reading, MA 01867), pp. 95–126, Santa Fe Institute, Addison-Wesley Publishing Company, 1994.

Barrass, S. and Zehner, B., "Responsive Sonfication of Well-logs," in *Proceedings of the 2000 International Conference on Auditory Display*, (Atlanta, Georgia), 2000.

Dombois, F., "Using Audification in Planetary Seismology," in *Proceedings of the 7th International Conference on Auditory Display*, (Hiipakka, J., Zacharov, N., and Takala, T., eds.), (Helsinki University of Technology, Espoo, Finland), 2001.

Wilson, C. M. and Lodha, S. K., "Listen: A Data Sonification Toolkit," in *Proceedings of the 1996 International Conference on Auditory Display*, (Palo Alto, CA), pp. 1119–1134, 1996.

Neilson, W. A., ed., *Webster's New International Dictionary*. Springfield, MA, USA: G. & C. Merriam, Co, unabridged second edition ed., 1947.

Kramer, G., ed., *International Conference on Auditory Display*, number XVIII in Studies in the Sciences of Complexity, (Jacob Way, Reading, MA 01867), Santa Fe Institute, Addison-Wesley Publishing Company, 1994.

Arsenault, L., "Iannis Xenakis's *Achorripsis*: The Matrix Game," *Computer Music Journal*, vol. 26, pp. 58–72, Spring 2002.

Cariani, P., "Temporal codes, timing nets and music perception," 2001. For Journal of New Music Perception, vol. 30, issue 2 (2001).

Childs, E., "The Sonification of Numerical Fluid Flow Simulations," in *Proceedings of the 7th International Conference on Auditory Display*, (Hiipakka, J., Zacharov, N., and Takala, T., eds.), (Helsinki University of Technology, Espoo, Finland), 2001.

Scaletti, C., "Using sound to extract meaning from complex data," in *Extracting Meaning from Complex Data: Processing, Display, Interaction II*, International Society of Optical Engineering, 1991.

Polansky, L., "Cocks Crow, Dogs Bark: New Compositional Intentions," *Leonardo Music Journal*, vol. 7, pp. 61–71, 1997.

Polansky, L., "Live Interactive Computer Music in HMSL, 1984-1992," *Computer Music Journal*, vol. 18, pp. 59 – 77, 1994.

Xenakis, I., *Formalized Music: Thought and Mathematics in Music.* Hillsdale, NY: Pendragon Press, 1992.

Voss, R. F. and Clarke, J., "$1/f$ noise in music: Music from $1/f$ noise," *J. Acoust. Soc. Am*, vol. 63, January 1978.

Madhyastha, T. M. and Reed, D. A., "A Framework for Sonification Design," in *Auditory Display*, (Kramer, G., ed.), (Jacob Way, Reading, MA 01867), pp. 267–289, Santa Fe Institute, Addison-Wesley Publishing Company, 1994.

Childs, E., "ACHORRIPSIS: A Sonification of Probability Distributions," in *Proceedings of the 8th International Conference on Auditory Display*, (Nakatsu, R. and Kawahara, H., eds.), (Kyoto, Japan), 2002.

Kaltenbrunner, M., "Y-Windows: Proposal for a Standard AUI Environment," in *Proceedings of the 8th International Conference on Auditory Display*, (Nakatsu, R. and Kawahara, H., eds.), (Kyoto, Japan), 2002.

Malham, D. G. and Myatt, A., "3-d sound spatialization using ambisonic techniques," *Computer Music Journal*, vol. 19, 1995.

Hayward, C., "Sound Synthesis Algorithms for Auditory Data Representations," in *Auditory Display*, (Kramer, G., ed.), (Jacob Way, Reading, MA 01867), pp. 369–404, Santa Fe Institute, Addison-Wesley Publishing Company, 1994.

Kramer, G., "Some Organizing Principles for Representing Data with Sound," in *Auditory Display*, (Kramer, G., ed.), (Jacob Way, Reading, MA 01867), Santa Fe Institute, Addison-Wesley Publishing Company, 1994.

Scaletti, C., "Sound Synthesis Algorithms for Auditory Data Representations," in *Auditory Display*, (Kramer, G., ed.), (Jacob Way, Reading, MA 01867), pp. 223–251, Santa Fe Institute, Addison-Wesley Publishing Company, 1994.

Pulkki, V., "Virtual source positioning using vector base amplitude panning," *Journal of the Audio Engineering Society*, vol. 45, June 1997.

Barrass, S., "Personify: a toolkit for perceptually meaningful sonification," in *Proceedings of the Australian Computer Music Conference ACMA '95*, (Melbourne, Australia), 1995.

