

Sonification Report: Status of the Field and Research Agenda

**Prepared for the National Science Foundation by members of the International
Community for Auditory Display**

Editorial Committee and Co-Authors

Gregory Kramer, Chair; Bruce Walker, Project Coordinator;
Terri Bonebright; Perry Cook; John Flowers; Nadine Miner; John Neuhoff

Co-Authors

Robin Bargar, Stephen Barrass, Jonathan Berger, Grigori Evreinov, W. Tecumseh Fitch,
Matti Gröhn, Steve Handel, Hans Kaper, Haim Levkowitz, Suresh Lodha,
Barbara Shinn-Cunningham, Mary Simoni, Sever Tipei

ABSTRACT

The purpose of this paper is to provide an overview of sonification research, including the current status of the field and a proposed research agenda. This paper was prepared by an interdisciplinary group of researchers gathered at the request of the National Science Foundation in the fall of 1997 in association with the International Conference on Auditory Display (ICAD).

Contents

CONTENTS	I
<hr/>	
1 EXECUTIVE SUMMARY	1
<hr/>	
2 INTRODUCTION	3
2.1 MOTIVATION	3
2.2 OBJECTIVES	4
<hr/>	
3 OVERVIEW OF KEY SONIFICATION COMPONENTS	5
<hr/>	
3.1 PERCEPTUAL RESEARCH IN SONIFICATION	5
3.1.1 PAST RESEARCH IN AUDITORY PERCEPTION	5
3.1.2 CURRENT RESEARCH TRENDS IN AUDITORY PERCEPTION	6
3.1.3 SUMMARY AND ANALYSIS OF PERCEPTUAL ISSUES IN SONIFICATION	7
3.2 TOOLS FOR SONIFICATION	8
3.2.1 CURRENTLY AVAILABLE SONIFICATION TOOLS	9
3.2.2 CURRENT TRENDS IN DEVELOPMENT OF SONIFICATION TOOLS	10
3.2.3 SUMMARY AND ANALYSIS OF CURRENT TRENDS IN SONIFICATION TOOLS	11
3.3 SONIFICATION APPLICATION AND DESIGN	12
3.3.4 SUMMARY OF SONIFICATION AND DESIGN	15
<hr/>	
4 GLOBAL ISSUES FOR ESTABLISHMENT OF A DISCIPLINE OF SONIFICATION	16
<hr/>	
4.1 RECOGNITION	16
4.2 COMMUNICATION AND COMMUNITY FORMATION	16
4.3 EDUCATION	17
4.4 LESSONS LEARNED FROM VISUALIZATION	18
4.5 SUMMARY OF GLOBAL ISSUES	19
<hr/>	
5 PROPOSED RESEARCH AGENDA	20
<hr/>	
5.1 GENERAL FUNDING RECOMMENDATIONS	20
5.2 PERCEPTION AND COGNITION	21
5.3 DEVELOPING SONIFICATION TOOLS	22
5.4 SONIFICATION APPLICATIONS, DESIGN AND THEORY	24
5.5 SUPPORT FOR THE FIELD AT LARGE	25
<hr/>	
6 GENERAL CONCLUSION	26
<hr/>	

REFERENCES

1 Executive Summary

Sonification is the use of nonspeech audio to convey information. The goal of this report is to provide the reader with (1) an understanding of the field of sonification, (2) an appreciation for the potential of sonification to answer a variety of scientific questions, (3) a grasp of the potential for sonification applications to facilitate communication and interpretation of data, and (4) specific ideas for productive support of sonification research.

The field is composed of the following three components: (1) psychological research in perception and cognition, (2) development of sonification tools for research and application, and (3) sonification design and application. In reviewing the current status of each of these components, some common themes become apparent. One is a trend toward research in high-level perceptual issues and development of corresponding complex tools. Another is the potential importance of multimodal displays. Finally, an overarching theme is the need for interdisciplinary research and interaction. By nature, the field of sonification is interdisciplinary, integrating concepts from human perception, acoustics, design, the arts, and engineering.

In order to establish a discipline of sonification, three global issues must be addressed. The first is the need for recognition of sonification as a valid area of research. The recognition and funding of sonification by the National Science Foundation (NSF) can play a major role in this validation. The second is communication within the sonification community. We propose support for coordinated workshops and conferences and a peer-reviewed journal of sonification. The final issue is the need to provide a curriculum for teaching sonification.

We recommend the following research agenda. Perception and cognition research should focus on dynamic sound perception, auditory scene analysis, multimodal interaction, and the role of memory and attention in extracting information from sound. The development of sonification tools should focus on providing the user with flexible control over data dimensions and sound parameters, facilitating data exchange to and from a variety of formats and display systems, and integrating a perceptual testing and evaluation framework. Applications and design research should focus on the formulation of a method for sonification design. In addition to funding promising flagship applications, task-dependent and user-centered approaches to sonification design should be supported. Timbre perception studies should be furthered and coupled with data-to-sound parameter-mapping research. Other worthy research topics in basic sonification theory and design research include aesthetics, metaphor, affect, and applications of gestalt formation.

A coordinated interdisciplinary research effort supported by moderate funding at the national level is necessary if sonification research is to prosper. The resultant advances in both basic research and technology development will contribute to scientific and commercial

applications, which will then feed back into the development of the field. National Science Foundation funding and leadership can help to accelerate this process.

2 Introduction

Sonification is defined as *the use of nonspeech audio to convey information*. More specifically, *sonification is the transformation of data relations into perceived relations in an acoustic signal for the purposes of facilitating communication or interpretation*. By its very nature, sonification is interdisciplinary, integrating concepts from human perception, acoustics, design, the arts, and engineering. Thus, development of effective auditory representations of data will require interdisciplinary collaborations using the combined knowledge and efforts of psychologists, computer scientists, engineers, physicists, composers, and musicians, along with the expertise of specialists in the application areas being addressed.

Success stories that predate the word “sonification” include the Geiger counter, sonar, the auditory thermometer, and numerous medical and cockpit auditory displays, particularly those designed to present data variations. More recent successes include software that enables blind chemists to examine infrared spectrographic data via auditory presentation (Lunney & Morrison, 1990) and the mapping of data-dependent auditory signals to ongoing processes in dynamic monitoring tasks such as anesthesiology workstations (Fitch & Kramer, 1994) or factory production controls (Gaver, Smith, & O’Shea, 1991). Potential future applications include novel ways of using sound to explore complex data sets, to supplement other sensory modalities for data communication, and to assist special user populations. These and other industrial and academic sonification efforts are discussed in this paper.

2.1 Motivation

Two recent developments indicate that support for sonification research is particularly relevant and timely: (1) the need to comprehend an abundance of data; and (2) increasingly powerful and available media technologies. Because of the increase in computing power over the past 20 years, scientists and researchers have generated ever-expanding amounts of data, which they need to interpret and understand. The recent rise in computational power is quickly transforming how we learn, communicate, and explore our world. New projects, such as the Sloan Digital Sky Survey (<http://www.sdss.org/>) and the Human Genome Project (<http://www.ornl.gov/hgmis/>), are yielding huge data sets that must be managed and explored. In addition, general computer users, including those in the applied social sciences, business, and government, must increasingly grapple with large, complex, abstract data sets for decision making and information discovery. The field of scientific visualization emerged to assist scientists and researchers in analyzing such large volumes of numerical data. However, scientific visualization techniques are often insufficient for comprehending certain features in the data (as demonstrated by the Voyager 2 data analysis and Quantum Whistle discovery discussed in the *Application and Design* section of this paper). Although scientific visualization techniques may

not yet be exhausted, some believe that we are approaching the limits of users' abilities to interpret and comprehend visual information. Audio's natural integrative properties are increasingly being proven suitable for presenting high-dimensional data without creating information overload for users. Furthermore, environments in which large numbers of changing variables and/or temporally complex information must be monitored simultaneously are well suited for auditory displays. Sonification research is well positioned to provide technology to assist scientists in comprehending the information and data-rich world of today and of the future.

Concurrent with the flood of data has been the emergence of powerful audio technologies, which are readily available across a wide range of computer platforms. Scientific visualization was in a similar situation a decade ago with the rapid development of computer graphics technology. The field of sonification is now in a position to leverage the new computer audio technology to solve many existing problems of scientific display. The wide availability of audio technology (e.g., in multimedia computers) makes auditory data representation a viable option for large numbers of users. Thus, there exists today a synergism between the widespread need for new data-comprehension methods and readily available technology that, with proper support and funding, can lead to a large number of users reaping the benefits conferred by the development of scientific sonification.

