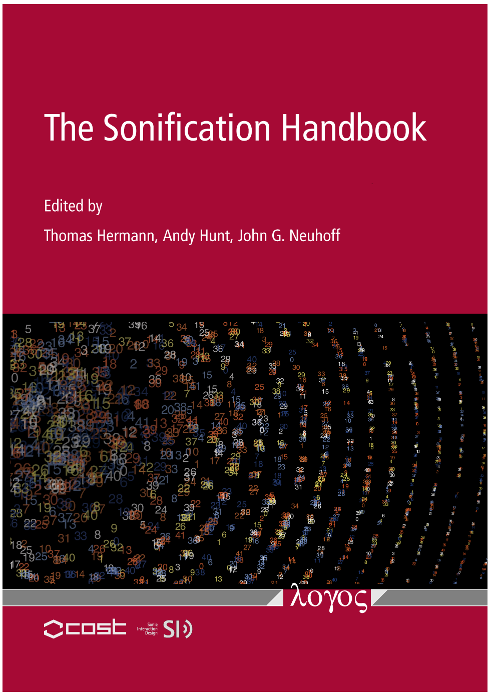

The Sonification Handbook

Edited by
Thomas Hermann, Andy Hunt, John G. Neuhoff

Logos Publishing House, Berlin, Germany
ISBN 978-3-8325-2819-5
2011, 586 pages

Online: http://sonification.de/handbook
Order: http://www.logos-verlag.com

Reference:
Hermann, T., Hunt, A., Neuhoff, J. G., editors (2011). *The Sonification Handbook*. Logos Publishing House, Berlin, Germany.

Chapter 1

Introduction

Thomas Hermann, Andy Hunt, John G. Neuhoff

The chapter provides a brief introduction to Sonification and Auditory Display, and suggests a vision on where sonification could be in 50 years. Finally the chapter explains the organization of the material in this volume and provides suggestions how to use the book and the media files.

Reference:
Hermann, T., Hunt, A., Neuhoff, J. G. (2011). Introduction. In Hermann, T., Hunt, A., Neuhoff, J. G., editors, The Sonification Handbook, chapter 1, pages 1–6. Logos Publishing House, Berlin, Germany.

Media examples: http://sonification.de/handbook/chapters/chapter1

Chapter 1

Introduction

Thomas Hermann, Andy Hunt, John G. Neuhoff

1.1 Auditory Display and Sonification

Imagine listening to changes in global temperature over the last thousand years. What does a brain wave sound like? How can sound be used to facilitate the performance of a pilot in the cockpit? These questions and many more are the domain of Auditory Display and Sonification. Auditory Display researchers examine how the human auditory system can be used as the primary interface channel for communicating and transmitting information. The goal of Auditory Display is to enable a better understanding, or an appreciation, of changes and structures in the data that underlie the display. *Auditory Display* encompasses all aspects of a human-machine interaction system, including the setup, speakers or headphones, modes of interaction with the display system, and any technical solution for the gathering, processing, and computing necessary to obtain sound in response to the data. In contrast, *Sonification* is a core component of an auditory display: the technique of rendering sound in response to data and interactions.

Different from speech interfaces and music or sound art, Auditory Displays have gained increasing attention in recent years and are becoming a standard technique on par with visualization for presenting data in a variety of contexts. International research efforts to understand all aspects of Auditory Display began with the foundation of the International Community for Auditory Display (ICAD) in 1992. It is fascinating to see how Sonification techniques and Auditory Displays have evolved in the relatively few years since the time of their definition, and the pace of development in 2011 continues to grow.

Auditory Displays and Sonification are currently used in a wide variety of fields. Applications range from topics such as chaos theory, bio-medicine, and interfaces for visually disabled people, to data mining, seismology, desktop computer interaction, and mobile devices, to name just a few. Equally varied is the list of research disciplines that are required to comprehend and carry out successful sonification: Physics, Acoustics, Psychoacoustics, Perceptual Research, Sound Engineering, Computer Science are certainly core disciplines that contribute to the research process. Yet Psychology, Musicology, Cognitive Science, Linguistics, Pedagogies, Social Sciences and Philosophy are also needed for a fully faceted view of the description, technical implementation, use, training, understanding, acceptance, evaluation and ergonomics of Auditory Displays and Sonification in particular. Figure 1.1 depicts an interdisciplinarity map for the research field.