McCabe, K. and Rangwalla, A., "Auditory Display of Computational Fluid Dynamics Data," in *Auditory Display*, (Kramer, G., ed.), (Jacob Way, Reading, MA 01867), pp. 327–340, Santa Fe Institute, Addison-Wesley Publishing Company, 1994.

McCabe, K. and Rangwalla, A., "Auditory Display of Computational Fluid

Dynamics Data," in *Auditory Display*, (Kramer, G., ed.), (Jacob Way, Reading, MA 01867), Santa Fe Institute, Addison-Wesley Publishing Company, 1994.

Berkhout, A. J., de Vries, D., and Vogel, P., "Acoustic control by wave field synthesis," *Journal of the Acoustical Society of America*, vol. 93, May 1993.

Schellenberg, G., "Name that tune: Identifying popular recordings from brief excerpts," *Psychonomic Bulleting & Review*, 1999.

Mendenhall, W. and Beaver, R. J., *Introduction to Probability and Statistics*. Belmont, California: Duxbury Press, 1994.

Pulkki, V., "Microphone techniques and directional quality of sound reproduction," in *112th Convention of the Audio Engineering Society*, (Munich, Germany), 2002.

Kramer, G. and et al., "Sonification Report: Status of the Field and Research Agenda," `http://www.icad.org/websiteV2.0/References/nsf.html`, 1997.

Chowning, J. M., "The Synthesis of Complex Audio Spectra by Means of Frequency Modulation," *Journal of the Audio Engineering Society*, vol. 21, 1973.

Pulkki, V., *Spatial sound generation and perception by amplitude panning techniques*. PhD thesis, Helsinki University of Technology, 2001.

Ronkainen, S., "Earcons in motion – defining language for an intelligent mobile device," in *Proceedings of the 7th International Conference on Auditory Display*, (Hiipakka, J., Zacharov, N., and Takala, T., eds.), (Helsinki University of Technology, Espoo, Finland), pp. 126–131, 2001.

Janata, P., "Neurophysiological mechanisms underlying auditory image formation in music," in *Elements of Musical Imagery*, (God/oy, R. I. and J/orgensen, H., eds.), pp. 27–42, Lisse: Swets & Zeitlinger, 2001.

Fernström, M., Griffith, N., and Taylor, S., “Bliain Le Baisteach - Sonifying a Year With Rain,” in *Proceedings of the 7th International Conference on Auditory Display*, (Hiipakka, J., Zacharov, N., and Takala, T., eds.), (Helsinki University of Technology, Espoo, Finland), 2001.

Tzelgov, J., Srebro, R., Henik, A., and Kushelevsky, A., “Radiation detection by ear and by eye,” *Human Factors*, vol. 29, 1987.

Fromm, J. E. and Harlow, F. H., “Numerical solutions of the problem of vortex street development,” *Physics of Fluids*, vol. 6, 1963.

Gilkey, R. H. and Anderson, T. R., eds., *Binaural and Spatial Hearing in Real and Virtual Environments*. Mahwah, NJ: Lawrence Erlbaum Assoc., 1997.

Squibbs, R., *An Analytical Approach to the Music of Iannis Xenakis: Studies of Recent Works*. PhD thesis, Yale University, 300 North Zeeb Road, Ann Arbor, MI 48103, November 1996. UMI Microform 9714309.

Joseph, A. and Lodha, S. K., “MUSART: Musical Audio Transfer Function Real-Time Toolkit,” in *Proceedings of the 8th International Conference on Auditory Display*, (Nakatsu, R. and Kawahara, H., eds.), (Kyoto, Japan), 2002.

Xenakis, I. and Kanach, t. S., *Arts/Sciences: Alloys, The Thesis Defense of Iannis Xenakis*. Hillsdale, NY: Pendragon Press, 1985.

Wickens, C. D., Gordon, S. E., and Liu, Y., *An Introduction to Human Factors Engineering*. New York: Longman, 1998.

Childs, E., “Using Multi-Channel Spatialization in Sonification: A Case Study with Meteorological Data,” in *Proceedings of the 9th International Conference on Auditory Display*, (Shinn-Cunningham, B., ed.), (Boston, MA), 2003.

Mayer-Kress, G., Bargar, R., and Choi, I., “Musical Structures in Data from

Chaotic Attractors," in *Auditory Display*, (Kramer, G., ed.), (Jacob Way, Reading, MA 01867), pp. 341–368, Santa Fe Institute, Addison-Wesley Publishing Company, 1994.