2.2 Objectives

The goal of this report is to provide the reader with (1) an understanding of the field of sonification, (2) an appreciation for the potential of sonification to answer a variety of scientific questions, (3) a grasp of the potential for sonification applications to facilitate communication and interpretation of data, and (4) specific ideas for productive support of sonification research. A key to the success of sonification research and applications is interdisciplinary interaction. We begin our discussion of sonification with an overview of three principal components of the field. First, we discuss the relevant psychological research related to auditory perception. Second, we discuss sonification tools for research and application development, with emphasis on sound synthesis and control software. Third, we review the status of sonification design and applications. In each of these sections, we review past sonification research and the current research trends as well as identifying the significant research needs. Thus prepared, we examine key global issues that currently impede sonification development. Specifically, we address issues in interdisciplinary collaboration, communication among researchers, education of researchers in sonification techniques, and the lessons to be learned from data visualization. Finally, we propose a comprehensive sonification research agenda, including example research questions and issues. By presenting the background and status of sonification research, discussing the major obstacles, and outlining the research opportunities and potential, we aim to provide a rationale for support and funding of the development of the emerging field of sonification.

3 Overview of Key Sonification Components

This section provides an overview of the three key sonification components: perception, research and development tools, and sonification design and applications. It is the combination of these components that creates the highly interdisciplinary field of sonification.

3.1 Perceptual Research in Sonification

As stated previously, sonification conveys information to a human listener by mapping data onto *perceived* relations in an acoustic signal. Thus, understanding the perception of sonified data is key to the success of any sonification application. This section reviews relevant perceptual research in sonification and outlines outstanding perceptual research issues.

3.1.1 Past Research in Auditory Perception

A rich history of research has provided valuable insight into the physiological, perceptual, and cognitive aspects of auditory perception for speech and relatively simple auditory events, such as pure tones and noise bursts. Much of this work has contributed to a functional knowledge base of various auditory thresholds, psychophysical scales, and models of auditory perception. In particular, there is a strong foundation of research regarding intensity, frequency, and temporal discrimination of static sounds (Hartmann, 1997; Moore, 1995, 1997). The determinants of pitch and loudness, the effects of masking (Gulick, Gescheider, & Frisina, 1989), and auditory localization abilities (Blauert, 1997) are also well understood. Recently, researchers have begun to extend their investigations into the perception of more complex, dynamic auditory patterns in speech and music, which is particularly relevant for sonification research (Bregman, 1990; Handel, 1989; McAdams & Bigand, 1993).

From this body of research, two basic features of auditory perception have been discovered that suggest sound can be effective for representing data in a variety of settings. First, auditory perception is particularly sensitive to temporal characteristics, or changes in sounds over time. Human hearing is well designed to discriminate between periodic and aperiodic events and can detect small changes in the frequency of continuous signals. This points to a distinct advantage of auditory over visual displays. Fast-changing or transient data that might be blurred or completely missed by visual displays may be easily detectable in even a primitive, but well-designed auditory display. Thus, sonification is likely useful for comprehending or monitoring complex temporal data, or data that is embedded in other, more static, signals. Second, unlike visual perception, perception of sound does not require the listener to be oriented in a particular direction. Auditory displays can therefore be used in situations where the eyes are already busy with another task. These characteristics make sound highly suitable for monitoring and alarm applications, particularly when these alarms may arise from many possible locations, or when visual attention may be diverted from the alarm location.

Other aspects of auditory perception bear on the promise of sound as a medium for data display and will help illuminate the optimal means of mapping data to specific dimensions of sound. These aspects include parallel listening (ability to monitor and process multiple auditory data sets), rapid detection (especially in high-stress environments), affective response (ease of learning and high engagement qualities) and auditory gestalt formation (discerning relationships or trends in data streams) (Kramer, 1994a). These complex phenomena build on existing psychoacoustic knowledge and are active areas in auditory perceptual research.

3.1.2 Current Research Trends in Auditory Perception

A research focus germane to the present discussion is the role of learning in auditory display efficacy. There are applications for which training is necessary to provide highly efficient performance. The special abilities of skilled sonar operators, for example, illustrate that learning can significantly enhance the efficiency with which auditory patterns can be discerned (Howard & Ballas, 1982). Assistive technology for blind users is another example where attainable skilled performance levels are more important than the level of performance achieved on first exposure (Earl & Levanthal, 1999). Since learning ties together basic perception with higher-level cognitive processes, a significant proportion of psychological studies are relevant. Further research is needed into how performance with auditory displays changes with practice.

Much of the past work in psychoacoustics has examined the perception of a single auditory dimension (e.g., pitch) in isolation. Auditory events in the real world, however, involve dynamic sounds that simultaneously change in frequency, intensity, and often location. The unique capabilities of the auditory system to use such covariation to define perceptual events and “scenes” (Bregman, 1990) can potentially be exploited to create meaningful and compelling data displays. Further basic research in dynamic sound perception is necessary to improve our understanding of these capabilities (Neuhoff, 1998).

Sample Research Question #1. A sample problem in auditory display design relates to determining the appropriate mapping between data and sound features in a sonification display. For some applications it may be desirable to create mappings between data and sound features that are realistic or “natural,” in the hopes that they will be immediately compelling and comprehensible (e.g., a synthesized engine sound for an aircraft display). However, “natural” sounds may, in some cases, lack the number of discernible parameters necessary to represent a data set with many variables. Thus, should designers synthesize entirely novel sounds that are structured and can be manipulated, or should they attempt to incorporate sounds that are somehow familiar to the listener?

It has been argued that properly designed sonifications have the potential to increase the amount of data a human can simultaneously process beyond that achievable with traditional visual display technology (Scaletti and Craig, 1990). To achieve this goal, we must gain an understanding of how many different auditory information streams can be monitored without

loss, as well as how memory load, attention mechanisms, and other cognitive processes affect information transfer. Although there is a vast literature on auditory attention and selective listening, the overwhelming majority of this research has been concerned with speech and language rather than sonification.

Memory for auditory events poses both benefits and limitations compared with visual memory. Highly salient musical patterns can be easily recognized and recalled even when subjected to radical transformations. This property can be used to enhance visual-based data mining and pattern search tasks. However, sonification can be limited by the temporal constraints of memory in continuous or dense data sets. Although we cannot perceive a lengthy sonification “at a glance,” a multimodal approach can tap the positive features of each of the component sensory domains. Research has demonstrated that sound can enhance a visual or haptic display by providing an additional channel of information (e.g., Wickens, 1984; Wickens, Gordon, & Liu, 1998). Although such multimodal displays have promise, more research on the effect of supplementary audio specifically for data representation is required. A related issue is cross-modal interaction. This research examines how a visual event can change the perception of a sound and vice versa. The well-known example of a ventriloquist illustrates how visual events can alter perception of location of an auditory source (Radeau, 1992; Thurlow & Jack, 1973). Other examples include altering the perception of a speech phoneme by dubbing it onto a video of a speaker saying a different phoneme (McGurk & McDonald, 1976), and the alteration of the perception of timbre of a musical note depending upon whether it is “seen” as plucked or bowed (Saldaña & Rosenblum, 1993).

Eventual applications of sonification will likely include multimodal displays in which auditory and visual information either supplement each other or provide independent information. Cross-modal interactions therefore will need to be considered. Increasing our knowledge of haptic-auditory interactions (which is currently much more sparse than our knowledge of audition and vision) will likely be critical in the design of effective multimodal displays for blind users. Finally, multimodal displays will not always provide the best application solution. In fact, the combination of visual and auditory information may prove less effective in some circumstances than information presented in one sensory modality (Tzelgov, Srebro, Henik, & Kushelevsky, 1987). Thus, experimentation involving the perceptual interactions of sonified data used in conjunction with other sensory modalities is important to the success of multimodal applications.

3.1.3 Summary and Analysis of Perceptual Issues in Sonification

Research in auditory perception has progressed from the study of individual auditory dimensions, such as pitch, tempo, loudness, and localization, to the study of more complex

phenomena, such as auditory streaming, dynamic sound perception, auditory attention, and multimodal displays. Many of the major current research areas in sonification are similar in that they focus on the identification of applications for which audition provides advantages over other modalities, especially for situations where temporal features are important or the visual modality is overtaxed. Applications are currently being explored for both normally sighted and visually impaired user populations. The main issues that will drive sonification research forward include (1) mapping data onto appropriate sound features, (2) understanding dynamic sound perception, (3) investigating auditory streaming, (4) defining and categorizing salience in general auditory contexts and understanding where highly salient sonic events or patterns can surpass visual representations in data mining, and (5) developing multimodal applications of sonification.

