A Musical Instrument based on
3D Data and Volume Sonification Techniques

Lars Stockmann, Axel Berndt, Niklas Röber

Department of Simulation and Graphics
Otto-von-Guericke University of Magdeburg
Lars.Stockmann@Email.de
{aberndt|niklas}@isg.cs.uni-magdeburg.de

Abstract. Musical expressions are often associated with physical gestures and movements, which represents the traditional approach of playing musical instruments. Varying the strength of a keystroke on the piano results in a corresponding change in loudness. Computer-based music instruments often miss this important aspect, which often results in a certain distance between the player, his instrument and the performance.

In our approach for a computer-based musical instrument, we use a system that provides methods for an interactive auditory exploration of 3D volumetric data sets, and discuss how such an instrument can take advantage of this music-based data exploration. This includes the development of two interaction metaphors for musical events and structures, which allows the mapping of human gestures onto live performances of music.

1 Introduction

Over the past years, computers have contributed to musical performances in several ways. Already in the late 1960s, computers have been employed to control analogue instruments. The GROOVE synthesizer developed by Max Mathews was one of the first computer controlled analogue synthesizers [14]. Since the introduction of the MIDI standard as a communication protocol, computers have been used as a means for conduction and arrangement in many music productions, but also as a bridge between input devices and synthesizers. In this context, computer have also been used to augment a performance by adding algorithmically generated notes that fit musical structures, as for example in *Music Mouse* [19] or MIDI composing software like *Bars & Pipes*[1]. *Intelligent instruments* like Music Mouse facilitate an easier, more intuitive, approach to the creation of music for the musically inexperienced. At the same time they offer new ways of creating music – even for professional musicians.

In today's productions, external synthesizers are often omitted. Their place is taken by *virtual instruments*, such as Native Instruments'[2] simulation of the B3 organ or the virtual acoustic and electric piano. Even standard consumer hardware is powerful enough for their deployment, and they are used to imitate any kind of instrument in realtime. In contrast to the achievements in sound synthesis, input devices other than MIDI-keyboards are still not common in mu-

sic production, although recently a new research area solely focussing on new musical interaction methods has been established. One example[3] that is planned to be commercially available in the near future is the *reacTable* system, which is described in [10, 12]. Like Crevois et al., who developed an instrument called *Sound Rose* (see [6]), Jordà et al. use a tangible interface as a new intuitive way for live music performances.

Computer-based instruments are designed in a way that a musical controller generates data that is passed to a computer and therein mapped to a single acoustic stimuli of a certain pitch and volume, or to parameters that somehow control an algorithmic composing. The advantage of this approach is that virtually any type of data can be used as input for these instruments. The mapping of arbitrary data to sound (including music) is part of another very important area of research, specifically sonification. It is often used in the development of *Auditory Display* systems, and employed to acoustically convey scientific data. While for a long time, sonification has merely been a part of visualization research, the techniques which were outlined by Gregory Kramer (see [13]) have been developed and successively improved to provide an enhancement, and at places even a superior alternative, to visual representations in science (e.g. [7]). Especially when it comes to the visualization of the inner and outer structures of 3D volumetric data sets. The auditory channel can be used to reduce the load of information that other-

---

[1]Bars & Pipes `www.alfred-j-faust.de/bp/MAIN.html`
[2]Native Instruments `www.nativeinstruments.de`

[3]An overview of some musical controllers can be found at `www-ccrma.stanford.edu/~serafin/NBF/Newport.htm`

wise has to be absorbed by the visual channel alone. The main challenge for sonification research is to find an expressive, intuitive, and comprehensible mapping from the data domain towards sound.

In our sonification system, we employ spatial interactions to facilitate an intuitive method for an auditory exploration of 3D volumetric data sets. It uses a strictly functional mapping of data to complex sounds, based on differences in pitch and volume. This system is the basis for a novel computer-based instrument that can be used without musical experiences. The instrument is designed out of two metaphors: The *Tone Wall* metaphor allows a performer to directly generate a melody, while the *Harmonic Field* is used for an computer-aided accompaniment. Both techniques can be used at the same time. It produces diverse sounds, and allows for a highly interactive performance. It can be shown that spatial interactions inherent a great potential for the use in computer-based instruments.

