
Zirkonium: Non-invasive software for sound spatialisation*

CHANDRASEKHAR RAMAKRISHNAN

Native Systems Group, ETH Zürich, Clausiusstrasse 59, CH-8006 Zürich, Switzerland
E-mail: cramakrishnan@acm.org

Zirkonium is a flexible, non-invasive, open-source program for sound spatialisation over spherical (dome-shaped) loudspeaker setups. By *non-invasive*, we mean that Zirkonium offers the artist spatialisation capabilities without forcing her to change her usual way of working. This is achieved by supporting a variety of means of designing and controlling spatialisations. Zirkonium accommodates user-defined speaker distributions and offers HRTF-based headphone simulation for situations when the actual speaker setup is not available. It can acquire sound sources from files, live input, or via the so-called *device mode*, which allows Zirkonium to appear to other programs as an audio interface. Control data may be predefined and stored in a file or generated elsewhere and sent over OSC. This paper details Zirkonium, its design philosophy and implementation, and how we have been using it since 2005.

1. INTRODUCTION

The Institut für Musik und Akustik (IMA) at ZKM | Zentrum für Kunst und Medientechnologie, Karlsruhe has as its primary mission commissioning, producing and presenting concerts of electroacoustic music. It was founded in 1989 and has been at its present location in the ZKM complex since 1997. IMA is home to composers' ateliers and studios for recording and producing electronic and acoustic music, as well as a concert space, the Kubus, a rectangular concert hall inside the iconic cubical building at the main entrance to ZKM. The Klangdom project was launched in 2004 in order to extend the existing capabilities of the Kubus and offer a more immersive sound-spatialisation system.

The Klangdom, as with all spatialisation systems, is made up of three components: loudspeaker positioning, a technique for routing sound to the speakers, and a controller for defining the routing. In the Klangdom, the Zirkonium software defines the latter two.

Zirkonium's primary task is offer the user a model of a room, let her place virtual sound sources in the space, and route sound to speakers to realise this placement. This paper describes how Zirkonium performs this. We begin by providing historical background on sound spatialisation approaches (Section 2). We then describe

the Klangdom (Section 3) and the motivation behind Zirkonium and how it handles spatialisation (Section 4), and discuss its implementation, providing details on certain interesting aspects (Sections 5 and Section 6). We have been using Zirkonium in production since 2005, and we relate our experience using Zirkonium in a variety of concert situations (Section 7). We conclude with an assessment of the lessons we have learned and suggestions of future work still to be done.

2. HISTORICAL BACKGROUND

Spatial distribution of sound events has played an important role in electroacoustic music since its very beginnings. As early as 1951, the studio at the Radio-diffusion-Télévision Française (RTF) employed a quadraphonic spatialisation system with two front channels, one channel in the back, and one above the listeners and developed a controller, the *pupitre d'espace*, used for live control of spatialisation (Zvonar 2005).

Richard Zvonar (Zvonar 2005) and Leo Küpper (Küpper 1984) offer fascinating histories of the development of spatialisation in music, each with his own focus. Through their research, we can see two broad tendencies in systems for sound spatialisation devised since the late 1970s. We label these the *acousmatic approach* and *simulation approach* (Table 1).

The acousmatic approach can be thought of as an extension of *musique concrète*. Similar to *musique concrète*'s focus on the capabilities of the tape machine as a source of generating musical material, the acousmatic approach focuses on the loudspeaker and its position as the way to organise sound in space. This is well illustrated by several different systems.

The acousmonium was developed at the Groupe de Recherches Musicales (GRM) in the mid-1970s. An acousmonium is made up of different types of loudspeakers distributed throughout a room. It is usually played live by a composer/performer who routes the audio of a piece (often 2-channel) to the different loudspeakers, taking advantage of the sound reproduction characteristics and physical placement of the loudspeakers to realise a spatialised live performance.

Another system with a similar philosophy, but a different realisation, is the sound dome of the sort

*The Klangdom project was generously funded by the Zukunfts-offensive initiative of Baden-Württemberg.

Table 1. Summary of spatialisation systems.