Through the use of complex nonspeech audio and sophisticated multimodal displays, sonification has the potential to advance basic research in cognition and perception in important ways. This research, in turn, will provide a foundation upon which solid tools and applications can be built and evaluated. Regardless of the particular perceptual research or area of application, advances in sonification depend upon the existence of flexible and usable tools for sound production and display, some of which are discussed in the next section.

3.2 Tools for Sonification

In many ways, a “sonification session” will resemble work with corresponding data visualization tools. Consider the following scenario. A scientist using a sonification system sits at a workstation and listens as her data is used to control a sound synthesis system. As various sound parameters change, she listens for patterns or anomalies. She focuses on one section of the data, looping a section or replaying various data points. She then decides to change which data variable controls which auditory variable and listens to the data again. One of her colleagues, a seismologist, finds it most effective to listen to a data set simply by frequency-shifting the data into the audible domain (see Hayward, 1994). He then filters out or compresses the dynamics of selected frequencies and listens again to the thump or rumble.

The tools for sonification include both hardware (e.g., audio recording, signal processing, and playback equipment) and software (e.g., sound synthesis, editing, and analysis tools). Sound hardware is now commonly found on desktop and personal computers, and a variety of inexpensive or free software sound tools also exist. The combination of these tools, along with custom and domain-specific hardware and software, provides a strong starting point for the development of sonification tools. This section gives an overview of currently available tools, current trends, and outstanding research needs.

3.2.1 Currently Available Sonification Tools

A variety of approaches exist for handling sound across hardware platforms that range from PCs to parallel supercomputers. Available software includes everything from free software for sound wave editing, to MIDI (Musical Instrument Digital Interface) controllers, music authoring programs, and sophisticated signal analysis and synthesis packages. While some tools provide a response in real time, others offer more detailed and flexible control off-line. Generally speaking, sonification tools do not require as much processing power or special-purpose hardware as do 3-D visualization tools.

For sound editing, many software packages provide basic cut, copy, paste, record, playback, and looping functions, as well as a variety of sound effects useful for the production of music and audio tracks. MIDI is a standard music protocol created in 1984 that provides sound control via basic performance gesture messages, such as note onsets and offsets and continuous controller changes. All commercial music synthesizers and nearly all personal computers now support MIDI. The most popular tools, called “sequencers,” allow the recording, playback, and manipulation of MIDI data on the computer screen.

General mathematical analysis programs, such as Matlab, Mathematica, or Maple, are sometimes used to do research on signal processing and sound synthesis algorithms. However, most commercially available signal analysis packages do not support real-time signal processing, and some do not support direct control of sound input or output.

There do exist a variety of software tools for sound analysis, synthesis, manipulation, and control, which have been developed by academic centers specializing in computer music, computer science, and electrical engineering. Examples of these tools include CSound (MIT), CMix (Princeton), CLM (Stanford CCRMA), and KYMA and SuperCollider (commercial products). Each of these systems provides a rich set of functionality for creating, manipulating, and controlling sound, but each has limitations based on the original design, target user community, or the need for specialized hardware.

Embedded within many of these tools is the capability for sound creation. Pure tones and sequences of tones are created at the lowest level through playback of sinusoids. MIDI provides the ability to instruct audio hardware or software to play a particular frequency tone with onset, offset, and duration control. Other methods are employed for creating more complex sounds. Traditional methods include additive synthesis (Mathews, 1969), subtractive synthesis (Pierce, 1992), and frequency modulation (Chowning, 1973).

An important recent development has been an increase in the availability of systems that synthesize spatial sound attributes over headphones (Carlile, 1996; Wenzel, 1992). The most effective such systems use Head-Related Transfer Functions (HRTFs) to simulate the acoustic cues used for spatial localization, including interaural time and intensity differences, and spectral cues arising from the listener’s own head and ears. A number of commercially available systems

now allow spatial auditory displays to be realized fairly easily. In general, these systems are reasonably good at simulating the horizontal position of a sound source, but their ability to simulate vertical position and distance is still much less robust (Wenzel, Arruda, Kistler, & Wightman, 1993).

3.2.2 Current Trends in Development of Sonification Tools

Three major trends characterize sonification tools research. First is the continued effort to provide access to low-level signal processing functions to nonexperts, resulting in more powerful hardware and software sound-processing packages. Unfortunately, there is not a large trend toward providing standardized tools, and tools developed in academic environments have been poorly documented and unsupported. Thus, even though low-level tools are available, tool users such as perception researchers and application designers still struggle to fit pieces of the puzzle together across the many diverse systems.

A second area of research relevant to sonification tools is sound synthesis. Synthesis techniques fall into two general classes of sound control methods— one based on creating, replicating, or transforming the effect of a sound; and the other based on modeling the physical properties of a sound-producing object. The two categories suggest possibilities of sonification with virtually any degree of abstraction in terms of the taxonomic or metaphoric approach to mapping the sound production method to the task. There are many strategies for synthesizing sound and relating control over these sounds to the data under consideration. The first sound synthesis category, essentially a perceptually based one, comprises numerous techniques including additive synthesis, nonlinear synthesis (such as amplitude or frequency modulation), filtering of digitally sampled sound, and Fourier-based analysis and resynthesis techniques. The second category of physical modeling uses digital waveguides to simulate the complex interactions of strings, pipes, and other vibrating and resonating objects and the forced air, plectra, hammers, springs, or scraped bows that typically instantiate or interact with these objects in musical situations (Karplus & Strong, 1983; Smith, 1996; Cook, 1995). Spectral synthesis has also been successful for synthesizing musical instruments (Serra, 1989). Realistic impulsive or percussive sounds have been synthesized by digital waveguides (Pierce & Van Duyne, 1997), with simplified physical models (Gaver, 1994; Van Den Doel & Pai, 1996), and through a physically-informed spectral additive modeling approach (Cook, 1997). Real-world sounds have been synthesized with a variety of techniques, including timbre trees for synthesizing swarms of bees and turbulence (Takala & Hahn, 1992), massively parallel additive synthesis (Freed, Rodet, & Depalle, 1993), and wavelet-based models for synthesizing sounds such as rain and car engines (Miner, 1998).

Research in spatial auditory displays is a third active area. In order to provide adequate perception of sound source elevation (i.e., vertical position) and to avoid unreasonable increases

in localization errors, individualized HRTFs must be used in current displays (Wenzel, Arruda, Kistler, & Wightman, 1993). Several laboratories are developing more efficient ways to custom fit these synthesis techniques to individual listeners. Other researchers are focusing on the development of more reliable methods for providing distance cues (Brungart, 1998; Duda and Martens, 1998), including cues for source movement (Jenison, Neelon, Reale, & Brugge, 1998), increasing the computational efficiency of the synthesis algorithms (e.g., Carlile, 1996), and investigating trade-offs between realism in the display and localization accuracy (Shinn-Cunningham, 1998).

Together with this applied synthesis and signal processing research a particularly relevant domain of current inquiry is perceptually based synthesis and signal encoding. Current research in perceptually based data reductive signal encoding exemplifies the applied value of audio perception studies. Similarly, incorporating knowledge of psychoacoustics with sonification-based models of timbre classification (Grey, 1978; Martens, 85; Barrass, 1996; Krimphoff, 1993; Wessel, 1979) offer a broad range of approaches, including multi dimensional scaling (Grey and Moorer, 1978; Hajda et al, 1999), neural networks (Wessel, 1997), and fuzzy classifiers (Kieslar, 1996) as groundwork for innovative methods of linking sound and data.

Synthesis has recently become computationally efficient, thus allowing for real-time control. Together with the broad range of synthesis and processing methods comes a wide range of parametric controls. Harnessing these methods will involve creating intuitive controls and controllers. The limitless possibilities need to be organized in such ways as to provide distinct approaches to fit specific tasks.

3.2.3 Summary and Analysis of Current Trends in Sonification Tools

Although there has been significant progress in sound production hardware and software, and widespread distribution for consumer multimedia purposes, tailoring these products to sonification research and development presents significant hurdles. Ideally, researchers who are neither composers nor audio engineers would be able to produce detailed and intelligible data-driven sounds, manipulate sonification designs, and evaluate human performance associated with these designs. Current tools are too complex, specific, and unwieldy to allow these activities.

We identify three major issues in the tool development area that must be tackled to create appropriate synthesis tools developed for use by interdisciplinary sonification researchers.