It is clear that in such an interdisciplinary field, too narrow a focus on any of the above isolated disciplines could quickly lead to "seeing the trees instead of understanding the forest". As with all interdisciplinary research efforts, there are significant hurdles to interdisciplinary research in Auditory Display and Sonification. Difficulties range from differences in theoretical orientations among disciplines to even the very words we use to describe our work. Interdisciplinary dialogue is crucial to the advancement of Auditory Display and Sonification. However, the field faces the challenge of developing and using a common language in order to integrate many divergent "disciplinary" ways of talking, thinking and tackling problems. On the other hand this obstacle often offers great potential for discovery because these divergent ways of thinking and talking can trigger creative potential and new ideas.

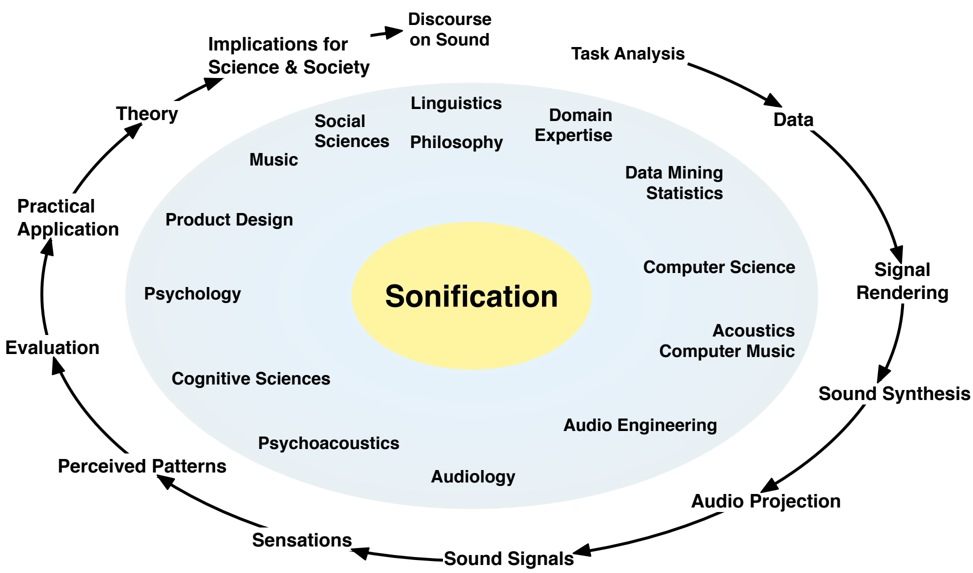

Figure 1.1: The interdisciplinary circle of sonification and auditory display: the outer perimeter depicts the transformations of information during the use cycle, the inner circle lists associated scientific disciplines. This diagram is surely incomplete and merely illustrates the enormous interdisciplinarity of the field.

1.2 The Potential of Sonification and Auditory Display

The motivation to use sound to understand the world (or some data under analysis) comes from many different perspectives. First and foremost, humans are equipped with a complex and powerful listening system. The act of identifying sound sources, spoken words, and melodies, even under noisy conditions, is a supreme pattern recognition task that most modern computers are incapable of reproducing. The fact that it appears to work so effortlessly is perhaps the main reason that we are not aware of the incredible performance that our auditory system demonstrates every moment of the day, even when we are asleep! Thus, the benefits of using the auditory system as a primary interface for data transmission are derived from its complexity, power, and flexibility.

We are, for instance, able to interpret sounds using multiple layers of understanding. For example, from spoken words we extract the word meaning, but also the emotional/health state of the speaker, and their gender, etc. We can also perceive and identify “auditory objects” within a particular auditory scene. For example, in a concert hall we can hear a symphony orchestra as a whole. We can also tune in our focus and attend to individual musical instruments or even the couple who is whispering in the next row. The ability to selectively attend to simultaneously sounding “auditory objects” is an ability that is not yet completely understood. Nonetheless it provides fertile ground for use by designers of auditory displays. Another fascinating feature is the ability to learn and to improve discrimination of auditory stimuli. For example, an untrained listener may notice that “something is wrong” with their car engine, just from its sound, whereas a professional car mechanic can draw quite precise information about the detailed error source from the same sound cue. The physician’s stethoscope is a similarly convincing example. Expertise in a particular domain or context can dramatically affect how meaning is constructed from sound. This suggest that – given some opportunity to train, and some standardized and informative techniques to hear data – our brain has the potential to come up with novel and helpful characterizations of the data.