System	Speaker distribution	Speaker type	Control	Type
Acousmonium	Mostly Frontal	Multiple Types	Mixer	Acousmatic
BEAST	Frontal Emphasis	Multiple Types	Mixer/Software	Acousmatic
Sound Dome	Spherical	Single Type	Mixer/Software	Acousmatic
Ambisonics	Spherical	Single Type	Software	Simulation
Wave Field Synthesis	Panels	Single Type	Software	Simulation

championed by Leo Küpper (Küpper 1984) and exemplified by the German Pavilion at the World Expo '70 in Osaka, Japan. A sound dome differs from an acousmonium in that it is made up of one type of speaker and specifies a specific distribution of the speakers in space (on the surface of a sphere). Nonetheless, both share the core spatialisation philosophy that the position of the loudspeaker determines the origin of the virtual sound source.

The technique of vector-base amplitude panning, or VPAB, represents a further development in this way of thinking (Pulkki 1997). As with the dome, it places the speakers on the surface of a sphere, but provides a more sophisticated scheme for creating a virtual sound source by routing the audio to the three speakers closest to the position of the virtual sound source.

The simulation approach, on the other hand, uses signal processing to simulate physical acoustics to produce the illusion of moving sound sources. Whereas the acousmatic approach can be realised in the analogue domain with just a mixer, the simulation approach requires the intercession of a computer. One of the early pioneers was John Chowning, who in his 1971 article (Chowning 1971) described techniques for simulating moving sound sources over a quadraphonic speaker setup. The *Spatialisateur* or *Spat* (Jot and Warusfel 1995) developed at IRCAM is an evolution of this idea and lets the user specify a sound source's position as well as its reverberation characteristics, which are also important for the perception of sound localisation.

Other recent work along the direction has focused on Wave Field Synthesis and Ambisonics. Both of these techniques have as their basis the wave equation. Wave Field Synthesis (Berkhout 1988) uses a large number of small loudspeakers to synthesise an approximation of the wavefront that would be observed were there a sound source at the specified position. Ambisonics similarly approximates the sound field at a point, but does so using ordinary speakers (Daniel 2001).

3. THE KLANGDOM

One of the challenges of the Klangdom project was choosing components that integrated well, conceptually and physically, into the traditions and architecture of IMA. A spatialisation system conceived with

the unconstrained freedom of a tabula rasa would certainly take a different form. But for the Klangdom, a goal of the design was to extend the existing capabilities of the Kubus, but at the same time to support the established day-to-day activities of IMA.

Composers and musicians who perform at IMA | ZKM come from a variety of backgrounds and represent a multitude of approaches to music. These different approaches can make contradictory demands on a spatialisation system. Some composers place much worth on having the system deliver the sounds they have composed as faithfully as possible to the loudspeakers. Other composers are more interested in hearing a convincing illusion of movement and are willing to tolerate their sounds being processed, filtered and reverberated to achieve that effect. This is parallel to the distinction between the acousmatic and simulation approaches to spatialisation.

The hardware and software of the Klangdom was designed to accommodate, as much as possible, both uses. The loudspeakers are distributed in a dome-shape, enveloping the audience. This configuration is suitable for VBAP as well as Ambisonics. The loudspeakers are suspended from three concentric rings around the Kubus. The first two rings are made of track, allowing the speakers to be easily moved, and the third ring can be raised or lowered. We chose MeyerSound UPJ-1P loudspeakers for their quality of sound reproduction, as well as the fact that they are normal, full-range concert loudspeakers. This offers the advantage that the speakers can be removed from the dome temporarily and used elsewhere if there is an acute need. This was motivated by the reality of our concert schedule.

The Klangdom was completed and first used in concert in 2005 (Ramakrishnan, Gossmann, Brümmer and Sturm 2006). The Zirkonium software, developed to control the Klangdom, has since undergone two major iterations. This paper describes the current, as of February 2009, version.

4. VISION

The goal behind Zirkonium is to create a spatialisation programme specifically for composers of electroacoustic music. It should not interfere with a composer's usual process, and should allow her to use

her preferred tools whenever possible. The model we envisioned is that of a series of services for spatialisation to be leveraged and utilised in combination with other programmes (Digital Audio Workstations, Max/MSP, SuperCollider, etc.), as opposed to a monolithic application that imposes itself into all aspects of the compositional process. Thus, there was a conscious design decision that Zirkonium focus solely on panning and the construction of time-based panning choreographies, while allowing other domains of spatialisation, such as reverberation, to be implemented elsewhere and still be incorporated.

To this end, Zirkonium defines a simple, straightforward panning model and provides an open infrastructure with entry points where other programmes may dock themselves. This enables a full spectrum of possibilities for sound and panning to be generated and controlled in real time or composed beforehand for storage and playback. Furthermore, since composers will want to perform their pieces in different locations, where different infrastructure may be available, Zirkonium lets composers describe the relevant elements of the concert space and adapts the spatialisation to those parameters, minimising setup overhead and obviating the need for interchange formats.

