

Interfaces for Musical Expression Based on Simulated Physical Models

Andrew Johnston

University of Technology, Sydney

Submitted to the Faculty of Engineering and Information Technology in Partial
Fulfilment of the Requirements for the Degree of Doctor of Philosophy

2009

Certificate of Authorship / Originality

I certify that the work in this thesis has not previously been submitted for a degree nor has it been submitted as part of requirements for a degree. I also certify that the thesis has been written by me. Any help that I have received in my research work and the preparation of the thesis itself has been acknowledged. In addition, I certify that all information sources and literature used are indicated in the thesis.

Andrew Johnston
October 2009

Contents

Contents	i
Acknowledgements	xi
Abstract	xiii
1 Introduction	1
1.1 Introduction	2
1.2 Background	2
1.3 What is a Virtual Instrument?	3
1.4 Musical Expression and Exploration	4
1.5 Significance of the Research	6
1.6 Structure of the Thesis	7
2 Literature Review	9
2.1 Introduction	10
2.2 Design Criteria	10
2.2.1 Ease of Use and Expressive Potential	10
2.2.2 Intimacy and transparency	12
2.2.3 Predictability and Consistency	16
2.2.4 Diversity	17
2.2.5 Personality	19
2.2.6 Notes in time vs. timbre / Micro vs. Macro level operations	20
2.3 Mapping	21
2.3.1 Limitations of Mapping	23
2.3.2 Instrumental Gestures	24
2.4 Applications of Physical Modelling Techniques in Music	26
2.4.1 Physical models in sound synthesis	26
2.4.2 Physical Models in Musical Interfaces	28
2.5 Interactions with Virtual Instruments	31
2.6 Audio-visual works for acoustic instrument & computer	32
2.6.1 Messa di Voce	32
2.6.2 The Singing Tree	34

2.6.3	The Metasaxophone	36
2.6.4	Hyperinstruments	39
2.7	Conclusion	41
3	Methodology	43
3.1	Introduction	44
3.2	Research Paradigm	45
3.3	Research Approach	47
3.3.1	Action Research	47
3.3.2	Design Science	50
3.3.3	Design Approaches	52
3.4	Research Methods	56
3.4.1	Identifying design criteria	56
3.4.2	User Studies	57
3.4.3	Grounded theory method	65
3.5	Conclusion	70
4	Virtual Musical Instruments	71
4.1	Introduction	72
4.2	Design process	72
4.3	Design criteria	73
4.4	Overview of the virtual instruments	76
4.4.1	Controlling the virtual instruments	76
4.4.2	Physical model as mapping layer	77
4.5	Partial Reflections I	81
4.5.1	Physical model	81
4.5.2	Mappings	82
4.5.3	Aesthetic Perspective	84
4.6	Partial Reflections II	86
4.6.1	Physical model	86
4.6.2	Mappings	86
4.6.3	Aesthetic Perspective	89
4.7	Spheres of Influence	91
4.7.1	Physical model	91
4.7.2	Mappings	93
4.7.3	Aesthetic Perspective	94
4.8	Discussion	95
4.8.1	Coupling	95
4.8.2	Continuous versus discrete excitation gestures	96
4.8.3	Sound	96
4.8.4	Effect of timbre change	97
4.9	Technical details	98
4.10	Conclusion	98