The paper is organized as follows: After an introduction to the sonification of volumetric data sets in Section 2, we advance by presenting our sonification system in Section 2.1. This includes some technical details regarding our implementation. We then elaborate in Section 2.2 how sonification and computer-based instruments connect, and how live music performances can benefit from an instrument that uses our sonification system. In Section 3 we describe musical data can be derived from spatial gestures in volumetric data sets. The *Tone Wall* metaphor (Section 3.1) specifies the pitch, loudness, and timbre space for melodic purposes. The *Harmonic Field* (Section 3.2) describes how volume data can be used to represent harmonies, broken chord play, and musical textures. Section 3.3 is concerned with a combination of both concepts for the presentation of a one man polyphonic performance. Finally the results are discussed in section 3.4, which also includes possible improvements for further research.

2 Volume Data Sonification

Data sonification is an underdeveloped, but growing field of research. In this section we describe how sonification can be applied to acoustically describe 3D volume data sets. Before we describe our method, we discuss several advantages that make sonification techniques at times superior to a more classic visual examination and presentation of scientific data sets. Examples are monitoring applications, or any type of unfocused operations and processes. The generated acoustic stimuli can be heard without paying direct attention. This yields an improved mobility. Furthermore, Kristine Jørgensen states that the presence of sound increases attention, and eases the perception by intentionally utilizing channel redundancy [11, 8]. A simple example is a flash light that is augmented with a sound while flashing. The proper use of acoustic stimuli in combination with the visual representation also generates a deeper sense of immersion, especially in interactive 3D environments [17]. Gregory Kramer stated that *'spatialized sound can, with limitations, be used to [...] represent three-dimensional volumetric data'* [13]. One reason is that spatialized sound provides a direct mapping to the physical 3D space.

3D volume data occurs in countless fields of research and is used to represent the inner and outer structure of objects or materials in a voxel representation. To find an expressive mapping of the these voxels to sound is one of the main challenges when designing a sonification system.

Since the development of powerful graphics accelerators, there has been much research on finding a good mapping in the visualization domain, but only a few attempts exist to exploit the possibilities of sonification to convey 3D volume data. Minghim and Forrest have suggested methods like the "Volume Scan Process", in which the density inside a *volume probe* is mapped to the pitch of a generated tone [16]. David Rossiter and Wai-Yin Ng traverse the voxels of a 3D volume and map their values to different instrument timbres, amplitudes and pitches [18]. Both systems are controlled through a quite simple mouse/keyboard interface. However, for the sonification of 3D volume data, interaction must not be seen as requirement, but as key aspect. In fact, it is the second most important aspect after the mapping. A direct exploration of the data by, e.g., moving the hand through an interactive 3D environment can provide the user with a better understanding of extent or local anomalies. Both examples of related work lack this ability of a responsive user interface for 3D input like a realtime tracking system, or need to compile the audio data before one can listen to it. The next passage outlines our sonification system, which focuses on direct interactions and an expressive mapping of the inner structure of 3D volume data.

2.1 Spatial Exploration of Volume Data

As mentioned before, a sonification system can greatly benefit from tracking devices that allow a direct exploration of the volume data. In the visualization domain, this is generally done using a certain *Viewpoint-Metaphor*, such as the ones presented by Colin Ware and Steven Osborne [23]. With respect to data sonification, the *eye in hand* metaphor can be easily transformed into the above described volume probe. Instead of a spherical or cubical shape, our approach uses the metaphor of a *chime rod*, as illustrated in Figure 1.

Figure 1: 3D Volume scan through chime rod

The rod can be moved freely through the 3D volume, and is controlled by an interactor that is connected to a 3D tracking device. The advantage of using a rod instead of a spherical or cubical shape is, that the pitch of a tone can be directly associated with a position along the rod. Together with an amplitude-modeling depending on the density value at a certain position, a complex tone is generated. This allows for an intuitive exploration of the inner structures of the volume data. Unfortunately, the system could not be implemented using a MIDI-controlled synthesizer. Instead, we devised our own sound synthesis. A sound is rendered depending on the density distribution of the volume that is in close vicinity of the chime rod. The listeners head is orientation-tracked, and the generated sound is spatialized to provide an additional localization cue for a more immersive experience.