5. ZIRKONIUM

At its core, all Zirkonium does is give the user a model of a room, let her place virtual sound sources in the space, and route sound to speakers to realise this request. When so described, it sounds simple, but there are many design decisions involved. As with the hardware, flexibility and pragmatism were the guiding principles in making these decisions. Thus, we defined the goal of offering this functionality in a minimalist way. This resulted in an acousmatic system. For many uses, this is sufficient. In other situations, Zirkonium can also be used as a basic building block and extended with simulations.

The philosophy of pragmatism and flexibility manifests itself in many places. Zirkonium uses VBAP for spatialisation, but it also defines an interface to allow for use of other algorithms, such as Ambisonics. Zirkonium does not add reverb, distance cues, and simulation of movement artefacts (e.g., Doppler shift), but it does support the user in doing this herself (see Section 7 for a further discussion).

Furthermore, Zirkonium has been designed to be used in other spaces, not just at ZKM. Users may define their own loudspeaker configurations that describe their local environment. And realising that the environment itself may not always be at the disposal of the artist, Zirkonium can simulate speaker configurations for headphone listening.

The following section, Section 6, discusses Zirkonium's spatialisation model – how it represents virtual

sound sources in space and converts that representation into sound. The section after that, Section 7, concerns itself with the process of defining and acquiring sound sources.

6. PANNING MODEL

Zirkonium takes a mono sound source and point in space and creates the illusion of the sound emanating from that point in space. This is achieved through the interaction of several components: the panning algorithm, VBAP; the distribution of the speakers in the room (speaker setup); and the specification of the desired position of the virtual sound source (figure 1). This section describes these three pieces of the puzzle.

6.1. VBAP

Equal power panning (EPP) is a standard technique for stereo panning while keeping the perceived loudness constant (Roads 1996: 460–1). Vector base amplitude panning is an extension of EPP to speakers distributed over the surface of a sphere (Pulkki 1997). In VBAP, virtual point sources are panned using a combination of the closest three speakers and scaled such that the perceived loudness does not change.

6.2. Speaker setup

The first step necessary to realise a spatialisation is the definition of a speaker setup. The speaker setup defines for Zirkonium the positions of the speakers in the room, information necessary to use VBAP (figure 2).

Speaker positions can be specified freely in 3-D Cartesian coordinates (X/Y/Z). As the speakers are added, Zirkonium does two things:

1. it projects the position of the new speaker onto the surface of a sphere, as required by VBAP, and




Figure 1. Positioning a sound in space with Zirkonium.




Figure 2. Defining speaker positions in Zirkonium.




Figure 3. Triangular meshes for two dome configurations: 48 speakers and 12 speakers.

2. it computes a triangular mesh that decomposes the individual speaker positions into triplets of speakers (figure 3).

Although Pulkki himself proposes an algorithm for computing a triangular mesh (Pulkki and Lokki 1998), we use Delaunay triangulation, a well-known and widely studied method from computational geometry (Bern and Eppstein 1992). In particular, we use the implementation provided by Triangle (Shewchuk 1996).

A further feature offered by the speaker setup is the ability to simulate it for headphone listening. This simulation is performed using a head-related transfer function (HRTF).

6.3. Positioning virtual sources

Positioning a virtual sound source involves taking a user-specified position and generating gain coefficients for each loudspeaker. The properties of VBAP imply




Figure 4. Projecting a Cartesian coordinate onto the sphere.

that at most three speakers per virtual sound source will have non-zero coefficients.

Virtual source positions may be specified in either Cartesian (X/Y) or spherical (azimuth/zenith) coordinates. One of the initially confusing aspects of Zirkonium is that source positions are two dimensional. This is because VBAP assumes that speakers are located on the surface of a sphere (Pulkki 1997). The surface of a sphere, though embedded in three dimensions, is two dimensional (i.e., a surface). Thus, only two coordinates are necessary to specify a point on it – the third coordinate is implicit (figure 4).

The first step in applying VBAP is to determine which triplet of speakers should be used. As pointed out in Pulkki's original article (Pulkki 1997), this can be easily done by computing the VBAP coefficients for each of the speaker triplets. Only one of these triplets will yield all non-negative coefficients. These are the ones that are then passed on to the mixer to realise the virtual sound source.