Nowadays we have access to enough computing power to generate and modify sonifications in real-time, and this flexibility may appear, at first glance, to be a strong argument for rapid development of the research field of sonification. However, this flexibility to change an auditory display often and quickly can sometimes be counter-productive in the light of the human listening system’s need of time to adapt and become familiar with an auditory display. In the real world, physical laws grant us universality of sound rendering, so that listeners can adapt to real-world sounds. Likewise, some stability in the way that data are sonified may be necessary to ensure that users can become familiar with the display and learn to interpret it correctly.

Sonification sets a clear focus on the use of sound to convey information, something which has been quite neglected in the brief history of computer interfaces. Looking to the future, however, it is not only sound that we should be concerned with. When we consider how information can be understood and interpreted by humans, sound is but one single modality amongst our wealth of perceptual capabilities. Visual, auditory, and tactile information channels deliver complementary information, often tightly coupled to our own actions. In consequence we envision, as attractive roadmap for future interfaces, a better balanced use of all the available modalities in order to make sense of data. Such a generalized discipline may be coined *Perceptualization*.

Sonification in 50 years – A vision

Where might sonification be 50 years from now? Given the current pace of development we might expect that sonification will be a standard method for data display and analysis. We envision established and standardized sonification techniques, optimized for certain analysis tasks, being available as naturally as today's mouse and keyboard interface. We expect sound in human computer interfaces to be much better designed, much more informative, and much better connected to human action than today. Perhaps sonification will play the role of enhancing the appreciation and understanding of the data in a way that is so subtle and intuitive that its very existence will not be specifically appreciated yet it will be clearly missed if absent (rather like the best film music, which enhances the emotion and depth of characterization in a movie without being noticed). There is a long way to go towards such a future, and we hope that this book may be informative, acting as an inspiration to identify where, how and when sound could be better used in everyday life.

1.3 Structure of the book

The book is organized into four parts which bracket chapters together under a larger idea. Part I introduces the fundamentals of sonification, sound and perception. This serves as a presentation of theoretical foundations in chapter 2 and basic material from the different scientific disciplines involved, such as psychoacoustics (chapter 3), perception research (chapter 4), psychology and evaluation (chapter 6) and design (chapter 7), all concerned with Auditory Display, and puts together basic concepts that are important for understanding, designing and evaluating Auditory Display systems. A chapter on Sonic Interaction Design (chapter 5) broadens the scope to relate auditory display to the more general use of sounds in artifacts, ranging from interactive art and music to product sound design.

Part II moves towards the procedural aspects of sonification technology. Sonification, being a scientific approach to representing data using sound, demands clearly defined techniques, e.g., in the form of algorithms. The representation of data and statistical aspects of data are discussed in chapter 8. Since sonifications are usually rendered in computer programs, this part addresses the issues of how sound is represented, generated or synthesized (chapter 9), and what computer languages and programming systems are suitable as laboratory methods for defining and implementing sonifications (chapter 10). The chapter includes also a brief introduction to operator-based sonification and sonification variables, a formalism that serves a precise description of methods and algorithms. Furthermore, interaction plays an important role in the control and exploration of data using sound, which is addressed in chapter 11.

The different Sonification Techniques are presented in Part III. Audification, Auditory Icons, Earcons, Parameter Mapping Sonification and Model-Based Sonification represent conceptually different approaches to how data is related to the resulting sonification, and each of these is examined in detail.

Audification (chapter 12) is the oldest technique for rendering sound from data from areas such as seismology or electrocardiograms, which produce time-ordered sequential data streams. Conceptually, canonically ordered data values are used directly to define the samples of a digital audio signal. This resembles a gramophone where the data values actually determine the structure of the trace. However, such techniques cannot be used when the data sets are arbitrarily large or small, or which do not possess a suitable ordering criterion.

Earcons (chapter 14) communicate messages in sound by the systematic variation of simple sonic 'atoms'. Their underlying structure, mechanism and philosophy is quite different from the approach of Auditory Icons (chapter 13), where acoustic symbols are used to trigger associations from the acoustic 'sign' (the sonification) to that which is 'signified'. Semiotics is here one of the conceptual roots of this display technique. Both of these techniques, however, are more concerned with creating acoustic communication for discrete messages or events, and are not suited for continuous large data streams.