Portability: Sonification scale places demands on audio hardware, on signal processing and sound synthesis software, and on computer operating systems. These demands may be more stringent than the requirements for consumer multimedia. Researchers dealing with problems that go beyond the limits of one system should be able to easily move their sonification data and tools onto a more powerful system. Thus, tools must be consistent, reliable, and portable across various computer platforms.

Similarly, tools should be capable of moving flexibly between real-time and nonreal-time sound production.

Flexibility: We need to develop synthesis controls that are specific and sophisticated enough to shape sounds in ways that take advantage of new findings from perceptual research on complex sounds and multimodal displays and that suit the data being sonified. In addition to flexibility of synthesis techniques, simple controls for altering the data-to-sound mappings or other aspects of the sonification design are also necessary. However, there should be simple “default” methods of sonification that allow novices to sonify their data quick and easily.

Integrability: Tools are needed that afford easy connections to visualization programs, spreadsheets, laboratory equipment, and so forth. Combined with the need for portability, this requirement suggests that we need a standardized software layer that is integrated with data input, sound synthesis, and mapping software and that facilitates the evaluation of displays from perceptual and human factors standpoints.

3.3 Sonification Application and Design

Sonification is evolving from a field of inquiry to a field of application. The question is no longer whether it works or even whether it is useful, but rather, how one designs a successful sonification. We begin with some examples of sonification that have proven successful in practical terms as well as in scientific terms. Next we look at the need for a method that can make it quicker and easier to produce a sonification that is likely to be useful. Finally we suggest research directions that will advance the development of sonification as a field of design practice. Throughout this section we suggest the basic features of perception and cognition that are relevant to sonification and stand to benefit from a sonification research agenda.

3.3.1 Some Successful Sonifications

Perhaps the most successful example of sonification is the Geiger-counter, which was invented by Hans Geiger in the early 1900s and is still in widespread use today. The Geiger counter is an instrument that “clicks” in response to invisible radiation levels, alerting one to danger that may go unnoticed with a visual display, as well as allowing a continual awareness of the degree of danger. Experiments show that people are better at monitoring radiation levels by audio than by visual display and that the audio display is also better than an audio+visual display (Tzelgov, Srebro, Hennik, & Street, 1987). Such applications capitalize on the listener’s ability to detect small changes in auditory events or the user’s need to have his eyes free for other tasks. A device similar in concept to the Geiger-counter, called the Pulse-oximeter, became standard equipment in medical operating theaters in the United States during the mid-1980s. The Pulse-

oximeter produces a tone that varies in pitch with the level of oxygen in a patient's blood, allowing the doctor to monitor this critically important information while visually concentrating on surgical procedures. The idea was extended to a six-parameter medical workstation by Fitch and Kramer. Medical students working with this workstation in a simulated operating room scenario were able to identify emergency situations more quickly with the audio display than with a visual display or a combined audio+visual display (Fitch & Kramer, 1994).

Sonification has also been successfully used in data analysis and exploration tasks, proving fruitful where visual techniques have not. During the Voyager 2 space mission there was a problem with the spacecraft as it began its traversal of the rings of Saturn. The controllers were unable to pinpoint the problem using visual displays, which just showed a lot of noise. When the data was played through a music synthesizer, a "machine gun" sound was heard during a critical period, leading to the discovery that the problem was caused by high-speed collisions with electromagnetically charged micrometeoroids (Kramer, 1994a, p. 35). Recently there was a very exciting moment when physicists Davis and Packard attributed an important discovery they have called the "Quantum Whistle" to their use of a sonification technique (Pereverzev et al., 1997). After months of unsuccessful study of visual oscilloscope traces for evidence of an oscillation predicted by quantum theory, Davis and Packard decided to listen to their experiment instead. What they heard was a faint whistling—the first evidence that these oscillations actually do occur. These cases illustrate the ability of the auditory system to extract underlying structure and temporal aspects of complex signals that are often important in scientific exploration and discovery.

Another promising area for sonification is sensory substitution for visually impaired users. There has been increased interest in augmenting haptic displays with sound for purposes of presenting graphical information. Schemes for auditory rendering of maps and diagrams embedded in text (Kennel, 1996; Gardner et al., 1996; Stevens & Edwards, 1997) are being developed, along with a number of approaches to rendering nontext Web-based content and geographical position to blind users. Meijer (1992) has developed means for scanning arbitrary visual images and presenting them in sound. In a more specialized scientific domain, Lunney and Morrison (1990) have developed an effective means of presenting infrared spectrometry data via several complementary mapping schemes for blind chemists and chemistry students. This is an example of "diagnostic" sonification specifically designed for a special needs population.

Educational applications are also promising. Studies show that most people can understand trends, clustering, correlations, and other simple statistical features of a data set just as well by listening to it as they could by reading a graph (Flowers, Buhman, & Turnage, 1996). There are indications that using sonification to present information to students in primary and secondary schools can provide a more engaging learning experience (Kramer, 1994b). Rhythm and music are used as a mnemonic device for teaching young students concepts such as the

alphabet and the number of days in each month. Similarly, it may be possible to harness the underlying components of this learning dynamic to assist students in grasping more sophisticated concepts such as common curves in calculus or distributions in statistics. Representing concepts and data through sound provides a means of capitalizing on strengths of individual learning styles, some of which may be more compatible with auditory representations than more traditional verbal and graphical representations. Adult education, training, and generalized information presentation may likewise benefit. The demands for information presentation in the scientific community are particularly acute.

The increasing number, size, and complexity of data-sets challenge existing visualization techniques. An example is the terabyte-sized data sets in seismology. Since seismic data is acoustic in nature, seismologists have often suggested listening to this type of data. Replaying the seismic recordings at audio rates makes it possible to overview 24 hours in just a few minutes. Early work in auditory presentation of seismic data showed that subjects could successfully discriminate between earthquakes and bomb blasts (Speeth, 1961). More recently, Hayward developed a number of sonification techniques for seismic analysis tasks and data types (Hayward 1994).

3.3.2 The Evolution of Design

By now it is clear that sonification works and can be very useful. What is not so clear is how to go about designing a successful sonification. At the first International Conference on Auditory Display in 1992, pioneering sonification researcher Sarah Bly drew attention to the lack of a theory of sonification as a “gaping hole impeding progress in the field” (Bly, 1994). To underline the point, she challenged experts to sonify a multidimensional data set for a classification task. There was significant disparity in the accuracy of judgments made with the three sonifications that were submitted.

An applied theory of sonification will make design and research more efficient. Optimally, such a theory will formulate a set of guidelines that will parallel guidelines that grow out of human computer interface and human factors research. As HCI research rests on applied perceptual studies, so too will a theory of sonification rest upon applied perceptual and cognitive research. High-level cognitive and perceptual issues such as auditory scene analysis and cross-modal interactions must inform design decisions. The implications and limitations of applying existing tools for creating music and conducting psychoacoustic or auditory research to sonification must be recognized, and validated design guidelines must be integrated into newly developed sonification tools. To create display designs appropriate for each specific application, the target environments must be carefully analyzed in terms of both human and technical standards, as well as the nature of the data and goals of the specific application.

Design guidelines that prove generic to various sonification display systems and tasks are gradually emerging. Theories of timbre perception inspire and influence many methods of sound synthesis and composition. It is these ‘perceptually informed’ approaches that form the foundation for meaningful correlation of auditory dimensions and display dimensions. For example, Grey's (1977) study of multidimensional scaling of timbre perception has been used by Barrass (1996) to create a three-dimensional sonification display. Design guidelines associated with other high-level cognitive issues such as the use of metaphor (Ballas, 1994; Kramer, 1994b; Walker & Kramer 1996) and semiotics (Bargar, 1994) are in the very early stages of development. As sonification theory emerges, it promises to prevent researchers from repeating the errors of prior efforts as methodical approaches to sonification design become available. Integrating all of these issues presents unique challenges. Nevertheless, there have already been a number of successful sonification applications, and the list continues to grow. These successes suggest that other successful applications can be developed and that design principles will emerge as the field matures.

3.3.4 Summary of Sonification and Design

Sonification has been successfully deployed in a broad range of application areas, such as the Geiger-counter and the Quantum Whistle. Despite the successes, however, we still do not really know how to design a sonification that we know in advance will work well for a specific task. Progress in sonification will require specific research directed at developing predictive design principles. Only in this way will the field advance from ad hoc experiments to a coherent field of design research and practice. We can draw on the existing literature in psychoacoustics, perceptual psychology, and the cognitive sciences. However, sonification also involves issues of representation, task dependency, and user-interface interaction. This research needs to be task centered and user centered and to integrate perceptual sciences. Among the suggested topics for sonification research are the roles of aesthetics, metaphor and affect in display design, applications of gestalt formation, the impact of context on identification of data structures, and audio interaction paradigms. Progress in sonification requires research by interdisciplinary teams with funding that is intended to advance the field of sonification directly, rather than relying on progress through a related but peripheral agenda.