In addition to position, a *spread* can be specified for any of the axes (that is, azimuth and zenith, or x and y). This samples the specified range and creates extra virtual sound sources to distribute the sound over an arc, line or area (figure 5). The result is rescaled such that the perceived loudness stays constant. The current implementation does *not* do any extra processing, such as de-correlation, to create the impression of a diffuse sound distributed over an area (Kendall 1995), though this is on the list of future work we would like to do.

The matrix arithmetic used to implement VBAP is carried out by the LAPACK routines for fast, efficient linear algebra (Anderson, Bai, Bischof, Blackford, Demmel, Dongarra, Du Croz, Greenbaum, Hammarling, McKenney and Sorensen 1999). The coefficients for the speakers and the total panning of various sources are all mixed together using the Apple Matrix Mixer Audio Unit. The Matrix Mixer Audio Unit efficiently implements its sample processing using routines highly optimised for each of the Apple-supported processors (G4, G5, Intel). By leveraging LAPACK and the Matrix Mixer Audio Unit, Zirkonium can efficiently implement VBAP.




Figure 5. Spreads in spherical and Cartesian coordinates.

6.4. Circumventing the spatialiser

Our experience from concerts has shown us that it is useful to occasionally circumvent the spatialiser. For this purpose, Zirkonium offers direct outs. Direct outs are simply sent to a specified channel of the audio device and are not positioned (figure 6).

7. DEFINING, STORING AND PLAYING SPATIALISATIONS

Zirkonium offers several different ways to create and play spatialisations (Table 2). It defines a file format and editor to define precomposed spatialisations. It also defines interfaces for controlling spatialisations in real time. Positions of sound sources may be specified via OSC. Sound sources themselves may be generated from other programs, such as Max/MSP, SuperCollider or a digital audio workstation (DAW) such as Logic, and passed on to Zirkonium.

7.1. Stored spatialisations

The Zirkonium file format is designed for composed spatialisations. In the editor, the user can specify movements by a start time (in absolute time from the beginning of the piece) and duration. A movement can alter the position of a sound source, its spread, or both.

Table 2. Options for defining sound sources and control data in Zirkonium.

Audio vs. control source	File	OSC	Audio unit
File	X	X	
Input	X	X	
Device mode		X	X




Figure 6. The architecture of the Zirkonium spatialiser.

The file also specifies sound sources to be spatialised (figure 7). The sound sources may be located in audio files, or come from live input. The audio files themselves may include any number of channels and be in a variety of formats, including common LPCM formats (AIFF, WAV, SND), as well as compressed formats (MP3, MPEG-4 Audio), and even more exotic formats such as SDII. The user defines how the channels of a file or live input map to sound sources. Individual channels may map to zero, one, or multiple virtual sound sources.

7.2. Device mode

An alternative way to acquire sound sources is via device mode. In device mode, Zirkonium appears as an audio device to other software. This gives any audio program, including Max/MSP, SuperCollider, and DAWs such as Logic or Digital Performer access to Zirkonium.

Though similar results may be achieved by other means – either using two computers and connecting the output of one computer to another running Zirkonium, or by using the Jack audio router (Letz, Fober, Orlarey and Davis 2004) – the device mode offers advantages (figure 8). For one, it is usually simpler to set up. And, furthermore, device mode does not add latency.

7.3. OSC control

All the aspects of a virtual sound source – its position, spread and gain – may be controlled by an OSC interface. The OSC namespace provides for several different ways of positioning a sound source: Cartesian coordinates, spherical coordinates, or placing a sound source at the position of the nearest speaker (Table 3).

OSC control is the best option when a complex algorithm defines the source position. It is the only




Figure 7. A Zirkonium spatialisation file.




Figure 8. The device mode.

Table 3. The Zirkonium OSC namespace.

OSC namespace	Description
/pan/az	Pan channel to specified spherical coordinate
/pan/xy	Pan channel to specified Cartesian coordinate
/pan/speaker/az	Pan channel to the speaker closest to the specified spherical coordinate
/pan/speaker/xy	Pan channel to the speaker closest to the specified Cartesian coordinate
/master/gain	Set master gain

option when the spatialisation must be controlled or modified live.

8. STRATEGIES FOR WORKING WITH ZIRKONIUM

The first concert with Zirkonium and the Klangdom was in April 2006. Since then, Zirkonium has been regularly in use, not only in the Kubus, but in concerts at other venues as well. It has been used to spatialise pieces originally written for other formats, it has been used for live spatialisation, and it has been part of the compositional process of new pieces designed specifically for the Klangdom. We relate some of our experiences from these different situations.