5	Interactions with the Virtual Instruments	99
5.1	Introduction	100
5.2	Method Summary	101
5.2.1	Participants	101
5.2.2	Data Collection	103
5.2.3	Qualitative Data Analysis	105
5.2.4	Analysis Example	106
5.3	Findings: Interactions with the Virtual Instruments	112
5.3.1	Modes of Interaction	112
5.3.2	Focus of Attention	122
5.3.3	Aspects of Musician - Virtual Instrument Interaction	128
5.4	Findings: Evaluating the Virtual Instruments	142
5.4.1	Observer's notes	142
5.4.2	Questionnaire Findings	149
5.4.3	Audience Interviews	155
5.5	Conclusion	157
6	Framework for Criteria-Based Design	161
6.1	Introduction	162
6.2	Framework for Criteria-Based Design	162
6.2.1	Instrumental interaction	163
6.2.2	Ornamental Interaction	165
6.2.3	Conversational Interaction	166
6.2.4	All Modes of Interaction	167
6.3	Conclusion	168
7	Implications, Ongoing and Future Work	171
7.1	Introduction	172
7.2	Taxonomies of Interaction with Virtual Instruments	172
7.2.1	Rowe's Taxonomy	173
7.2.2	Winkler's Performance Models	175
7.2.3	Chadabe's Continuum	176
7.2.4	Modes of Compositional Engagement	178
7.3	Physical Models and Musical Interfaces	179
7.3.1	Mapping Between Acoustic Sounds and Physical Forces	179
7.3.2	Using Simulated Physical Models to Control Audio Synthesis	179
7.3.3	Other Audio Synthesis Methods	180
7.3.4	Instrumental Control and Technical Feedback	182
7.4	Partial Reflections III: Exploring Conversational Interaction	182
7.4.1	Responding to Two Performers Simultaneously	183
7.4.2	A More Complex Physical Model	183
7.4.3	Altering Physical Model Structure During Performance	185
7.5	Conclusion	187

8 Conclusion	189
References	195
A Memos	205
B Coding Scheme	225
C Episode Report	231
D Interview with Ben Marks	257
E Guidelines for Observers	265
F Consent Forms	269
G Questionnaire	273
H Transcript of Audience Study	275
I DVD Track Listing	279
J Software CD	281
K Publications	283

List of Figures

2.1	Graph illustrating degree of mapping transparency for player and audience (Fels et al. 2002, p. 116). Expressive instruments will be transparent to both audience and player - positioned in the top right section of the graph.	13
2.2	Block diagram of a virtual instrument using a dynamic independent visual-mapping layer (Momeni and Henry 2006, p.50).	30
2.3	A performance of <i>Messa di Voce</i> (the <i>Bounce</i> module). Here the vocalist Jaap Blonk is making cheek-flapping sounds which cause a stream of bubbles to be shown on screen as if emerging from his head. (<i>Reproduced with permission.</i>)	33
2.4	A still from the <i>Fluid</i> module of <i>Messa di Voce</i> showing a glowing fluid or plasma emerging from the performers' mouths as they vocalise. (<i>Reproduced with permission.</i>)	34
2.5	Photograph of the <i>Singing Tree</i> at the Haus der Musik, Vienna. (Photograph by the author.)	35
2.6	Diagram showing mappings from the Matasaxophone sensors to parameters of ExBow physical models (Burtner and Serafin 2002, p. 138)	38
2.7	Diagram showing the relationship between live audio and acoustic and sensor data streams from the Metasaxophone in the piece <i>S-TranceS</i> . Note that although not shown here, the bowed string sounds are also convolved with one another, which adds further to the sonic complexity of the piece.	39
3.1	Overview of the research project. Processes are shown as clouds, concrete outcomes as boxes.	45
3.2	Crotty's four elements of research (Crotty 1998, p. 5).	46
3.3	Action research cycle.	48
3.4	Overview of the design science approach to Information Systems research (Hevner et al. 2004, p.80). While the emphasis is on 'business needs', the paradigm is of course also applicable in other domains. . .	52
4.1	Block diagram of a virtual instrument using a dynamic independent visual-mapping layer (Momeni and Henry 2006, p.50).	77

