BETWEEN MAPPING, SONIFICATION AND COMPOSITION: RESPONSIVE AUDIO ENVIRONMENTS IN LIVE PERFORMANCE

Christopher L. Salter

Faculty of Fine Arts
Concordia University
Interactive Performance
and Sound, Hexagram
Institute
Montreal, Canada
csalter@gmx.net

Marije A.J. Baalman

Technische Universität
Berlin
Institute for Audio
Communication
Berlin, Germany
marije@nescivi.nl

Daniel Moody-Grigsby

Design and Computation
Arts
Concordia University
Montreal, Canada
dangrigsby@gmail.com

ABSTRACT

This paper describes recent work on a large-scale, interactive theater performance entitled *Schwelle* as a platform to pose critical questions around the conception, design and implementation of what is commonly labeled responsive audio environments. The authors first discuss some principal issues in the design of responsive audio environments specifically within the domain of stage performance, addressing existing human-computer interaction paradigms and discussing three key areas: sensing, mapping and data sonification. Next, we discuss larger questions of composition in relation to these key areas, suggesting that potential strategies cross three different domains: mapping within algorithmic composition, data sonification techniques, and time-based evolutionary processes emerging from dynamical systems theory. We then examine in detail the recent work on *Schwelle*, which employs real time, distributed sensor data to drive a continuous dynamical system-based composition engine. The project's conceptual and technical challenges are discussed as well as audience evaluation and feedback from the first presentation in Berlin in February 2007. This presentation lead the authors to re-iterate the design and build an additional state-system layer into the dynamical system in order to generate more perceivable sonic structures on both the macro as well as meso levels for the audience/listener. Finally, we conclude with a set of issues that may act as a framework for future research focused on compositional strategies for larger scale, distributed, networked-based sensor environments.

1. INTRODUCTION

This paper aims to pose some critical aesthetic and technical questions concerning the design and implementation of compositional systems in what is increasingly referred to in the literature as responsive sound environments. Responsive sound environments encompass areas ranging from urban sound installations [1, 2, 3] to

Figure 1. The dancer, Michael Schumacher, during the performance of *Schwelle* in Montreal (photo by Anke Burger).

workaday models of auditory display [4]. We wish here, however, to focus on one small space of this potentially rich field: the design of such environments within a live performance context. The rubric of live performance not only describes familiar live performance forms such as dance, theater, instrumental or electronic music and similar genres, but also crosses into the arena of performative installations, environments and architecture. To limit our scope, we will focus on stage performance as the central application, presenting the case study of a recent large-scale interactive dance theater performance entitled *Schwelle* which premiered in Berlin at the Transmediale Festival for Art and Digital Culture in February 2007 and was subsequently revised for presentation in Montreal at Place des Arts/Cinquieme Salle series in May 2007.

Although there is no agreed upon use of the term, interactive or responsive sound environments generally refer to a system which “regenerates a soundscape dynamically by mapping ‘known’ gestures to influence diffusion and spatialization of sound objects created from evolving data” [5]. Outside of a musical context, the original use of the term derives from the work of computer scientist Myron

Kruger in the 1970s. Kruger more broadly defined responsive environment as a “physical space in which a computer perceives the actions of those who enter and responds intelligently through complex visual and auditory displays” [6]. Since a central component of such environments is the use of sensors or capture devices to gather data from the physical world and relay this to a computer in real time, technical and artistic research in the area of responsive sound environments has been mainly explored within the domain of human-computer interaction (HCI) with a musical slant. Consequently, if, as Kruger stated, the intelligent response of the computer occurs at the level of *auditory* display, we would then assume that questions about compositional output would be a key focus within musical HCI circles. A review of the literature, however, finds little to support this statement, with the vast majority of work in music oriented HCI focused on the development of either the control/sensing *device* itself (the “interface” or “instrument”) or the creation of mapping strategies for synthesis algorithms that “sonify” data derived from such devices [7].