8.1. Spatialising existing pieces

Most of the performances with Zirkonium have involved pieces that were not composed specifically for the Klangdom. Many of these were composed for either 4- or 8-channel playback. We have found that Zirkonium is a good environment for working with such pieces, particularly in cases where it is possible to access the DAW sessions that were used to create the final mix.

We illustrate the general technique by way of the example of Ludger Brümmer's originally quadraphonic composition *Glasharfe*. The original version already contained sound sources moving within the four channels as well as quadraphonic reverberation. The challenge for the Klangdom version was to keep the movement and spatial sensation of the original, while adapting it for a more immersive environment.

To do this, we went back to the DAW session used to create *Glasharfe*. The 4-channel version is a mix-down of 12 quadraphonic tracks. For the Klangdom, we made four sub-mixes, each involving three of the original tracks. In Zirkonium, we created four virtual quadraphonic groups. Each group maintains the standard spatial relationship of a quadraphonic speaker setup, but can be moved around and placed in different parts of the Klangdom (figure 9). This




Figure 9. Schematic for the *Glasharfe* spatialisation. Quadruplic groups may be moved, rotated, brought closer to the centre of the dome, or pushed towards the edge, but, usually, remain positioned on the corners of a square.

way, the movements as originally composed still make sense, but the final result is much more enveloping. Adding additional movement of the quadruplic groups and briefly breaking the relationship between the channels within a group are also useful effects.

Since quadruplic and octophonic pieces are common in electronic music, it is important to have techniques for translating them to the Klangdom. We have found the technique described here to be successful.

8.2. Movement and distance effects

As the position of a sound source is changed, Zirkonium alters the panning to reflect the change, but does not render the physical effects of movement such as Doppler shift. Zirkonium also restricts panning to direction and does not account for the distance of sound sources, which affects volume and reverberation.

Zirkonium provides tools for the user to remedy this herself. Multi-channel reverberation is quite complicated and highly subject to individual taste. Reverb over 48-channels is all the more complicated. But there are a number of plug-ins that produce reverb for 4, 5 or 8 channels. Some users may have even built their own reverb in Max/MSP or SuperCollider.

The Zirkonium solution is to leverage one of these reverbs. The user creates four (or five or eight) virtual sound sources in Zirkonium and positions them to retain their quadruplic distribution. A 4-channel reverb provides the reverberation, and Zirkonium can be used to rotate or draw the channels into or

away from the centre. Distance can be handled similarly.

The composer Todor Todoroff did exactly this for his piece *Around and above, weightless ...*.

8.3. Real-time control

With live control, a performer can use the Zirkonium as an instrument and *play* a spatialisation. We have built a very simple live controller for spatialisation using the JazzMutant Lemur interface. This controller lets a user create several rings of virtual sound sources and control the height in the dome and rotational speed of each ring. Despite its simplicity, it has been used to great effect by the composers Gilles Gobeil and Robert Normandeau.

Another use for real-time control is to algorithmically determine the positions of sound sources. For Alvin Curran's *TransDADAExpress 2*, Frank Halbig programmed a Max/MSP controller that dynamically positioned channels based on chaotic oscillators. The audio for the piece was played back from file, but the positions of the sound sources was determined in real time and controlled via OSC.

9. FUTURE WORK

As with any software project, Zirkonium is continually evolving, and we are always trying to incorporate better solutions for the problems we encounter.

A particularly major gap is Zirkonium's poor integration with DAWs. This is the preferred environment for many electroacoustic composers, and better tools for controlling Zirkonium from them would be an appreciated addition.

Extending the spatialisation capabilities of Zirkonium is another present goal. First, it would be nice to support Ambisonics as well as VBAP, since the basic premises of Zirkonium are compatible with Ambisonics as well. Furthermore, we would like to add decorrelation to better create the impressions of diffuse areas of sound as opposed to point sources.

10. CONCLUSION

At this point, it seems appropriate to reflect on the lessons we have learned from three years of using and revising Zirkonium, and to offer some advice to media technologists building spatialisation systems and to composers who use them.

First, we recommend, of course, checking out Zirkonium (www.zkm.de/zirkonium). It is open source and available without cost. As we have presented here, we think it does a good job of solving the problems it sets out to address. But it is not going to be the solution for everyone. All software involves design decisions and compromises, and those made

by Zirkonium may not be appropriate for you. Still, we think we can extract some general advice, applicable to all spatialisation systems.