4.2	Mapping from the volume of the source sound to force acting on the simulated physical model.	80
4.3	The mass-spring model for <i>Partial Reflections I</i> . Circles represent masses and connecting lines represent springs. The smaller, shaded sphere's position is fixed, all others are mobile masses.	82
4.4	Ben Marks plays the virtual instrument for <i>Partial Reflections I</i> . Note the bright glow of some of the spheres indicating that they are currently having force exerted upon them. (Photo: Ros Hodgekiss)	85
4.5	Ben Marks plays the virtual instrument for <i>Partial Reflections I</i> . (Photo: Ros Hodgekiss)	85
4.6	The mass-spring model for <i>Partial Reflections 2</i> . Circles represent masses and connecting lines represent springs. The smaller, shaded sphere's position is fixed, all others are mobile masses.	87
4.7	Photograph of performer and <i>Partial Reflections II</i> virtual instrument just before playing commences. (Photo: Ros Hodgekiss)	87
4.8	Photograph of Ben Marks performing with <i>Partial Reflections II</i> virtual instrument. Note that due to the slow camera shutter speed, the photograph shows six masses on screen. Actually, to the naked eye, one mass was orbiting around the central point when this photo was taken. (Photo: Ros Hodgekiss)	90
4.9	Photograph of Ben Marks performing with <i>Partial Reflections II</i> virtual instrument. Here, two masses are orbiting the fixed central point. (Photo: Ros Hodgekiss)	90
4.10	Screenshot of <i>Spheres of Influence</i> in its resting state.	92
4.11	Diagram of the physical model for the <i>Spheres of Influence</i> virtual instrument. Shaded circles are fixed masses, non-shaded circles are mobile masses and lines represent springs. The fixed masses around the outside are closest to the viewer, the fixed mass in the centre is in the far distance. The mobile masses (non-fixed) are positioned in between the outer fixed masses and the fixed central mass. None of the fixed masses are visible on-screen.	92
4.12	Photograph of Ben Marks performing with the <i>Spheres of Influence</i> virtual instrument. (Note: This photograph was taken before the mapping between the harmonic structure of the acoustic sound and the colour of on-screen spheres was enabled.) (Photo: Ros Hodgekiss)	94
4.13	Screenshot from the <i>Spheres of Influence</i> virtual instrument, illustrating how different visual colours are produced in response to acoustic sounds with different harmonic structures.	95
5.1	Demographic information collected from the musicians who participated in the study.	102

5.2	Whiteboard sketch showing the modes of interaction. Note that 'ornamental' mode was still characterised as an 'effect' at this stage of the analysis.	110
5.3	An arrangement of computer punch-cards labelled with modes of interaction and virtual instrument characteristics. These were used to help clarify and make coherent the emerging theory.	111
5.4	Another arrangement of punch-cards, showing a number of desirable and undesirable virtual instrument characteristics as described by the musicians in the study.	111
5.5	Three modes of interaction mark boundary points of a map of interactions with a virtual instrument.	117
5.6	Virtual instruments which support conversational interaction facilitate a shifting balance of power between musician and virtual instrument. .	136
5.7	Sample page of notes taken by an observer present at the sessions. .	143
5.8	Summary of findings by the observers of the musicians' interactions with <i>Partial Reflections I</i> . '1' indicates the observer found clear evidence the criteria was met; '0' indicates that the observer did not find evidence the criteria was met; '1?' indicates that evidence was found but was incomplete or unconvincing (these are not counted in the totals).	145
5.9	Summary of findings by the observers of the musicians' interactions with <i>Partial Reflections II</i> . '1' indicates the observer found clear evidence the criteria was met; '0' indicates that the observer did not find evidence the criteria was met; '1?' indicates that evidence was found but was incomplete or unconvincing (these are not counted in the totals).	146
5.10	Summary of findings by the observers of the musicians' interactions with <i>Spheres of Influence</i> . '1' indicates the observer found clear evidence the criteria was met; '0' indicates that the observer did not find evidence the criteria was met; '1?' indicates that evidence was found but was incomplete or unconvincing (these are not counted in the totals).	147
5.11	List of questionnaire responses. Numbers indicate degree of concurrence with statements (1 = strongly disagree, 2 = disagree, 3 = neutral, 4 = agree, 5 = strongly agree).	153
5.12	Mean values and standard deviations for each question.	154
7.1	Overview of scanned synthesis. The performer manipulates a dynamic system (usually a mass-spring model) to change the shape of the audio waveform, thus directly affecting the timbre of generated sounds (Verplank et al. 2000, p. 368).	181
7.2	The physical model for PR3 was made up of 48 masses arranged in a circle.	184

7.3	Screenshot showing effect on the physical model when a middle C is sounded on an acoustic instrument. As with the other virtual instruments, the mass which was currently having force exerted upon it was also made to glow in order to aid transparency of operation.	185
7.4	During performance the structure of the physical model was altered. This screenshot shows the model after a number of links have been cut and the tension in some springs relaxed.	186
8.1	Overview of the research project. Processes are shown as clouds, concrete outcomes as boxes.	190