We would like to argue that neither the current HCI focus nor traditional “algorithmic” computer music approaches to playback based composition sufficiently address the complexities of composing or sound designing for large scale, multi-input, sensor-based environments. This compositional challenge becomes all the more apparent when one deals with the difficult tasks of real time sonic manifestation of complex data, balancing legibility (listener perceivability of understanding what is mapped to what) with musical richness and complexity (timbral and rhythmic variation, dynamics, density, clustering) within the context of live performance where long term time evolution and duration are paramount to spectator/listener experience. Instead, our aim here is to develop a position and framework that exploits the tensions among mapping, sonification and composition for sensor-based responsive audio environments, suggesting that potential strategies for composition cross three different areas: mapping within algorithmic composition, sonification, and time-based evolutionary processes emerging from dynamical systems theory.

We first describe the areas of sensing, mapping and sonification as they are currently being explored in the HCI/New Instruments for Musical Expression (NIME) context and the tensions therein. Secondly, we describe work on a large-scale dance-theater project *Schwelle* as a case study for live performance-based responsive sound environments that couples compositional strategies with the use of computationally based dynamical systems fed by multiple streams of sensor data in order to generate what Roads calls *macro* and *meso* sonic structures over time [8]. Finally, we conclude with a set of questions that may act as a framework for future research focused on compositional strategies for larger scale, distributed, networked-based sensor environments.

2. SENSING AND INSTRUMENTS

The vast majority of literature in the area of sensor-based control and manipulation of real time audio has arisen within the NIME perspective. This is not surprising given that the New Interface for Musical Expression conference emerged in 2001 as a breakout group from the SIGCHI constituency to focus more exclusively on issues of musical import. The central focus of this approach lies specifically in the area of designing new (usually) sensor augmented devices that elicit more nuanced and expressive forms of human-computer interaction than traditional interface devices like keyboards or screens. NIME as well has explored issues as diverse as “simultaneous multiparametric control, timing and rhythm and training” [9]. This approach cuts across both areas of research and teaching. Stanford University’s CCRMA HCI courses such as Human Computer Interaction Theory and Practice: Designing New Devices and HCI Performance Systems: Music Controller Design and Development primarily emphasize the chains of Gesture→ Sensor→ Sound and Human→ Sensor→ Microcontroller→ PC→ Loudspeaker [10]. Similarly, the central goal of NYU’s ITP course in New Interfaces for Musical Expression is the formalization of a general set of design issues for musical expression that were articulated in the 2001 NIME workshop [11].

Central to the NIME tenet is that the sensor-augmented device or *controller* can be seen as a new kind of instrument. While this instrument model varies depending on the context (recent controllers include everything from cell phones to giant metal stretched strings [12, 13], the general NIME orthodoxy is based on a model of musical expression that involves the following characteristics: (1) real time sensor input for the real time control of musical parameters, (2) techniques for the conditioning, analysis and feature extraction of sensor signals and (3) mapping strategies for creating relationships between input and output parameters. In many ways, these three steps directly align with the gesture→ sensor→ sound = musical expression model that is at the core of the standard NIME approach, most specifically through interaction with a sensor-augmented musical device by way of gestural interaction.

3. MAPPING

Along with sensing-device modalities, the holy grail of HCI approaches to music is mapping. Mapping involves “the liason or correspondence between control parameters (derived from performer actions) and sound synthesis parameters” [14]. The techniques by which this has been accomplished have been extensively described in the literature [15, 16, 17] and thus, we will not rehearse them here. What is important to note is that mapping operates over both instrumental and algorithmic contexts. While instrumental mapping focuses most specifically on the mapping of real time input data (i.e., sensor information) to sound synthesis parameters, as defined above, algorithmic mapping suggests that the mapping techniques themselves are

entangled within higher level structural processes (computational, mathematical, etc); “the mapping of gestures to sounds may be considered the composition itself” [14]. According to Doornbusch [18] “mapping in algorithmic composition is different from mapping in instrumental design because composition is a process of planning and instruments are for real time music production.” The process of algorithmic composition thus, operates over multiple time scales, from what Curtis Roads called the *meso* scale of individual “sequences, combinations and transmutations” that are generated by sound synthesis to the *macro* scale, which involves the larger temporal-structural form and organization of musical or sonic events [8].