The realtime tracking is achieved using a Polhemus FASTRAK that allows four sensors to be connected. The input data is processed in the client PC that, besides the sonification and sound rendering also performs the visualization. Figure 2 shows the setup of our sonification system.

Figure 2: Setup of the sonification system

Using sonification and visualization at the same time does not only induce the afore mentioned redundancy that eases perception of the data by dispensing the information on two channels, but also allows for multi variate data to be presented directly without the need of switching between different representations. However, it is a crucial aspect of the system that the visualization, which requires a powerful hardware, does not interfere with the audio streaming, even if the system is not equipped with the latest graphics accelerator. Thus, we make great use of multi-threading running the visualization on a low priority to ensure that the audio stream is never interrupted. A scheme of the whole sonification system is illustrated in Figure 3

Figure 3: Schematics of the sonification system

For the sound processing and output in a multi-threading environment we use an audio API that is specially designed for realtime audio applications [20] and revised it for our purposes. The results were promising and beared the idea to introduce music elements into the system. In the next section we elaborate on how sonification methods and computer-based instruments are connected and show how our system can contribute to the research field of the latter.

2.2 Volume Sonification in a Musical Environment

Hunt and Hermann who advance the research of the *model based sonification* impose interaction with the physical world to be a natural cause for acoustic feedback [9]. This feedback is used to gather information about an object. E.g., a bottle of water that is shaken reveals information about its contents. This cause-and-effect chain can not only be used to convey abstract information like the number of messages in an e-mail inbox of a mobile phone [24] but is also a powerful paradigm for computer-based instruments. In the broadest sense, one could consider these instruments as merely a special case of sonification: The sonification of interaction itself. In a musical improvisation interaction can be seen as an expression of emotion and mood.

A computer that is asked to improvise could, of course, not use mood or emotion as basis for its performance, but arbitrary, or specially arranged data. Using music to convey data can have some advantages. Often sonification suffers from annoyance. Paul Vickers and Bennett Hogg state that 'Sonification designers concentrated more on building systems and less on those systems' æsthetic qualities' [22]. Accoustic stimuli that abide by the rules of music are generally more appealing for the listener than sounds that use arbitrary pitch and timbre. It may even stimulate the interactive exploration of data, as the listener self-evidently becomes a music performer by interacting with the dataset. She or he will try to achieve the most pleasant musical result. A distinct variation in the data means a distinct variation in music. Its location can be memorized more easily when the performer 'explores it intentionally' because she or he feels that this particular variation fits best in the current music progression.

However, finding a meaningful mapping of arbitrary multi-dimensional data to music must be considered highly challenging. Some approaches can be found in projects like the Cluster Data Sonification or the Solar Songs by Marty Quinn. In his Image Music[4] sonification, the user can interactively explore a 2D image through music. However, nothing has been done yet in the domain of 3D volume data. Furthermore, the said examples are not intended for live music performances. The interaction is limited to mouse input that does not meet the high responsiveness demanded by a music performer.

Besides the mapping, the method for interacting with the system is crucial for its efficiency. Like the afore mentioned sonification system computer-based instruments mostly use either mouse/keyboard interaction, or are designed to be played with MIDI-keyboards. These demand a certain skill in order to be adequately handled. Systems using the elements of direct interaction as a means for acoustic excitation are scarce. Instruments like the *Fractal Composer* introduced by Chapel, for example, provide a mouse driven graphical user interface [5]. The system composes music using the MIDI protocol in realtime that depends on parameters which are set by the user. She or he has no direct control over the melody or harmony that is generated. This induces a big distance between the performer and the instrument. She or he can only influence the composition on a fairly high level. These systems are referred to as *interactive instruments* [4] or *active instruments*[5]. In contrast, the reacTable and the Sound Rose mentioned earlier are collaborative instruments that use direct interaction.