List of Tables

2.1	Summary of the taxonomy of effective instrumental gestures (Cadoz 1988, Cadoz and Wanderley 2000). These terms will be used in chapter 4 to describe the mapping between performers' acoustic sounds and the resulting behaviour of the virtual instruments.	25
4.1	Mass-spring model simulation parameters for <i>Partial Reflections I</i> . Note that the values shown here are unit-less, as the choice of units has no impact on the simulation. Simulation parameters are proportional in that a force value of 10 applied to a mass of 100 will have the same effect as a force of 1 applied to a mass of 10, provided the simulation rate is unchanged (Henry 2006). Thus while it is possible to conceive the mass as being in kilograms and the link length in meters if desired, this has no intrinsic meaning in the simulation.	82
4.2	Summary of mappings from live audio to physical model in <i>Partial Reflections I</i>	83
4.3	Summary of mappings from physical model behaviour to audio synthesis in <i>Partial Reflections I</i>	84
4.4	Mass-spring model simulation parameters for <i>Partial Reflections II</i> . . .	86
4.5	Summary of mappings from live audio to physical model in <i>Partial Reflections I</i>	88
4.6	Summary of mappings from physical model behaviour to audio synthesis in <i>Partial Reflections II</i>	89
4.7	Mass-spring model simulation parameters for <i>Spheres of Influence</i> . . .	92
4.8	Summary of mappings from live audio to physical model in <i>Spheres of Influence</i>	93
4.9	Summary of mappings from physical model behaviour to audio synthesis in <i>Spheres of Influence</i>	93
6.1	A framework of design criteria informed by practice and research. . . .	163
8.1	A framework of design criteria informed by practice and research. . . .	193

Acknowledgements

I am indebted to a great number of people who supported and encouraged me during this project. To my supervisors, Linda Candy and Ernest Edmonds, thank you for giving me the freedom to explore and the intellectual and moral support to ensure that I did, eventually, produce a thesis.

To my colleagues Michael, Jim, Laurel and Bruce, thank you for the many lunches during which we discussed family, philosophy, food, travel and general gossip. Thanks for the comradeship, the ideas and the laughs.

To all the researchers at the Creativity and Cognition Studios, past and present, thanks for providing such a vibrant, creative and interesting environment in which to work. There is nothing more rewarding than hanging around with smart people who are passionate about the things that really matter.

Special thanks to Ros for being such a good friend and a willing 'guinea pig' in the early stages.

To all the musicians who participated in the study, thank you for being so willing to give up your valuable time and for providing such insightful comments and feedback. Without your help this project would not have been possible.

Ben Marks – mate your contribution has been enormous. Thanks for the inspiration and support, for being so open to new ideas and for your energy and generosity of spirit.

I have been incredibly fortunate to have been surrounded by inspiring musicians all my life. Foremost amongst them has been my father. Dad, thanks for the music – and so many other things. Mum, for the love, unstinting support and for being such a lovely Nana I will always be grateful.

Finally, to my beautiful girls: Emma, Zoe and Mia – you always put a smile on my face. Thank you for supporting, enduring and inspiring this work. Love always.

Abstract

This thesis is concerned with the design of interactive virtual musical instruments intended to augment acoustic instruments in live performance. For the purposes of this work, a virtual musical instrument is defined as a computer system designed to facilitate musical expression and/or exploration. The aim of the research is to develop understanding of the nature of virtual instruments and how musicians interact with them. The approach has been to use participatory design techniques to develop a series of virtual instruments for use in live performance and then to examine closely the experiences of musicians who use them.

An interaction design strategy which uses simulated physical models to mediate between the sounds produced by acoustic instruments and computer generated sounds and visuals has been developed. In this approach, a simple physical system is modelled in software and characteristics of acoustic sounds are mapped to forces and other parameters which affect the model. In response the model moves in ways that are physically realistic. These movements are then used as parameters to control video and audio synthesis.

Using a research approach which draws on action research, design science and participatory design, a series of virtual instruments which use this interaction technique were developed and used in live performances. A set of initial design criteria which guided development were identified. In order to refine these criteria and better understand the impact that using these virtual instruments has on musicians' music-making, a series of user studies were conducted. A number of expert musicians used the virtual instruments and discussed their experiences. These sessions were video-recorded, transcribed and analysed using grounded theory techniques.