Until we understand more about what makes sonification successful, the field will remain mired in ad hoc trial-and-error design. Predictive principles for sonification must be developed which are founded on research specifically on sonification issues. This kind of research requires multidisciplinary teams funded specifically to carry out studies of sonification.

4 Global Issues for Establishment of a Discipline of Sonification

Sonification research and practice has a scattered history. During the NSF workshop on sonification we identified recognition for the field, communication, and education as global issues that need special attention if we are to establish a connected, coherent discipline of sonification. In addition, lessons learned from the development of the field of visualization over the past 20 years may be very instructive for the development of the field of sonification. In the next four sections we discuss these issues and how they can be addressed by the NSF.

4.1 Recognition

Sonification research is typically carried out as part of the activities in some other program, for example, in a psychology department, a human-computer interaction lab, or an engineering school. Although this diversity of disciplinary backgrounds gives the field hybrid diversity and richness, sonification research is typically considered peripheral to these fields, making it difficult to gain support or recognition for sonification efforts. The topical organization of the administration of agencies such as NSF and university academic departments discourages funding work across disciplines and makes it difficult to effectively evaluate research proposals with interdisciplinary components. In addition, past experience of many sonification researchers confirms that university administrators and funding agencies often fail to appreciate the necessity of interdisciplinary funding. It is too easy for such research to “fall through the cracks” in the funding system, being seen as peripheral or inappropriate by any of its component disciplines. This situation, together with the “channelization” of publications, impacts peer review at all levels (including tenure decisions), and funding can be denied to research that is not at the center of a specific discipline.

The establishment of sonification as a credible and recognizable research field lays a strong supporting foundation for research proposals and grant applications in this area. But this is a chicken-and-egg type of problem: without support for research we cannot build the body of knowledge into a discipline, and without a developed discipline it is extremely difficult to obtain research support. A clear statement from the NSF that recognizes sonification as a scientific discipline would directly address this problem.

4.2 Communication and Community Formation

Technical communication within the field as well as broader communications with funding agencies, institutions, potential users, and the general public is important for the development of a discipline of sonification. Community formation, including conferences and workshops, Web sites, and other activities that encourage collaboration, is also essential to developing the field.

Currently, the main communications in the sonification community occur through ICAD (International Community for Auditory Display) conferences, ICAD conference proceedings, the ICAD Web site, and the ICAD e-mail list-server. The ICAD conferences were initiated at the Santa Fe Institute in 1992 as biannual events in the United States. Now conducted under its own nonprofit auspices, the International Conference on Auditory Display is an annual international venue for sonification and other auditory display research. With the maturation of the Web site and list-server, an international community of researchers has begun to form. Other communications channels that touch on sonification research are the ACM's SOUND list-server, SIG-GRAPH, SIG-CHI (Computer-Human Interface), and UIST (User Interface Software and Technology) conferences, and, in cooperation with ICAD, the Acoustical Society of America and European Acoustical Union.

To date, however, there is no dedicated journal of sonification, making it difficult to publish peer-reviewed articles about sonification. Publications must be submitted to journals specialized to other disciplines, such as electrical engineering, computer science, perception, computer music, virtual environments, or scientific or industrial application areas. Although the appearance of articles in these diverse journals indicates a wide interest in the topic, each article reaches only a fraction of the sonification community, and the reports often focus more on the application than on the sonification. A journal of sonification would provide an important venue for publication to support academic and researcher careers, allow articles to be specifically about sonification, make it easier to track developments, and further distinguish and unify the field. The contents of such a journal would include recent results, tips and techniques, tools, Web sites, reviews, conferences, employment opportunities, a bibliography, and many other important resources. The journal could be either paper or on-line in format. The on-line format has advantages for international distribution as well as providing the possibility to distribute audio material.

The establishment of a journal of sonification would require paid staff and significant funding. In the future, such a journal may be a worthwhile candidate for grant funding. Until such a journal can be established, sonification results can be published in special issues of existing journals, which offer an opportunity to expose segments of the general scientific community to sonification research. We recommend the continued financial support of ICAD efforts, including workshops, conferences, and Web activities.

4.3 Education

Students interested in sonification currently have to tailor an academic program from diverse units in music, engineering, computer science, physics, psychology, and the arts. The lack of cross-disciplinary interaction among departments and the lack of a clear idea of what classes to include are massive obstacles facing a student interested in sonification. Often, classes

outside of one's specialty are not approachable by the uninitiated. The problem can be addressed by the development of a curriculum for sonification. This curriculum would provide guidance for students who wish to research this area, would provide a framework for classes devoted to sonification, and would further define the discipline. Already some institutions and individuals are taking this direction. For example, the Australian Centre for Arts and Technology has a course in sonification, an auditory display course is being offered as part of the Human-Computer Interaction at the University of Glasgow, and a sonification course is offered at the University of Illinois, Urbana-Champaign. Further support of doctoral dissertations focused on sonification would provide a valuable stimulus to research.

4.4 Lessons Learned from Visualization

Important parallels exist between sonification and the more established field of scientific visualization. Thus, proponents of sonification stand to learn from the challenges faced in the early days of visualization, and the manner in which these challenges were addressed.

Visualization has been employed for centuries as a powerful way to present information. Charts and graphing techniques (Tufte, 1992) and maps and cartography (Tufte, 1990) were well-established ways to represent information long before our time. The real boost for visualization was the development of computer graphics (Foley, van Dam, Feiner, and Hughes, 1990). The progression of hardware and algorithms has provided users with increasingly powerful interactive 3-D graphics and corresponding breakthroughs in the representation of spatially indexed data. Meanwhile, sophisticated interaction methods have made it possible to explore dynamic data sets efficiently (Tufte, 1997).

Although researchers were accustomed to the use of graphs and other basic visualization techniques, there was still a long delay between the first data visualization experiments and broad acceptance of these scientific tools. Researchers, concentrated in their own application areas, were reluctant to take the time to use a new technique until it was efficient for them to do so. It was the gradual accumulation of success stories (see, for example, McCormick, DeFanti & Brown, 1987) that paved the way for significant research support, journals, and ongoing international conferences. Widely employed techniques, such as isocontours and isolines, evolved into a knowledge base common among visualization researchers. More complex techniques, such as glyphs, evolved more recently, building upon the existing research infrastructure and taxonomies. Such taxonomies do not currently exist in the field of sonification.

Data visualization did not evolve, however, without serious difficulties. In a sometimes ad hoc manner, scientific applications were driven by the data exploration needs of scientific disciplines. Although significant parallel efforts in the use and evaluation of visual information by artists and cognitive scientists have progressed, little integration of this information with the visualization community has occurred. Ideally, the sonification community can learn from the

visualization community's mistakes and embrace cross-disciplinary interactions from the outset, including those with cognitive scientists and artists. In order to develop sonification in the most effective manner, collaborations must involve scientists from disciplines that could benefit from sonification: psychophysicists, human factors researchers, and composers, as well as specialists in sonification design.

Sonification efforts must be carefully evaluated with appropriate user validation studies, taking into account application-specific goals as well as aesthetic considerations. The absence of such studies in the early days of visualization slowed its acceptance. Without this multidisciplinary approach, the field of sonification will mature slowly or not at all; instead, applications of sonification will be developed occasionally on an ad hoc basis, but no theoretical framework guiding effective sonification will result.

Like the general public, the research community is “visually biased,” in part because of the power of our visual systems and the associated long history of graphical presentation of data, and in part because of the current prevalence of sophisticated digital graphics and data visualization techniques. Even so, this challenge forms the basis for the eventual acceptance of sonification. Consider this statement by a sonification researcher:

“People spend up on stereo graphics supercomputers and CAVEs and yet say ‘sound won’t add anything.’ The only convincing argument is a working sonification. I’ve found that once they have experienced it, there is an almost immediate mindshift, and sonification becomes a natural part of the interface (just like the switch from silent movies).”