In relation to compositional processes, however, mapping is not without its detractors. In a NIME keynote from 2001, Joel Chadabe acknowledges a disconnect, or discrete relationship between the various structures of an electronic musical instrument versus an acoustic instrument [19]. While he acknowledges that as “instruments become more complex to include large amounts of data, context sensitivity and music as well as sound-generating capabilities, the concept of mapping becomes more abstract and does describe the more complex realities of electronic instruments”, Chadabe still focuses on an instrument model as the fundamental organizational mode for composition with electronic systems. A perhaps more germane criticism comes from computer scientist/artist Marc Downie who claims that mapping has become a catch all phase whose meaning has become “vague almost past the point of usefulness” and whose “predictive and explanatory power has long left us” [20]. Coming from the area of synthesis character generation with AI agent-based techniques, Downie’s critique stems from the argument that with the fast “binding” of “sensor” to “output”, “mapping deflates the awesome power of the algorithmic before it can appear.” In a sense, mapping flattens out the richness of computational complexity, confusing control parameters with the internal structure of a particular process. The task in developing richer interactive possibilities is to seek out new kinds of environments that involve unsupervised types of machine learning in order to “train” mappings and “induce them out of interactions” rather than create *a priori arbitrary relationships between input and output*. Chadabe and Downie’s arguments suggest that we clearly need new ways of thinking about the complexities of interactive musical systems that move beyond instrumental paradigms as well as the controller-based mapping approaches that are the core of the NIME literature and instead, towards algorithmic systems that deal with real time data in a complex, time evolutionary manner at both meso and macro levels.

4. SONIFICATION AND COMPOSITION

Another avenue that seems to suggest techniques for dealing with responsive environment compositional questions with complex, real time data is sonification. Sonification is primarily described as “the mapping of numeri-

cally represented relations in some domain under study to relations in an acoustic domain for the purpose of interpreting, understanding or communicating relations in the domain under study” [21]. Due to its scientific origins and context, sonification dispenses with an instrument model altogether, focusing instead on the auditory display of complex, real time data sets, and, in particular, providing a perceptual framework for dealing with abstract, non-representational (ie., non visual) data. Only recently, structural studies have started to explore which mapping techniques are suitable for certain types of data streams [22].

Although sonification uses sound as the technique and medium for representing data, the vast majority of the literature emphasizes non-musical applications in areas such as the representation of complex phenomena (chaotic or self-similar non linear systems) [23] auditory icons and audio interfaces [24] and non-speech audio representations of physiological data, in addition to other applications. What is clear is that while sonification moves away from instrument-control models, it does so at potential (or considerable) aesthetic expense. In an unpublished text, Larry Polansky, however, makes a useful distinction between sonification for scientific versus artistic purposes. Polansky notes that the closest sonification comes to composition is when mathematical or physical processes assist in generating new musical forms that are time varying. Sonification is not the same as algorithmic composition, although both techniques might utilize similar mathematical or physical processes. Using the example of computing π through a stochastic algorithm, Polansky posits two models. The first aims at using two pitches in order to sonify the statistical procedure, revealing how close or far we are from the number based on perceived consonance or dissonance. This technique aims at illustrating or representing a mathematical process in order to “elucidate” or reveal information about the unfolding of that process over time. The second example, however, uses mathematical processes (in the example, probabilistic techniques) in order to inspire new kinds of time varying compositional structures; not to “hear the Gaussian distribution as much as we want to use the Gaussian distribution to allow us to hear new music”. In this sense, mathematical processes become more of a manifestation or embodiment of the process. Characteristics and qualities that are essential to reveal information used in sonification such as “clarity”, “efficiency” and “economy” may be only of secondary or of no interest within a n aesthetic context [25].