---

[4]Design Rhythmics Sonification Research Lab `www. drsrl.com/`

Indeed the tangible interface is very intuitive though these attempts are momentarily limited to two dimensional space. Besides the afore mentioned reacTable and Sound Rose The "Morph Table" system that uses morphing techniques presented in [25] is a good example how this interface can be used for music generation [2]. However, the music is also controlled on a rather high level. The system generates transitions between a source- and a target pattern which is applied on precomposed melodies and rhythms. It is not possible to create a melody directly. Furthermore, it is limited to two dimensions.

Chadabe describes a system called *Solo* that uses modified theremin's (see [21]) as 3D input devices to guide the system [3]. Again, the melody is generated algorithmically. The performer controls variables like tempo and timbre. The computer is used for sound synthesis. Thus, this approach is similar to that described in [5] and [2] as the performer has only a global influence on the generated music. However, we think that 3D input devices can be used to intuitively control both, melody and accompaniment. Where the former is generated through a direct mapping of the position to pitch while the latter could benefit from semi automatic composition or precomposed elements. This not only opens the path for diverse improvisations but also can be considered more immersive than just influencing certain aspects of music that is otherwise algorhythmicaly generated.

Our system for interactive exploration of 3D volume data is applicable in that it provides the necessary degrees of freedom to have both aspects in one instrument as well as the responsiveness demanded for a live performance. This makes it possible to develop metaphors for music and sound generation. Two are described in the next section.

3 Volumetric Music

Along the lines of traditional musical instruments, computer-based musical instruments have to find intuitive performative metaphors for musical events. A typical example: To strike one key on the piano means playing its corresponding pitch. The keystroke velocity regulates its loudness. The following sections will describe and discuss this mapping of spatial gestures to musical events and structures, in analogy to the previously discussed image and volume data sonification techniques. The volumetric data represents thereby the medium of interaction and defines the basis for a music processing.

3.1 Tone Wall

A question that arises is: How can different tones be represented in the 3D space? A very intuitive way is a mapping along the vertical axis: low pitches go down, high pitches go up.

But an expressive performance necessitates more than the on/off switching of simple tones. It must be possible to form them. One of the most important means therefore is dynamics (i.e., loudness). In correspondence to the keystroke velocity on the piano, we consider the tone space as a wall. The deeper the performer/interactor punches through that virtual wall (in z-direction) the louder the tone will be played. Short punches produce staccato notes, whereas to hold a tone, the interactor remains in the wall for as long as desired.

An additional parameter is the punch velocity that affects the attack and onset behavior of the tone. A fast punch causes a short attack (a very direct beginning of the tone), and a more percussive onset performed in a slow velocity results in a softer tone at the beginning independent of its dynamic level.

Thus, the y- and z-axis open up the complete bandwidth of expressive tone forming known from keyboard instruments, like the piano, and the punch velocity is a new means to specify details of the tone beginning. However, it would be unwise to not additionally exploit the potentials lying in the x-axis. Many instruments allow the player to vary its timbre to a certain extent, for which the x-axis is predestined. Different timbres can be blended from left to right, e.g. from a very dark sinusoidal waveform over relaxed, clear sound characteristics up to brilliant and very shrill sounds. There are no limitations in sound design in comparison to traditional musical instruments. The complete *Tone Wall* concept is illustrated in figure 4.

For more timbral variances and freedom, it is possible to fill the *Tone Wall* with volumetric data of varying density. It can be employed as static behavior or react on interactions, e.g. like particles that are charged with kinetic energy when they are hit by the interactor device. Due to the freedom to apply any sound synthesis methods, the *Tone Wall* interface is not restricted to pitch based melodic structures, but also for more complex sound structures and noises for contemporary music styles.

3.2 Harmonic Field

In contrast to the *Tone Wall* concept, which specifies an interface to create basic musical events, the *Harmonic Field* is already a pre-composed musical environment, which can be freely explored by the performer.