One of the goals of the present report is to inform the National Science Foundation and the general scientific and industrial communities of the existence and potential of data sonification. The single most important event in the creation and definition of the visualization field was the establishment of the McCormick, DeFanti, and Brown commission and the publication of its report in the July 1987 issue of *Computer Graphics*. That report created the necessary momentum for funding (both industrial and governmental) in the area of scientific visualization and helped launch visualization into the successful discipline it is today. While we believe that sonification is many years behind data visualization, we hope that the current white paper will serve sonification by leading to a report similar to that produced by McCormick, Defanti, and Brown (1987).

4.5 Summary of Global Issues

We have identified three global issues that need to be addressed to establish a discipline of sonification. The first is recognition of sonification as a valid area of research. This recognition will enable researchers to include sonification as a component in multidisciplinary proposals. The recognition of sonification by the NSF can play a major role in this validation. The second issue is communication within the sonification community. We propose support for

coordinated workshop and conference activity and a peer-reviewed journal of sonification. These would gather together the fragmented information that currently exists, would improve the coherence within the community, and would help foster academic careers that depend on publications. The final issue is the need to provide a curriculum for teaching sonification. This curriculum would provide guidance for students and teachers and would further define the discipline. Support for the major changes proposed here needs to come not only from individual researchers, but also from sources outside specific academic and professional areas. Agencies such as NSF that span diverse disciplines can play a key role in advancing young, interdisciplinary fields such as sonification. In addition, the development of the field of scientific visualization may provide both guidance and encouragement to researchers, policy-makers, and funding agencies involved in the present development of the field of sonification.

5 Proposed Research Agenda

Based upon this analysis of the field of sonification and its component disciplines, we recommend the following research agenda. For ease of discussion, these recommendations are divided according to the established disciplines of psychology (perception and cognition), computer science (sound synthesis and tool building), and the emerging discipline of sonification (design, applications, and theory). We also make overall recommendations regarding issues of education, communication, and research. Sample research questions are embedded within each section to lend specificity to the more general discussion.

5.1 General Funding Recommendations

Because sonification research is in a formative stage, diversity of ideas is crucial, and funding policies for sonification research should take into consideration the need to encourage innovation. Multiple grants of modest proportions, evaluated via a streamlined review process, will stimulate innovative research more effectively than will supporting a few larger grants with the same total funds. Funding must be adequate, however, to support multidisciplinary teams and to support research project development and evaluation.

Funding bodies should also recognize that sonification research could contribute to applications in a wide variety of disciplines. In response to the information revolution, some funding agencies are beginning to develop new, broadly conceived and interdisciplinary initiatives for research and development to enhance the exploration and communication of data via new technology. The National Science Foundation's *Knowledge & Distributed Intelligence Program* is one such example. We strongly encourage such cross-disciplinary programs to support sonification research.

To accomplish these goals, the development of alternatives to traditional basic research funding sources might be valuable. Such alternatives could include coordination between basic and applied funding sources, or interdisciplinary initiatives in individual agencies. Mechanisms could be developed within funding agencies for piggy-backing sonification research onto larger projects within specific disciplines that could benefit from effective sonification. For example, if a large project in seismology has already passed review, there should be a way of obtaining incremental funding for investigating the benefits of sonifying the seismic data via a separate application and review procedure.

5.2 Perception and Cognition

Ongoing research in auditory perception that is of particular relevance to sonification includes (1) dynamic sound perception, (2) auditory scene analysis, (3) auditory memory, and (4) the role of attention in extracting information from sound. Other pertinent research includes investigation of the normal variation in perceptual and cognitive abilities and strategies in the human population, differences in cognitive representations of auditory displays for sighted and blind individuals, and the role of learning and familiarity in visual display efficacy. We recommend coordination between research in perception and in sonification, with sonification research dollars being focused on those problems most closely associated with display design and use.

For example, a crucial issue in sonification is how distinctions among basic acoustic properties of sound events, such as their energy envelope and spectral content, affect selective attention to event streams. Analysis of multivariate data may sometimes require focused attention on individual variables, and at other times require divided attention to allow the listener to detect similarities and contrasts in trends of different variables or streams. A research question that is already important in auditory scene analysis is how tone frequencies influence perceptual grouping and attention (Bregman, 1990), but other more complex acoustic variables have received little attention. Study of these issues is necessary for the formulation of design guidelines for constructing efficient multivariate data displays using sounds. Much of this effort could take place within ongoing research projects in auditory scene analysis.

Second, substantial questions have been identified in the area of multimodal perception. Specifically, effective sonification depends on the ability to evaluate displays that utilize different modalities and to identify the best modality for presenting a particular type of information. It is clear from existing research that our senses differ in their underlying abilities, but further research is necessary to identify what data features are perceptually most salient to each sense and to determine how to use this knowledge in designing effective displays. Multimodal interactions (e.g., between visual and auditory displays) are poorly understood, yet critically affect most sonification applications. When does redundant presentation (that is,

presenting information in more than one modality) improve the ability to extract data (i.e., cross-modal synergy)? When does information presented in one modality interfere with the perception of information in another modality (i.e., cross-modal interference)? How can the total amount of information perceived across all modalities be maximized? Only by careful investigation of these issues can we optimize displays for the type of information conveyed. Funding perception laboratories to study sonification-specific issues in multimodal perception would be one way to encourage such progress.

Sample Research Question #2. To what extent (and under what circumstances) can the presentation of redundant auditory and visual data representation increase display efficiency? A large body of memory research suggests that use of multiple encoding strategies leads to more durable and accurate memory. While claims for the efficacy of multimedia education often point to such research, we lack a systematic evaluation of how redundant sound and sight might provide synergy in multimodal presentations. One potentially promising domain for use of redundant presentation is in teaching about data patterns (e.g., showing students features in economic trends and key relationships between indices). Will memory for key features and patterns be better if they are both heard and seen? Or, in contrast, will multimodal presentation result in unwanted interference and cross-talk, hampering understanding?

5.3 Developing Sonification Tools

The most significant advances in software that will form the basis of sonification research and applications are being made by centers devoted to research in computer music and acoustics. Research in sound synthesis, signal analysis, and the perception of complex sound is generally associated with departments of music, electrical engineering, and computer science, or occasionally with psychology departments. It is essential that sonification research leverage these efforts. Therefore, our general recommendation for tool development includes supporting efforts to adapt sound synthesis software to the needs of sonification research, in particular as follows:

1. Control: provide efficient, effective, accessible, parametric controls of the sound that constitutes the display medium.
2. Mapping: allow the design of new sonifications "on the fly" by giving the user flexible, intuitive control over which data dimensions control which sound parameter;
3. Integrability: facilitate data importation from a variety of formats to allow data from many different disciplines to be sonified.
4. Synchrony: allow easy integration with other display systems such as existing visual monitors, virtual reality systems, or assistive technology devices.
5. Experimentation: integrate a perceptual testing framework with the overall sound synthesis and mapping functions.

The software emerging from the music and acoustics research centers is generally powerful and complex, but intimidating to the novice. While powerful tools are clearly necessary, equally important are tools that are simple to use and that encourage casual exploration of sonification. Growth of the field requires an easy-to-use system that enables scientists and students to sonify their own data sets in an interactive manner.

Whether tools are simple or sophisticated, consistency, reliability, and portability across platforms must be maintained. Funds should be allocated to develop sonification software guidelines, including some basic user-interface guidelines, standard means of importing the data to be sonified, and a common terminology for the functions of the software. Wherever possible, high-level tools should be developed in conjunction with, or with input from, applications experts (e.g., scientists, blind users, process control specialists) to ensure that the tools are relevant and comprehensible to the widest possible user population.

Finally, one of the major needs in the sonification community is identification of the acoustic parameters tied to perceptual attributes of complex sounds, such that the sound attributes have natural, intuitive, and veridical relationships to the data being represented. Both recent psychoacoustic research (investigating the perception of complex auditory phenomena) and research in acoustics and computer science (developing high-level models for synthesizing complex, real-world sounds) have direct application to sonification. Further development of synthesis techniques and algorithms that provide high-level handles corresponding to high-level descriptions of sounds should be encouraged. Examples of such correspondences include high-level handles analogous to naturally occurring changes in real-world sounds (Gaver, 1994) or simultaneous control of multiple acoustic variables to produce a single, complex perceptual variable (Kramer, 1996).

Two final criteria for tool development funding must be mentioned. Whenever possible, tools should enable real-time interaction with the data. Interactive data exploration has been found to be a significant feature in determining the usefulness of data visualization; we subjectively find the same dynamic at work in data sonification. Also, one important criterion for all tools-related funding should be effective distribution and support. Code must be easy to modify and maintain, and explicit plans for wide distribution should be an integral part of any proposal in this area.