Parameter Mapping Sonification (chapter 15) is widely used and is perhaps the most established technique for sonifying such data. Conceptually, acoustic attributes of events are obtained by a 'mapping' from data attribute values. The rendering and playback of all data items yields the sonification. Parameter Mapping Sonifications were so ubiquitous during the last decade that many researchers frequently referred to them as 'sonification' when they actually meant this specific technique.

A more recent technique for sonification is Model-Based Sonification (chapter 16), where the data are turned into dynamic models (or processes) rather than directly into sound. It remains for the user to excite these models in order to explore data structures via the acoustic feedback, thus putting interaction into a particular focus.

Each of these techniques has its favored application domain, specific theory and logic of implementation, interaction, and use. Each obtains its justification by the heterogeneity of problems and tasks that can be solved with them. One may argue that the borders are dilute – we can for instance interpret audifications as a sort of parameter mapping – yet even if this is possible, it is a very special case, and such an interpretation fails to emphasize the peculiarities of the specific technique. None of the techniques is superior per se, and in many application fields, actually a mix of sonification techniques, sometimes called hybrid sonification, needs to be used in cooperation to solve an Auditory Display problem. Development of all of the techniques relies on the interdisciplinary research discussed above. These 'basis vectors' of techniques span a sonification space, and may be useful as mindset to discover orthogonal conceptual approaches that complement the space of possible sonification types.

Currently there is no single coherent theory of sonification, which clearly explains all sonification types under a unified framework. It is unclear whether this is still a drawback, or perhaps a positive property, since all techniques thus occupy such different locations on the landscape of possible sonification techniques. The highly dynamic evolution of the research field of auditory display may even lead to novel and conceptually complementary approaches to sonification. It is a fascinating evolution that we are allowed to observe (or hear) in the previous and following decades.

Finally, in Part IV of this book the chapters focus on specific application fields for Sonification and Auditory Display. Although most real Auditory Displays will in fact address different functions (e.g., to give an overview of a large data set and to enable the detection of hidden features), these chapters focus on specific tasks. Assistive Technology (chapter 17) is a promising and important application field, and actually aligns to specific disabilities, such as visual impairments limiting the use of classical visual-only computer interfaces. Sonification can help to improve solutions here, and we can all profit from any experience gained in this field. Process Monitoring (chapter 18) focuses on the use of sound to represent (mainly online) data in order to assist the awareness and to accelerate the detection of changing states. Intelligent Auditory Alarms (chapter 19), in contrast cope with symbolic auditory displays, which are most ubiquitous in our current everyday life, and how these can be structured to be more informative and specifically alerting. The use of sonification to assist the navigation (chapter 20) of activities is an application field becoming more visible (or should we say: audible), such as in sports science, gestural controlled audio interactions, interactive sonification etc. Finally, more and more applications deal with the interactive representation of body movements by sonification, driven by the idea that sound can support skill learning and performance without the need to attend a located visual display. This application area is presented in chapter 21).

Each chapter sets a domain-, field-, or application-specific focus and certain things may appear from different viewpoints in multiple chapters. This should prove useful in catalyzing increased insight, and be inspiring for the next generation of Auditory Displays.

1.4 How to Read

The Sonification Handbook is intended to be a resource for lectures, a textbook, a reference, and an inspiring book. One important objective was to enable a highly vivid experience for the reader, by interleaving as many sound examples and interaction videos as possible. We strongly recommend making use of these media. A text on auditory display without listening to the sounds would resemble a book on visualization without any pictures. When reading the pdf on screen, the sound example names link directly to the corresponding website at `http://sonification.de/handbook`. The margin symbol is also an active link to the chapter's main page with supplementary material. Readers of the printed book are asked to check this website manually.

Although the chapters are arranged in this order for certain reasons, we see no problem in reading them in an arbitrary order, according to interest. There are references throughout the book to connect to prerequisites and sidelines, which are covered in other chapters. The book is, however, far from being complete in the sense that it is impossible to report all applications and experiments in exhaustive detail. Thus we recommend checking citations, particularly those that refer to ICAD proceedings, since the complete collection of these papers is available online, and is an excellent resource for further reading.