Figure 4: Tone Wall

It defines a number of regions (as illustrated in figure 5) with their own harmonic content, e.g. a C major harmony in the grey area (harmony 1), a minor in the yellow (harmony 2), a cluster chord in the area of harmony 5, and so on. The performer can move his focus via a head-tracking interaction over the regions to change the harmony that is currently played; he literally looks to the harmonies to play them.

Each harmonic area defines a density gain towards the peak in its center. The density allocation can, of course, also feature more complex shapes, define multiple peaks, holes and hard surfaces. The values can be used for fading techniques, such as those described in [1]; high density can be implemented with a louder volume than low density. But the harmonic field is not restricted to static tones only. Chords can be ornamented by arpeggiated figures and compositional textures can be defined. Instead of using a simple in/out fading, the *texture density* can be adapted: very simple, transparent textures at lower density areas and rich in detail figures at higher densities.

Since harmonic areas can overlap, we applied a number of transition techniques—other than fading that does not satisfy in any situation. Held chords are transitioned part by part. Each part is moving stepwise towards its targeted pitch, where the steps are chosen according to the underlying scale of the harmony (e.g., major, minor, or chromatic scale). Instead of a stepwise movement, the transition can also be done by linear glissando. The transitional pitch is an interpo-

Figure 5: Harmonic Field

lation of the pitches of each harmonic area according to their density weightings. The goal pitch is reached when the old harmonic area is left, or a hole within with zero density is found. With complex cluster-like harmonies, the resulting metrumless clouds do wake associations with György Ligeti's *Clock and Clouds* for women's choir and orchestra.

Compositional textures, in any respect, are not metrumless. They are well-defined sequences of pitches/events in a certain tempo and rhythm. In the case of different tempi, the transitional tempo is an interpolation depending on the density weighting. Since the textures are repetitive, the morphing techniques of Wooller and Brown [25], and the interpolation technique of Mathews and Rosler [15] can be applied to combine the figural material.

However, generative textures were not included at the current state. Therefor, transition techniques for generative algorithms have to be developed and are classified as future work.

3.3 Poly Field

When performing music, it is always desirable of being able to handle both, melodic and harmonic data, simultaneously. Thus, both interfaces, the *Tone Wall* and the *Harmonic Field*, have to be accessible and controllable by one person at the same time.

This is achieved by employing two input devices which can be controlled independently. The user plays melodic gestures on the *Tone Wall* using hand and arm gestures and thereby controls the harmonic progression on the *Harmonic Field* through head gestures and a simple *look*. Furthermore, tilting the head can be used to steer timbral aspects of the *Harmonic Field* play.

Since it turned out to be of some difference to play melodic figures that harmonize with the *Harmonic Field* play, a further quantization is implemented to the *Tone Wall*. The scale that is playable on the *Tone Wall* is matched to the current harmonic base and the punch height is quantized to this scale.

3.4 Discussion

As with all musical instruments, it is necessary to invest a certain amount of practice to learn the intuition and motoric sensitiveness for a confident expressive play. The intuitive correspondence between gestural and musical events, especially in the case of the *Tone Wall* interface, turned out to be very supportive for a steep training curve. Nonetheless, a few practical issues have to be discussed.

The interaction with the *Tone Wall* is subject to a motoric limitation; it is quite exhausting to create fast pace melodies with a proper play over a long period of time. Tracking latencies (ranging between 8–10 ms) and sampling artifacts (interaction sample rate is 60 Hz with two interactors) also slightly interfere with the play and the possible speed of interaction.

Because of the absence of any visual reference points, it is at times difficult to meet the intended pitches. A calibration, according to the size of the performer, can lower this problem; his body can provide several reference points.

For playing melodic intervals, the interactor has to leave the wall, jump over the unwanted pitches, and punch back into it. Moving the interactor within the wall would trigger pitches in-between. Thus, melodic intervals are always adherent with short pauses. A legato articulation is not possible within this approach. Therefore, an interactor speed dependency has to be incorporated: a pitch is only played if the interactor's velocity is below a certain threshold. Pitches can be skipped by faster movements even within the wall. Since this raises the problem of creating fast pace melodies, this mode has to be detachable, e.g. by a button on the hand interactor.