Sample Research Proposal #3. Building upon a pre-existing real-time sound synthesis package, researchers will create a “sonification shell”: a generalized and easy-to-use framework for sonification research with facilities for importing data, choosing mappings between variables and sounds interactively, navigating through the data set, synchronizing the sonic output with other display media, and performing standard psychophysical experiments to evaluate the resultant sonification system. The system will be easy to get started with, providing a graphical user interface and “canned displays” with standard hooks for different types of variables (e.g., temporal vs. spatial variables, periodic vs. aperiodic), allowing novice researchers to begin immediately sonifying their data. However, it will not restrict more advanced users from tinkering “under the hood” and developing their own synthesis algorithms or complex mappings or nonstandard experimental protocols. The system will be developed in a platform-independent, object-oriented language (such as Java) for portability and easy maintenance and modification and will be distributed free via the Internet.

5.4 Sonification Applications, Design and Theory

Sonification will gain significant momentum once several specific applications become widely used. However, until there are intuitive, efficacious applications, skeptics will adhere to current display solutions. We suggest, then, that it is essential to identify and fund a small number of promising sonification applications. A well-publicized request for proposals could generate such applications. In addition, a design competition would encourage sonification researchers to promote their successful, but perhaps unpublicized efforts. Results of such projects could be documented and made available through a cross-referenced archive of successful (and perhaps not-so successful) designs. Issues to consider in evaluating such flagship projects would include the following:

1. Is the sonification an effective alternative or useful complement to other display approaches?
2. Does the data in this application lend itself to effective sonification?
3. Are there acoustic cues that can reliably convey this information?

Problem-driven projects that are carefully evaluated should be given a high priority. However, exhaustive testing of all data, acoustic parameters, and perceptual parameters is impractical for most multidimensional tasks. A reasonable alternative would be to evaluate performance and preference with a judiciously chosen subset of parameter values.

Effective sonification design will require a theoretical foundation. Carefully constructed theory will form the basis for efficient research efforts, supporting progressive improvements and avoiding duplication of efforts. Theory-building for sonification design will necessitate the input of professionals from various fields. Thus, support for interdisciplinary teams is critical and will encourage a task-oriented, human-centered approach to design. The theoretical framework should address questions such as the following:

1. Is there a psychologically-based or application-supported natural taxonomy of sonification techniques?
2. What types of data or tasks lend themselves naturally to effective sonification?
3. Which acoustic cues and data mappings are intuitive and facilitate the presentation of complex, multidimensional displays?
4. What factors limit how well information can be extracted from a sonification?

Research should be funded to investigate overarching questions associated with other theoretical considerations. Task-dependent and user-centered approaches to sonification design should be supported. Research should be conducted to explore predictably effective data to sound parameter mappings, with such research being informed by studies in timbre perception and associated synthesis methods. Guidelines should be developed that make effective displays easier to design and fit naturally with how people work and learn. Effective uses of metaphor, affect, aesthetics, semiotics, and music theory will be an integral part of these guidelines. Likewise, basic theory will help researchers avoid predictable design problems such as display-induced anomalies. Establishing guidelines that can be productively applied to a broad range of sonification tasks and display systems will optimize the development of a methodical approach to sonification design.

It is likely that theories will initially evolve from practice. However, as the field matures, funding should be provided for investigations to discern general principles and techniques that are effective, and under what circumstances those principles apply. Such a theoretical foundation will help to drive research, allowing a more coordinated and comprehensive account of both the benefits and limitations of sonification.

5.5 Support for the Field at Large

Sonification research must be supported by communication and community building efforts and by efforts to stimulate research. Groups such as ICAD, the International Community for Auditory Display, and existing research centers interested in sonification, should receive modest funding to hold small workshops, develop Web-based community building projects, and, above all, gather and disseminate sonification research results. One resource that is currently needed by the sonification community is an up-to-date repository of sonification examples, suggested technical or design standards, approaches to evaluation and display metrics, and other reference material. The existing ICAD Web site provides a model for such an effort, but some funding would be required to effectively cross-reference and collate, and then maintain, all of the existing research materials.

6 General Conclusion

A real need for more effective means of making sense of data is well documented. Most of the infrastructure required for auditory representations of data already exists, as a result of advances in computer technology and auditory research. A coordinated research effort, supported by moderate funding at the national level, is necessary if sonification research is to take advantage of this fertile environment. The resultant advances in both basic research and technology development will contribute to both scientific and commercial applications, which will then feed back into the development of the field. National Science Foundation funding and leadership can help to accelerate this process.

References

- Ballas, J. (1994) Delivery of information through sound, in *Auditory Display: Sonification, Audification and Auditory Interfaces*, G. Kramer, ed. 79-94. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. XVIII. Reading, MA: Addison Wesley.
- Bargar, R. (1994) Pattern and reference in auditory display, in *Auditory Display: Sonification, Audification and Auditory Interfaces*, G. Kramer, ed. 151-165. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. XVIII. Reading, MA: Addison Wesley.
- Barrass, S. (1996). *Sculpting a Sound Space with Information Properties*, Organised Sound, 1, 2, Cambridge University Press.
- Blauert, J. (1997). *Spatial hearing: The psychophysics of human sound localization*. Cambridge, MA: MIT Press.
- Bly, S. (1994) Multivariate data mappings. In *Auditory display Sonification, audification and auditory interfaces*, G. Kramer, ed. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. XVIII, 405-416. Reading, MA: Addison Wesley.
- Bregman, A. S. (1990). *Auditory scene analysis*. Cambridge, MA: MIT Press.
- Brungart, D. S. (1998). Control of perceived distance in virtual audio displays. *Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 1101-1104.
- Carlile, S. (1996). *Virtual auditory space: Generation and applications*. New York: R. G. Landes.
- Chowning (1973). The synthesis of complex audio spectra by means of frequency modulation. *Journal of the Audio Engineering Society*, 21(7), 526-534. Reprinted in C. Roads & J. Strawn (eds.), *Foundations of computer music*. Cambridge, MA, MIT Press.
- Cook, P. R. (1995). "A Hierarchical System for Controlling Synthesis by Physical Modeling," International Computer Music Conference, Banff, 1995.
- Cook, P. R. (1997). Physically inspired sonic modeling (PhISM): Synthesis of percussive sounds. *Computer Music Journal*, 21(3). 38-49.
- Duda, R. O., & Martens, W. L. (1998). Range dependence of the response of a spherical head model. *Journal of the Acoustical Society of America*, 104, 3048-3058.
- Earl, C. & Levanthal, J. A. (1999) Survey of Windows screen reader users: Recent improvements in accessibility. *Journal of Visual Impairment and Blindness: American Foundation for the Blind*, v93 #3 (in press).
- Fitch, T., & Kramer, G. (1994). Sonifying the body electric: Superiority of an auditory over a visual display in a complex, multivariate system. In G. Kramer (ed.), *Auditory display: Sonification, audification and auditory interfaces. Proceedings of the First International Conference on Auditory Display (ICAD) 1992*, 307-326. Reading, MA: Addison-Wesley.
- Flowers, J. H., Buhman, D. C., & Turnage, K. D. (1996). Cross-modal equivalence of visual and auditory scatterplots for exploring bivariate data samples. *Human Factors*, 39(3), 341-351.
- Foley, J., van Dam, A., Feiner, and S., Hughes, J., (1990) *Computer graphics principles and practice*. 2nd ed. Addison Wesley, New York.
- Freed, A., Rodet, X., & Depalle, P. (1993). Synthesis and control of hundreds of sinusoidal partials on a desktop computer without custom hardware. In *Proceedings of the International Conference on Electronic Engineering Times*, 1024-1030. Santa Clara, CA: DSP Associates.
- Gardner, J., Lundquist, R., Sahyun, S. (1996) Triangle: A Practical Application of Non-Speech Audio for Imparting Information Proceedings of ICAD 96. International Community for Auditory Display (Web document, available at www.santafe.edu/~icad).
- Gaver, W. W. (1994). Using and creating auditory icons. In G. Kramer (ed.), *Auditory display: Sonification, audification and auditory interfaces. Proceedings of the First International Conference on Auditory Display (ICAD) 1992* (pp. 417-446). Reading, MA: Addison-Wesley.
- Gaver, W. W., Smith, R. B., & O'Shea, T. (1991). Effective sounds in complex systems: the ARKola simulation. In *Proceedings of CHI '91*. New York: ACM.