5. COMPOSITION FOR RESPONSIVE SOUND ENVIRONMENTS IN LIVE PERFORMANCE

Formulating compositional strategies for responsive sound environments within a live theater or dance context presents a particular set of technical and perceptual challenges. First, unlike more open ended presentation contexts like walk-through sound installations or site specific environments [26], theater or dance performance (not to mention musi-

cal works) takes place over a defined and specified time duration. Furthermore, the audience/listener is present for this overall duration. Because of this, the perception of macro-structures (i.e., structural patterns, repetitions, sequences, motifs, themes and variations) becomes an essential element in the audience's understanding of how the performance evolves over time.

Second, as the audience is not directly involved in the manipulation of a set of controls, instruments or interfaces, the mapping model of direct gestural interaction to sound (or other forms of output) has become the default technique in establishing the effect of feedback between the human performer and the computer. Not surprisingly due to its emphasis on gestural control, the NIME model of sensing-mapping has also migrated to interactive performance contexts as most of these models also involve control of musical or visual parameters through sensor-augmented gesture and movement of performers. Due to the nature of what arts researcher Scott de La Hunta calls the "invisibility" and "transparency" of mapping various forms of input to output in a performance situation, "performer movement/action is used to trigger some sort of event (sonic, visual, robotic, etc.) in the space around or in some proximity to the performer" [27].

This direct, 1:1 feedback model many times results in either reducing the performer's range of bodily movements (leading to semaphore type of gestures) to activate sensors and synthesis structures to "visualize" the hidden mappings for the audience/listener or, more importantly, as Downie described it, deflating and simplifying more potentially interesting algorithmic compositional structures because such mappings might not be *immediately perceivable* by audience members and thus, not fulfill the feedback-based expectation that interaction seems to suggest. The consequences of focusing almost exclusively on the short term level of mapping-sound generation-feedback is that most interactive performance works fail to exploit the thing that makes musical or dramatic performance compelling and affective to the listener: the unfolding of an event over a durational time frame.

While few, there are examples of live performance events which have used multiple live feeds of sensor data for the control of both macro and meso compositional structures. Cage's *Variations V* and *Variations VII* both attempted to utilize multiple live feeds of data (in *Variations V*, capacitance and photoelectric sensing and in *Variations VII*, photoelectric and live telephone voice feeds) [28]. Other works such as Tod Machover's 1999 Brain Opera have also relied on simultaneous sensor based systems for the control of musical output. In the case of Cage, Machover or other similar interactive music models, however, the sensor interaction caused triggering of already pre-composed structures that would be altered within a finite set of boundaries as determined by the composer/author rather than through a set of more unpredictable processes generated by "unsupervised machine learning." While there has been work in moving away from the arena of triggering pre-composed sequences and towards the use of live

sensor data to drive longer term, time based processes [29] this work also relies principally on a controller/gesture-based model of the individual body as instrument and not the complexities of an environment.

It is our contention that live performance involving sensor augmented responsiveness presents a particularly robust context in which to research new techniques for computer assisted composition in that explicit attention must be paid to the organization and structuring of sound over longer time scales. The relationship between the structural components must be considered from both compositional and dramaturgical angles. Issues of mapping, therefore, must be subsumed into algorithmic processes which in turn are organized by even higher level narrative or dramaturgical structures. At the same time, one must explore the emerging levels of improvisation that occur between live performers, the environment (understood through sensing techniques), audience and computationally assisted structures.

The compositional challenge is one of using algorithmic complexity to generate potentially interesting and compelling patterns that function both at long term time scale while also manifesting interesting behaviors as a direct result of current inputs of sensor data at the meso scale. In this sense, Doornbusch is incorrect in stating that composition concerns preplanned structures while instruments focus on real time music production. Within a live performance context working with continuous data feeds from performers, objects, the controlled stage environment and potentially the audience's movement or behavior itself, we encounter both planned and simultaneously, real time structures.