The same approach could be useful to reduce the *wah*-effect when playing a pitch. The punch always hits the low dynamics area at the wall surface first, and the loud dynamics afterward. Hence, each tone fades in, even with fast punches that do only effect a more direct tone attack. Although the interaction sampling rates

used lower this effect, a velocity dependent sampling of the interactor would make the dynamic level more accessible.

However, all performative means of expression are available and easy to perform—dynamics and emphasis, articulation, (de-)tuning, timbral and articulational (glissando, triller etc.) effects.

For the *Harmonic Field* the composer is free to define any chords, assign them to any timbral instrumentation and figurative ornamentation, and combine them by overlapping. He can actually define any compositional and timbral texture and it can be explored freely by the player. The player, however, is fixed to this predefined set, unable to create new chords and textures interactively during the performance. Furthermore, the three-dimensional space cannot be explored adequately using head-orientation alone, i.e. looking at a harmonic area from a relatively fixed position, which allows only an exploration in 2D. The player should be able to move freely in 3D space. This raises conflicts with the *Tone Wall* metaphor. A possible solution is to position the *Tone Wall* always in front of the player and reposition it when the player moves through the *Harmonic Field*.

However, the combination of the *Harmonic Field* with the *Tone Wall* interface open up a very large musical bandwidth with more timbral freedom than any traditional musical instrument can offer. The three-dimensional setup of harmonic structures and their density-dependent ornamentation textures are also unique and provides an inspiring platform especially for performing contemporary music.

4 Conclusion, Future Work

In this paper we presented a gesture based approach towards virtual musical instruments. We introduced the conceptual basis, which is a novel interaction mechanism developed for the interactive auditory exploration of volumetric data sets. For their sonification we devised the musical metaphors of the *Tone Wall* and the *Harmonic Field*, and conceived their sonic behavior in a way that the interaction with them produces musical events and aesthetic structures, like tones, melodies, timbre effects, chords, and textures. We discussed assets and drawbacks of these metaphors and outlined advancements.

3D interaction devices open up a multitude of new possibilities for the design of computer-based instruments. Their big potential lies in their intuitive association with physical human gestures and musical events, for which the interaction with virtual volume data turned out to be the medium of choice. Future work includes the development of further metaphors and the integration of serial and generative concepts. The volumetric interaction interface also opens up a promising possibility for the conduction of music.

The musical volume representation concept is also a novel view on musical structure and elements, enabling new compositional forms and means of expression. Here lies the biggest potential of new computer-based instruments. It is unnecessary to imitate traditional instruments to create music that is performed better with the real ones. If one wants to play a piano, violin, trombone etc. the real ones perform always better. New instruments should not imitate them, but stand for a confident self-reliance to open up new possibilities for new music to constitute their right to exist.

References

[1] A. Berndt, K. Hartmann, N. Röber, and M. Masuch. Composition and Arrangement Techniques for Music in Interactive Immersive Environments. In *Audio Mostly 2006: A Conf. on Sound in Games*, pages 53–59, Piteå, Sweden, oct. 2006. Interactive Institute, Sonic Studio Piteå.

[2] A. R. Brown, R. W. Wooller, and T. Kate. The Morphing Table: A collaborative interface for musical interaction. In A. Riddel and A. Thorogood, editors, *Proceedings of the Australasian Computer Music Conference*, pages 34–39, Canberra, Australia, july 2007. Australian National University Canberra.

[3] J. Chadabe. Interactive Music Composition and Performance System. United States Patent Nr. 4,526,078, july 1985. filed sep. 1982.

[4] J. Chadabe. The Limitations of Mapping as a Structural Descriptive in Electronic Instruments. In *Proceedings of the Conference on New Instruments for Musical Expression (NIME-02)*, Dublin, Ireland, may 2002.

[5] R. H. Chapel. *Realtime Algorithmic Music Systems From Fractals and Chaotic Functions: Towards an Active Musical Instrument*. PhD thesis, University Pompeu Fabra, Department of Technology, Barcelona, Spain, sept. 2003.