Expect to handle standard formats. Composers will continue to compose for stereo, quadrasonic, 5.1, and octophonic configurations. Research into formats for description and interchange of spatial sound is an ongoing topic (Kendall, Peters and Geier 2008), but no standard format has yet been agreed upon. So have a good solution for incorporating content in standard formats to create an appealing spatialisation with little work.

And advice for composers: compose for the format in which you expect your work will primarily be heard. If that format is four channels, then optimise for four channels. But keep your source material (ProTools session, or whatever) in a form that you can easily produce sub-mixes of the entire piece. These sub-mixes can then be placed in space to create a convincing result.

Take advantage of the spatialisation capabilities of the format you compose for. If you are targeting a 4-channel system, then go ahead and add movement, reverberation, and so on. These effects can be leveraged when translating to a pluriphonic rendering. Movement through space offered by, for example, a sound dome or WFS system is a powerful effect, but the ability to use space itself, and place channels that would otherwise overlap in different locations, is an even more powerful effect.

REFERENCES

- Anderson, E., Bai, Z., Bischof, C., Blackford, S., Demmel, J., Dongarra, J., Du Croz, J., Greenbaum, A., Hammarling, S., McKenney, A. and Sorensen, D. 1999. *LAPACK Users' Guide*. Philadelphia, PA: Society for Industrial and Applied Mathematics.
- Berkhout, A. J. 1988. A Holographic Approach to Acoustic Control. *Journal of the Audio Engineering Society* **36**(12): 977–95.
- Bern, M. and Eppstein, D. 1992. Mesh Generation and Optimal Triangulation. In Ding-Zhu Du and Frank Hwang (eds.), *Computing in Euclidean Geometry*. Singapore: World Scientific.
- Chowning, J. M. 1971. The Simulation of Moving Sound Sources. *Journal of the Audio Engineering Society* **19**: 2–6.
- Daniel, Jérôme. 2001. Représentation de champs acoustiques, applications à la transmission et à la reproduction de scènes sonores complexes dans un contexte multimedia. PhD thesis, Université Paris VI.
- Jot, J. and Warusfel, O. 1995. Spat ~ : A Spatial Processor for Musicians and Sound Engineers. *CIARM: International Conference on Acoustics and Musical Research*. Ferrara, Italy. Available online at: <http://catalogue.ircam.fr/articles/textes/Jot95a>.
- Kendall, G. 1995. The Decorrelation of Audio Signals and its Impact on Spatial Imagery. *Computer Music Journal* **19**(4): 71–87.

- Kendall, G., Peters, N. and Geier, M. 2008. Towards an Interchange Format for Spatial Audio Scenes. *Proceedings of the 2008 International Computer Music Conference, Belfast, UK*.
- Küpper, L. 1984. Musik und Technik heute – Die Parameter des Klangraumes in der Musik. *Ars Electronica Festival-Katalog 1984 (Band 01)*. Available online at http://90.146.8.18/de/archiv_files/19841/1984_016.pdf.
- Letz, S., Foer, D., Orlarey, Y. and Davis, P. 2004. Jack audio server: MacOS X port and multi-processor version. Paper presented at the Sound and Music Computing '04 Conference, Paris, France.
- Pulkki, V. 1997. Virtual Sound Source Positioning Using Vector Base Amplitude Panning. *Journal of the Audio Engineering Society* **45**(6): 456–66.
- Pulkki, V. and Lokki, T. 1998. Creating Auditory Displays to Multiple Loudspeakers using VBAP: A Case Study with Diva Project. *International Conference on Auditory Display*. Glasgow: British Computer Society.
- Ramakrishnan, C., Gossmann, J., Brümmer, L. and Sturm B. 2006. The ZKM Klangdom. *Proceedings of the 2006 International Conference on New Interfaces for Musical Expression*. Paris: IRCAM, 140–3.
- Roads, C. 1996. *The Computer Music Tutorial*. Cambridge, MA: MIT Press.
- Shewchuk, J. R. 1996. Triangle: Engineering a 2D Quality Mesh Generator and Delaunay Triangulator. In Ming C. Lin and Dinesh Manocha (eds.), *Applied Computational Geometry: Towards Geometric Engineering*. Berlin: Springer-Verlag, Berlin. Lecture Notes in Computer Science (1148): 203–222.
- Zvonar, R. 2005. A History of Spatial Music. *eContact!* 7.4. Université Concordia. Accessed on 12 February 2009 at: http://cec.concordia.ca/econtact/Multichannel/spatial_music.html.