6. SCHWELLE: A CASE STUDY

Schwelle is an evening length theatrical event which explores the varying threshold states of consciousness that confront human beings in everyday life experience, such as the onset of sleep or the moments before physical death. The three-act project has been in development between August 2005-January 2007 in Berlin, Montreal and Amsterdam and has involved a number of cultural and academic institutions. The project had its premiere in February 2007 at the Transmediale Festival For Art and Digital Culture in Berlin followed by performances in Montreal in May. Future performances will follow in Shanghai and elsewhere in Europe and North America during 2008 and 2009.

Part II of the project consists of a theatrical performance that takes place between a solo dancer/actor and a responsive room. The exerted force of the performer's movements are captured by wireless accelerometers located on the two arms and chest. Additionally, rhythmic changes in the color temperatures and dimming curves of lighting are picked up by strips of photoelectric cells and wirelessly transmitted via RF to a central server.

The continuously generated data from both the performer and environment is then used to influence the time

Figure 2. Michael Schumacher in *Schwelle* in Berlin (photo by Thomas Spier).

evolutionary behavior of a dynamically changing composition/sound design that attempts to give the impression of a living, breathing room for the spectator. The responsive sound environment that is a key element of the *Schwelle* performance consists of a multi-channel surround auditory environment whose sonic behaviour is determined continuously over different time scales, depending on the current input, past input and the internal state of the system generated by performer and environment in partnership with one another.

The conceptual and technical details for this work have already been addressed elsewhere [30] and thus, we will focus on issues of sound design and composition within the constraints of a live performance context. First, the notion of responsiveness in the context of *Schwelle* as a theatrical event signifies two things: (1) the room is responsive in the sense that the environment provides information that is subsequently manifested (but not illustrated) through specific sonic structures having distinct timbral and rhythmic relations that behave differently depending on different “states” of the room and (2) the time evolution of this continually generating sound design has a certain set of identifiable patterns that emerge over the duration of the performance. Thus, from the start the tension exists between musical structures that continuously evolve based on what is happening in the environment and, at the same time, the dramaturgical need to give those musical structures a clear sense of pattern for the audience.

The challenge working within the aesthetic rather than informational/display framework of data sonification involves developing a compositional infrastructure that both responds and is responsive to the behavior of the performer and room in co-production with one another. The compositional approach to *Schwelle* thus included the division of musical relationships that were either to be associated with human influence or with the behavior of the room ambience.

Figure 3. *Schwelle* System Diagram.

6.1. *Schwelle*: Compositional System

The room composition for Part II of *Schwelle* is organized and built around 16 layers of sound structures generated within the SuperCollider3 programming environment [31] with the following qualitative characteristics (in parentheses, the parameters which can be modulated of these sounds):

- Continuous background noise (frequency, amplitude, modulation speed (‘activity’))
- Clouded events (density, frequency range, amplitude range, duration, amplitude)
- Regular, discrete events (frequency, tempo, amplitude, duration)

Faced with the complexity of creating a sonic environment that was identifiable to the audience/listener as a “character” (i.e., where sound would serve a dramatic function), we arrived at the underlying sonic language by building individual instruments in SuperCollider3 and listening to each layer of sound separately in order to identify qualitative changes. These sounds and their resulting musical parameters such as frequency, duration, rhythm, etc., were developed based on questions that emerged from *Schwelle*’s overarching dramatic structure: what kind of range of affective or emotional behaviors would the room environment exhibit over the course of the performance and how could such behaviors, “attributes” or states be manifested to the audience/listener through qualitative phenomena like density and thickness, legato-like continuous structures versus transient “events”, clustered versus particulate sounds and smooth versus jittery textures, among others. Based on this principle, no external “musical” structures in the sense of background music or music that would attempt to reinforce a particular mood or heighten dramatic tension were utilized.

In choosing the types of underlying meso structures, we took into account how such qualitative changes were coupled with psychoacoustic issues such as frequency and amplitude masking, critical bandwidth, relationships between sounds, etc. For example, one of the first steps

taken was to listen to each layer of the sound in order to determine what the limits of the amplitude and the frequency occurrence should be. Other SuperCollider instruments were designed utilizing specific synthesis techniques that would consequently generate particular kinds of aesthetic effects in the audience/listener (e.g., subsonic or sounds drifting above and below perceivable frequency thresholds, higher transients that would produce interesting aliasing effects, larger, Xenakis-influenced sound masses whose pitch wavering was stochastically determined).