[6] A. Crevoisier, C. Bornand, A. Guichard, S. Matsumura, and C. Arakawa. Sound Rose: Creating Music and Images with a Touch Table. In *NIME '06: Sixth meeting of the International Conference on New Interfaces for Musical Expression*, pages 212–215, Paris, France, 2006. IRCAM—Centre Pompidou.

[7] W. T. Fitch and G. Kramer. Sonifying the body electric: Superiority of an auditory over a visual display in a complex multivariate system.

In G. Kramer, editor, *Auditory Display: Sonification, Audification, and Auditory Interfaces*, Boston, MA, USA, 1994. Addison-Wesley.

[8] C. Heeter and P. Gomes. It's Time for Hypermedia to Move to Talking Pictures. *Journal of Educational Multimedia and Hypermedia*, winter 1992.

[9] A. Hunt and T. Hermann. The Importance of Interaction in Sonification. In *ICAD 04 — Tenth Meeting of the International Conference on Auditory Display*, Sydney, Australia, july 2004.

[10] S. Jordà, M. Kaltenbrunner, G. Geiger, and R. Bencina. The reacTable. In *Proceedings of the International Computer Music Conference*, Barcelona, Spain, 2005. International Computer Music Association.

[11] K. Jørgensen. On the Functional Aspects of Computer Game Audio. In *Audio Mostly 2006: A Conf. on Sound in Games*, pages 48–52, Piteå, Sweden, oct. 2006. Interactive Institute, Sonic Studio Piteå.

[12] M. Kaltenbrunner, S. Jordà, G. Geiger, and M. Alonso. The reacTable: A Collaborative Musical Instrument. In *Proceedings of the Workshop on "Tangible Interaction in Collaborative Environments" (TICE), at the 15th International IEEE Workshops on Enabling Technologies*, Manchester, U.K., 2006.

[13] G. Kramer, editor. *Auditory Display: Sonification, Audification, and Auditory Interfaces*. Addison-Wesley, Boston, MA, USA, 1994.

[14] M. V. Mathews. The Digital Computer as a Musical Instrument. *Science*, 142:553–557, nov. 1963.

[15] M. V. Mathews and L. Rosler. Graphical Language for the Scores of Computer-Generated Sounds. *Perspectives of New Music*, 6(2):92–118, Spring–Summer 1968.

[16] R. Minghim and A. R. Forrest. An Illustrated Analysis of Sonification for Scientific Visualisation. In *IEEE Conference on Visualization*, Atlanta, USA, oct. 1995.

[17] Niklas Röber and Maic Masuch. Playing Audio-only Games: A compendium of interacting with virtual, auditory Worlds. In *Proceedings of 2nd DIGRA Gamesconference*, Vancouver, Canada, 2005.

[18] David Rossiter and Wai-Yin Ng. A system for the complementary visualization of 3D volume images using 2D and 3D binaurally processed sonification representations. In *Proceedings of the 7th conference on Visualization*, pages 351–354, San Francisco, USA, 1996. IEEE Computer Society Press.

[19] Laurie Spiegel. Music Mouse. `http://retiary.org/ls/programs.html`, 2004.

[20] Lars Stockmann. Designing an Audio API for Mobile Platforms. Internship report, 2007.

[21] L. S. Theremin. Method of and Apparatus for the Generation of Sounds. United States Patent Nr. 73,529, dec. 1924.

[22] Paul Vickers and Bennett Hogg. Sonification abstraite/sonification concrète: An 'Æsthetic perspective space' for classifying auditory displays in the ars musica domain. ICAD 06 – 12th International Conference on Auditory Display, Juni 2006.

[23] Colin Ware and Steven Osborne. Exploration and virtual camera control in virtual three dimensional environments. *SIGGRAPH Comput. Graph.*, 24(2):175–183, 1990.

[24] J. Williamson, R. Murray-Smith, and S. Hughes. Shoogle: Excitatory Multimodal Interaction on Mobile Devices. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 121–124, New York, USA, 2007. ACM.

[25] R. W. Wooller and A. R. Brown. Investigating morphing algorithms for generative music. In *Third Iteration: Third International Conference on Generative Systems in the Electronic Arts*, Melbourne, Australia, dec. 2005.