Dombois, F., "Auditory Seismology: On Free Oscillations, Focal Mechanisms, Explosions and Synthetic Seismograms," in *Proceedings of the 8th International Conference on Auditory Display*, (Nakatsu, R. and Kawahara, H., eds.), (Kyoto, Japan), 2002.

Fitch, W. T. and Kramer, G., "Sonifying the Body Electric: Superiority of an Auditory over a Visual Display in a Complex, Multivariate System," in *Auditory Display*, (Kramer, G., ed.), (Jacob Way, Reading, MA 01867), pp. 307–326, Santa Fe Institute, Addison-Wesley Publishing Company, 1994.

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

Tillman, B., Bharucha, J. J., and Bigand, E., "Implicit Learning of Tonality: A Self-Organizing Approach," *Psychological Review*, vol. 107, pp. 885–913, 2000.

Patel, A. D., Gibson, E., Ratner, J., Besson, M., and Holcomb, P. J., "Processing Syntactic Relations in Language and Music: An Event-Related Potential Study," *Journal of Cognitive Neuroscience*, vol. 10, pp. 717–733, 1998.

Faïta, F. and Besson, M., "An event-related potential (ERP) study of musical expectancy: Comparison of musicians with nonmusicians," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 21, pp. 1278–1296, 1995.

Janata, P., Tillmann, B., and Bharucha, J. J., "Listening to polyphonic music recruits domain-general attention and working memory circuits," *Cognitive, Affective and & Behavioral Neuroscience*, vol. 2, pp. 121–140, 2002.

Appleton, J., "Machine Songs III: Music in the Service of Science – Science in the Service of Music," *Computer Music Journal*, vol. 16, Fall 1992.

Harley, J., "Formal Analysis of The Music of Iannis Xenakis by Means of *Sonic Events*: Recent Orchestral Works," in *Présences de Iannis Xenakis*, (Solomos, M., ed.), 16 place del la Fontaine aux Lions, 75019 Paris, France: Centre de documentation de las musique contemporaine, 2001. `http://www.cdmc.asso.fr`.

Bartels, J., *Collection of geomagnetic planetary indices K_p and derived daily indices A_p and C_p, for the years 1932 to 1961.* Amsterdam: North-Holland Pub. Co., 1962.

Ben-Tal, O., Berger, J., Cook, B., Daniels, M., Scavone, G., and Cook, P., "SONART: The Sonification Application Research Toolbox," in *Proceedings of the 8th International Conference on Auditory Display*, (Nakatsu, R. and Kawahara, H., eds.), (Kyoto, Japan), 2002.

Pampin, J., "ATS: A System for Sound Analysis, Transformation and Synthesis Based on a Sinusoidal Plus Critical Band Noise Model and Psychoacoustics," 2003. PAPER/PANEL 1, National Conference of the Society for Electro-Acoustic Music in the United States.

Møller, H., Sørensen, M. F., Hammershøi, D., and Jensen, C. B., "Head-related transfer functions of human subjects," *Journal of the Audio Engineering Society*, vol. 43, May 1995.

Bennett, J. C., Barker, K., and Edeko, F. O., "A new approach to the assessment of stereophonic sound system performance," *Journal of the Audio Engineering Society*, vol. 33, May 1985.

Pulkki, V., "Creating auditory displays to multiple loudspeakers using vbap: A case study with diva project," in *International Conference on Auditory Display*, (Glasgow, Scotland), 1998.

Gosman, A., Launder, B., Lockwood, F., and Reece, G., "National Development Progam in Computer Assisted Learning," Mechanical engineering dept., Imperial College of Science and Technology, London SW7 BX, 1975.

Ben-Tal, O., Berger, J., Cook, B., Daniels, M., Scavone, G., and Cook, P., "SONART: The Sonification Application Research Toolbox," in *Proceedings of the 8th International Conference on Auditory Display*, (Nakatsu, R. and Kawahara, H., eds.), (Kyoto, Japan), 2002.

Nesbitt, K. V. and Barrass, S., "Evaluation of a Multimodal Sonification and Visualization of Depth of Stock Market Data," in *Proceedings of the 8th International Conference on Auditory Display*, (Nakatsu, R. and Kawahara, H., eds.), (Kyoto, Japan), 2002.

Nakatsu, R. and Kawahara, H., eds., *Proceedings of the 8th International Conference on Auditory Display*, (Kyoto, Japan), 2002.

Hiipakka, J., Zacharov, N., and Takala, T., eds., *Proceedings of the 7th International Conference on Auditory Display*, (Helsinki University of Technology, Espoo, Finland), 2001.