In searching for computational models that would yield potentially interesting patterns with the existing meso sonic structures, we turned to work in the area of dynamical systems (i.e. mathematical models of dynamic processes, based on a set of differential equations that describe the dynamic behaviour of the system [32]. It is important to make this distinction, as the label “dynamic system” is often misused in other work (e.g. [2] where systems based on cellular automata and fractals are labeled as dynamic systems. The aim of the dynamical system deployed was to create a time varying model where the compositional structure would depend on the multiple feeds of real time sensor data input from the room and the performer. The sensor data is gathered, statistically analyzed and dynamically scaled as described in the previously cited paper [30]. This data is then fed into the dynamical system which is based on early ideas of J.F. Herbart [33], who developed a theory for the strength of ideas, due to sensory impressions, as a function of time and the strength of the impression. It is this metaphor, which we also apply in *Schwelle*: the room’s physical behaviour makes certain sonic ideas stronger within the room over a longer time scale. Each type of sensor input has a separate scaling factor to determine the amount of influence it has on the system. Within the multiple layers of sound, each layer consists of one single sound object (object in the sense of a sound that stands on its own), whose intensity (amplitude) and amount of occurrence is dependent on how strong the “idea” within the system is at a certain moment in time. The Herbart system functions in such a way that there is no direct one-to-one mapping, but that its output also depends on the past of the system. This makes sense musically, as it reflects how we perceive and understand the temporal evolution of music¹.

6.2. Schwelle: Evaluation and Reception

Gathered from audience response ascertained during the initial Berlin performances of *Schwelle*, however, while the resulting behavior of the system produced a continuously time evolving behavior with the individual sonic layers, the Herbart system alone did not suffice to achieve structures and patterns which were clearly distinguishable enough from each other. Audience members surveyed noted that the dynamic range of the overall sound environment sounded in fact “too static” oddly enough because

of continuous change. Based on the audience feedback and our intent for the system to generate more perceivable compositional patterns at the macro scale, we have now implemented a state-based system on top of the existing Herbart model that more clearly influences the patterning of the meso sound structures over time. Specifically, we have more clearly defined five qualitative sound states that align with different emotional states that take place in the room during the performance, such as “meditation”, “repressed anger”, “transient anger”, “sleep/dream” and “agitation.” For each of these states, we have designed the parameter limits in which each sound layer would shift to reflect that state. Some sound layers are more present in some layers than in others, or have different frequency ranges, durations and so on within certain layers.

The composition then moves through a statespace, which is implemented as a 3 dimensional geometrical space. The current moment gets assigned a set of coordinates, and the output of the Herbart system determines how it moves through the space with each output of the Herbart system connected to a translation vector. Formally this can be written as:

$$\tilde{\mathbf{p}}_n = \tilde{\mathbf{p}}_{n-1} + \alpha \tilde{\mathbf{h}} \mathbf{V} \quad (1)$$

where $\tilde{\mathbf{p}}_n$ is the position in the statespace at moment n , α is a speed factor, $\tilde{\mathbf{h}}$ is the vector with the output from the Herbart system, and $\mathbf{V}$ is a matrix containing the translation vectors.

In the space a sphere is defined with a certain radius (R), and if the current state moves out of the sphere ($|\tilde{\mathbf{p}}_n| > R$), it depends on where it crosses the border, as on the sphere different regions are defined as transitions to other states in the macro composition. Consequently, we can map the Herbart system outputs that contribute to a certain layer to vectors for moving through the space, so that on the way to a certain state, already more of that state is present within the piece; the “idea” for that sound is already getting stronger before moving into that state where the idea reach its maximum intensity. Thus, the aim is to have two kinds of compositional systems operating in conjunction with one another: the dynamical Herbart system which influences the time evolutionary behavior of the meso sound structures (what types of sounds are strongest and when) and the state system, which controls the larger macro clustering and repetition of the meso structures as the performance evolves. Based on subsequent conversations with audience members in attendance at the May 2007 Montreal performances, it was clear that the addition of the state system substantially contributed to a much more perceivable sense of compositional pattern as well as a clearer experience of the room’s sonic behavior. Audience members stated that they sensed that the room’s behavior was not “random” but instead felt unpredictable yet, “logical.” The fact that audience members experienced a much more “legible” composition but could not explain exactly why points to the successful marrying of both the dynamical system as well as the discrete state engine.