- Grey, J. M. (1977) Multidimensional perceptual scaling of musical timbres, *Journal of the Acoustical Society of America* 63, 1493-1500.
- Grey, J. M., and Moorer, J. A. (1978) Perceptual evaluations of synthesized musical instrument tones, *Journal of the Acoustical Society of America* 62(2). 454-462.
- Gulick, W. L., Gescheider, G. A., & Frisina, R. D. (1989). *Hearing: Physiological acoustics, neural coding, and psychoacoustics*. Oxford University Press.
- Hajda, et al (1999) Hajda, J. M., Kendall, R. A., Carterette, E. C., & Harshberger, M. L. Methodological issues in timbre research. In *I. Deliège & J. Sloboda (Eds.), Perception and Cognition of Music*. London: Psychology Press (in press).
- Handel, S. (1989). *Listening: An introduction to the perception of auditory events*. Cambridge, MA: MIT Press.
- Hartmann, W. M. (1997). *Sounds, signals, and sensation: Modern acoustics and signal processing*. New York.: Springer Verlag.
- Hayward, C. (1994). Listening to the earth sing. In G. Kramer (ed.), *Auditory display: Sonification, audification and auditory interfaces. Proceedings of the First International Conference on Auditory Display (ICAD)*, 369-404. Reading, MA: Addison-Wesley.
- Howard, J. H., & Ballas, J. A. (1982). Acquisition of acoustic pattern categories by exemplar observation. *Organizational Behavior and Human Decision Processes* 30(2), 157-173.
- Jenison, R. L., Neelon, M. F., Reale, R. A., & Brugge, J. F. (1998). Synthesis of virtual motion in 3D auditory space. *Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 1096-1100.
- Karplus, K., & Strong, A. (1983). Digital synthesis of plucked-string and drum timbres. *Computer Music Journal* 7(2), 43-55.
- Kennel, A. (1996) AudioGraf: A diagram reader for blind people, *Proceedings of ASSETS'96 Second Annual ACM Conference on Assistive Technologies*, 51-56. Vancouver: ACM Press.
- Kramer, G., ed. (1994a). *Auditory display: Sonification, audification, and auditory interfaces. Proceedings of the First International Conference on Auditory Display (ICAD) 1992*. Reading, MA: Addison-Wesley.
- Kramer, G. (1994b) Some organizing principles for representing data with sound. In G. Kramer (ed.) *Auditory display: Sonification, audification and auditory interfaces Proceedings of the First International Conference on Auditory Display (ICAD)*. 185-221. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. XVIII, 85-221. Reading, MA: Addison Wesley.
- Kramer, G. (1994c) Sound and communication in virtual reality. In F. Biocca (ed.), *Communication in the age of virtual reality*. 259-276. New York: Lawrence Earlbaum.
- Kramer, G. (1996) Mapping a single data stream to multiple auditory variables: A subjective approach to creating a compelling design. In *Proceedings of ICAD 96*. Web publication by the International Community for Auditory Display (Web document, available at www.santafe.edu/~icad).
- Krimphoff, J. (1993) Analyse acoustique et perception du timbre. Thesis, Université de Maine.
- Lunney, D., & Morrison, R. (1990). High technology laboratory aids for visually handicapped chemistry students. *Journal of Chemistry Education* 58, 228.
- Martens, W. L. (1985). Palette: An environment for developing an individualized set of psychophysically scaled timbres. *ICMC*, 85. 355-365.
- Mathews, M. (1969). *The technology of computer music*. Cambridge, MA: The MIT Press.
- Meijer, P.B.L. (1992) An Experimental System for Auditory Image Representations, *IEEE Transactions on Biomedical Engineering*, Vol. 39(2). 112-121.
- McAdams, S., & Bigand, E. (eds.) (1993). *Thinking in sound: The cognitive psychology of human audition*. Oxford: Clarendon Press.

- McCormick, B. H., DeFanti, T.A., & Brown, M. D., (1987) Visualization in scientific computing. *Computer Graphics* 21(6), special issue.
- McGurk, H., & McDonald, T. (1976). Hearing lips and seeing voices. *Nature*, 264. 746-748.
- Miner, N. E. (1998). *Creating wavelet-based models for real-time synthesis of perceptually convincing environmental sounds*. Unpublished doctoral dissertation, University of New Mexico.
- Moore, B. C. J. (ed.) (1995). *Handbook of perception and cognition: Vol. 6. Hearing*. New York: Academic Press.
- Moore, B. C. J. (1997). *An introduction to the psychology of hearing*. 4th ed. Orlando, FL: Academic Press.
- Neuhoff, J. G. (1998). A perceptual bias for rising tones. *Nature* 395(6698), 123-124.
- Pereverzev, S. V., Loshak, A., Backhaus, S., Davis, J. C., and Packard, R. E. (1997) Quantum oscillations between two weakly coupled reservoirs of superfluid ^3He , *Nature* 388, 449-451.
- Pierce, J. R. (1992). *The science of musical sound* (2nd ed.). New York: W. H. Freeman and Company.
- Pierce, J. R. & Duyne, S. A. V. (1997). A passive non-linear digital filter design which facilitates physics-based sound synthesis of highly nonlinear musical instruments. *Journal of the Acoustical Society of America* , 101(2), 1120-1126.
- Radeau, M. (1992). Cognitive impenetrability in auditory-visual interaction. In *Analytic Approaches to Human Cognition*, 41-55. Amsterdam: Elsevier Science Publishers B.V.,
- Saldaña, H. M., and Rosenblum, L. D. (1993). Visual influences on auditory pluck and bow judgments. *Perception and Psychophysics* 54 (3), 406-416.
- Scaletti, C., & Craig, A. B.. (1990). Using sound to extract meaning from complex data. In E. J. Farrel (ed.), *Extracting meaning from complex data: Processing, display, interaction* 1259, 147-153. SPIE.
- Serra, X. (1989). *A system for sound analysis/transformation/synthesis based on a deterministic plus stochastic decomposition*. Unpublished doctoral dissertation, Stanford University.
- Shinn-Cunningham, B. G. (1998). Applications of virtual auditory displays. *Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 1105-1108.
- Smith, J. O. (1996). Physical modeling synthesis update. *Computer Music Journal* 20(2).44-56
- Speeth, S. D., (1961) Seismometer sound. *Journal of Acoustic Society of America* 33(7) s. 909-916.
- Stevens, R. D., Edwards, A. D. N. and Harling, P. A. (1997). Access to mathematics for visually disabled students through multi-modal interaction. *Human-Computer Interaction* (Special issue on Multimodal Interfaces), 12 (1&2) pp.47-92.
- Takala, T., and Hahn, J., (1993) Sound rendering. *Computer Graphics Proc. SIGGRAPH'92*, 211-220.
- Thurlow, W. R., & Jack, C. E. (1973). Certain determinants of the "ventriloquism effect." *Perceptual and Motor Skills* 36(3, Pt. 2), 1171-1184.
- Tufte, E. (1990). *Envisioning information*. Graphics Press, May 1990.
- Tufte, E. (1992). *The visual display of quantitative information*. Graphics Press, Reprint.
- Tufte, E. (1997). *Visual explanations: Images and quantities, evidence and narrative*. Graphics Press.
- Tzelgov, J., Srebro, R., Henik, A., & Kushelevsky, A. (1987). Radiation detection by ear and by eye. *Human Factors* 29(1), 87-98.
- Van Den Doel, K., and Pai, D. K. (1996). *Synthesis of Shape Dependent Sounds with Physical Modeling*. *Proceedings of the International Conference on Auditory Displays*, Palo Alto, CA (Web document, available at www.santafe.edu/~icad).
- Walker, B., and Kramer, G. (1996) Mappings and metaphors in auditory displays: an experimental assessment, *Proceedings of ICAD 96*. International Community for Auditory Display (Web document, available at www.santafe.edu/~icad).

- Wenzel, E. M. (1992). Localization in virtual acoustic displays. *Presence, 1*, 80-107.
- Wenzel, E. M., Arruda, M., Kistler, D. J., & Wightman, F. L. (1993). Localization using nonindividualized head-related transfer functions. *Journal of the Acoustical Society of America* 94, 111-123.
- Wessel, D.L. (1979) Timbre space as a musical control structure. *Computer Music Journal* 3(2), 45-72.
- Wickens, C. D. (1984). Processing resources in attention. In R. Parasuraman & R. Davies (eds.), *Varieties of attention*, 63-101. New York: Academic Press.
- Wickens, C. D., Gordon, S. E., & Liu, L. (1998). *An introduction to human factors engineering*. New York: Addison Wesley Longman.