¹ in fact, Herbart applied his theory also to the understanding of “Tonlehre” [34, 35]

7. CONCLUSION AND FUTURE WORK

We have described some positions on research into the development of more robust compositional models for real time, sensor driven responsive sound environments in the context of live stage performance. While the existing NIME research in gestural, controller-based mapping to sound is valuable from both technical and artistic standpoints (e.g., choosing appropriate devices, conditioning and analysis techniques for real time data) we feel it does not sufficiently address more complex compositional questions of how one designs and composes sound for sensor-driven responsive environments over longer time evolutions. This paper has attempted to articulate some of the issues and pose a concrete example of ongoing artistic work that attempts to address them. It also acts as a position statement for new research starting up with the two principle authors and Marcelo Wanderley at McGill University's Input Devices and Musical Interaction Laboratory which will focus on the development and use of large scale, distributed wireless sensor networks for live stage performance contexts in which issues of sound design for potentially n-dimensional, ad-hoc sensor networks will be addressed.

8. REFERENCES

- [1] M. d'Inverno, J. Eacott, H. Lorstad, and F. Oloffson. The intelligent street, responsive sound environments for social interaction. In *ACM SIGCHI International Conference on Advances in Computer Entertainment Technology, ACE 2004, published ACM*, 2004.
- [2] Eduardo R. Miranda, K. McAlpine, and S. Hoggar. Dynamical systems and applications to music composition: A research report. In *Proceedings of Journees d'Informatique Musicale (JIM97), Lyon, France*. French Society for Musical Informatics (SFIM), 1997.
- [3] Son-o-house. <http://www.noxarch.com>.
- [4] E. Mynatt, M. Back, R. Want, M. Baer, and J. Ellis. Designing audio aura. In *ACM SIGCHI Conference on Human Factors in Computing Systems, 1998, published ACM*, 1998.
- [5] D. Livingstone and E.R. Miranda. Composition for ubiquitous responsive environments. In *Proceedings of International Computer Music Conference (ICMC2004), Miami, USA*. International Computer Music Association - ICMA, 2004.
- [6] M. Kruger. Responsive environments. In *Proceedings of the American Federation of Information Processing Societies, USA., 46: 423-429*. AFIPS, 1977.
- [7] NIME. <http://www.nime.org>.
- [8] Curtis Roads. *Microsound*. MIT Press, 2001.
- [9] Nicola Orio, Norbert Schnell, and Marcelo Wanderley. Input devices for musical expression: Borrowing tools from HCI. In *Workshop (NIME-01) during ACM CHI'01 - Seattle, USA, April 2001*, April 2001.
- [10] Michael Gurevich, Bill Verplank, and Scott Wilson. Physical interaction design for music. In *Proceedings of International Computer Music Conference (ICMC2003), Singapore*, 2003.
- [11] Gideon D'Arcangelo. Creating a context for musical innovation: A nime curriculum. In *Proceedings of the 2002 Conference on New Instruments for Musical Expression (NIME-02), Dublin, Ireland*, 2002.
- [12] Global string. <http://www.sensorband.com/atau/globalstring/>.
- [13] Greg Schiemer and Mark Havryliv. Pocket gamelan: Tuneable trajectories for flying sources. In *Proceedings of the 2006 Conference on New Instruments for Musical Expression (NIME-06), Dublin, Ireland*, 2006.
- [14] Eduardo Miranda and Marcelo Wanderley. *New Digital Musical Instruments: Control and Interaction Beyond the Keyboard*. A-R Editions, 2006.
- [15] A. Hunt and M.M. Wanderley. Mapping performer parameters to synthesis engines. *Organised Sound, Cambridge: Cambridge University Press*, pages 97–108, 2002.
- [16] A. Hunt, M.M. Wanderley, and M. Paradiso. The importance of mapping in musical instrument design. *Journal of New Music Research*, pages 429–440, 2003.
- [17] A. Hunt and R. Kirk. *Trends in Gestural Control of Music*, chapter “Mapping Strategies for Musical Performance”. IRCAM-Centre Pompidou, 2000.
- [18] P. Doornbusch. Composers' views on mapping in algorithmic composition. *Organised Sound, Cambridge: Cambridge University Press*, pages 145–156, 2002.
- [19] J. Chadabe. The limitations of mapping as a structural descriptive in electronic musical instruments. In *Proceedings of the 2002 Conference on New Instruments for Musical Expression (NIME-02), Dublin, Ireland*, 2002.
- [20] M. Downie. Choreographing the extended agent. <http://www.openendedgroup.com/ideas/ideas/thesis/marcdphd.zip>, 2005.
- [21] C. Scaletti. *Auditory Display: Sonification, Audification, and Auditory Interfaces*, Ed. Gregory Kramer, chapter “Sound synthesis algorithms for auditory data representations”. Addison-Wesley, 1994.

- [22] Alberto De Campo. Toward a sonification design space map. In *Proceedings of the ICAD 07-13th International Conference on Auditory Display, Montreal, Canada, June 26-29, 2007*.
- [23] R. Bargar. *Auditory Display: Sonification, Audification, and Auditory Interfaces*, Ed. Gregory Kramer, chapter "Pattern and Reference in Auditory Display". Addison-Wesley, 1994.
- [24] W. Gaver. *Auditory Display: Sonification, Audification, and Auditory Interfaces*, Ed. Gregory Kramer, chapter "Using and Creating Auditory Icons". Addison-Wesley, 1994.
- [25] L. Polansky. Manifestation and sonification. <http://eamusic.dartmouth.edu/~larry/sonification.html>, 2002.
- [26] Garth Paine. Reeds - a responsive sound installation. In *Proceedings of the ICAD 04-10th Meeting of the International Conference on Auditory Display, Sydney, Australia, 2004*.
- [27] Invisibility/corporeality. http://www.noemalab.org/sections/ideas/ideas_articles/delahunta.html.
- [28] Variations VII. <http://www.medienkunstnetz.de/works/variations-vii/>.
- [29] J. Ryan and C. Salter. Tgarden: Wearable instruments and augmented physicality. In *Proceedings of the 2003 Conference on New Instruments for Musical Expression (NIME-03), Montreal, CA, 2003*.
- [30] M. Baalman, D. Grigsby, and C. Salter. Schwelle: Sensor augmented adaptive sound design for live theatrical performance. In *Proceedings of the 2007 Conference on New Instruments for Musical Expression (NIME-07), New York, NY, 2007*.
- [31] J. McCartney. Supercollider. <http://www.audiosynth.com>, <http://supercollider.sourceforge.net>.
- [32] Gene F. Franklin, J. David Powell, and Abbas Emami-Naeini. *Feedback Control of Dynamic Systems*. Addison-Wesley Publishing Company, 3rd edition edition, 1994.
- [33] Johann Friedrich Herbart. *Kleinere Abhandlungen*, chapter "De Attentionis Mensura causisque primariis" (orig. published 1822). E.J. Bonset, Amsterdam, 1969.
- [34] Johann Friedrich Herbart. *Kleinere Abhandlungen*, chapter "Psychologische Bemerkungen zur Tonlehre" (orig. published 1811). E.J. Bonset, Amsterdam, 1969.
- [35] Johann Friedrich Herbart. *Kleinere Abhandlungen*, chapter "Psychologische Untersuchungen. Erstes Heft" (orig. published 1839). E.J. Bonset, Amsterdam, 1969.