The results of the study identified three modes of interaction with the virtual instruments: instrumental, conversational and ornamental. Musicians interacting with the virtual instruments in instrumental mode emphasise the importance of being in control and being able to trust that the instrument will respond consistently. When musicians use a virtual instrument ornamentally, they surrender detailed control of the generated sound and visuals to the computer, allowing it to create audio-visual layers that are added to the musicians' sound. The more complex, (and difficult to design for) conversational interaction involves the sharing of control between the musician and the virtual instrument. The balance of power is in flux, allowing the virtual instrument to talk back to the musician, reflecting and transforming the sonic input in ways that move the performance in new musical directions.

The contributions of this thesis are therefore:-

- a set of virtual musical instruments which use a unique interaction paradigm in which simulated physical models mediate between live sounds produced on acoustic instruments and computer generated sounds and visuals;
- a theory of musician-virtual instrument interaction; and
- a set of design criteria informed by practice and user studies.

Chapter 1

Introduction

1.1 Introduction

This thesis is concerned with the design and use of interactive virtual musical instruments which respond to, augment and ‘talk back’ to musicians playing acoustic instruments. The virtual musical instruments are ‘virtual’ because they do not exist physically: they are not made of brass, wood or other material substance, but instead are computer programs which respond to acoustic sounds via a microphone and generate their own sounds and visuals.

The aim of the research is to develop understanding of the nature of virtual instruments and how musicians interact with them. The approach has been to use participatory design techniques to develop a series of virtual instruments for use in live performance and then to examine closely the experiences of musicians who use them. The end results are:

1. The development of a series of virtual instruments which use a unique interaction paradigm;
2. A theory of musician-virtual instrument interaction; and
3. A set of virtual instrument design criteria informed by 1 and 2 above.

1.2 Background

The work described in this thesis came about because I have been drawn to both music and computing throughout my life. Musically, I have played the bass trombone, having completed a music degree and played professionally with a number of groups including symphony orchestras, big bands, early music ensembles and theatre bands. Computationally, I’ve completed a Masters degree in Information Technology and currently hold the position of lecturer at the University of Technology Sydney, where I predominantly teach software development and agile methods.

For some time I resisted using computers to a significant degree in my music practice because I felt, perhaps wrongly, that they could not provide anything approaching the kind of nuanced expressivity that traditional instruments afford. Computer music seemed to me to have a predominantly ‘mechanical’ or mathematical aesthetic, not because that was what was necessarily desired, but because computers produced these kinds of sounds relatively easily.

For this reason, my first inclination was to develop software to enhance the development of skills on traditional instruments. That is, software which might provide feedback to a musician playing an acoustic instrument in order that they might discover something about their sound and/or technique which they were unaware of previously. However, as the project progressed I began to realise that increasing hardware and software capabilities had made the expressive use of computers in

live performance situations far more viable than I had naively supposed. In particular, I felt that Cyrille Henry's creative use of simulated physical models (Henry 2004a, Momeni and Henry 2006) demonstrated how digital technologies could be used live to create expressive music and visuals with a distinctly non-digital aesthetic. This inspired me to shift the focus away from software for teaching and learning towards developing virtual musical instruments for use in live performance.

Having made this shift, I was fortunate to have a close relationship with composer/trombonist Ben Marks who was interested in composing music for the emerging virtual instruments as they developed. Ben's participation in this project was invaluable. Together we worked to develop the performance work *Partial Reflections*, which was presented at several concerts during the course of the project. Our collaboration gave life to the virtual instruments and is the creative core of this thesis.

Ben and I were pleased with the artistic results of our collaboration, but I was intrigued by a number of questions. How did using the virtual instruments alter the way musicians thought about their music-making? How did they conceive of the virtual instruments while they were interacting with them? What characteristics of the virtual instruments enhanced their experience and which detracted from it? How could they be made more effective? Attempting to address these questions led to a close examination of the experiences of a number of highly experienced musicians as they used the virtual instruments. These user studies complement the creative work and provide insight into the impact of software of this kind on musical creativity.

1.3 What is a Virtual Instrument?

In this thesis I will use the terms *virtual musical instrument* or, more simply, *virtual instrument* to describe a computer system designed to facilitate musical expression and/or exploration that responds to the actions of a performer by producing sounds.¹ While not included in my deliberately broad definition, it should be noted that many virtual instruments (including those created during this research project) also produce some kind of visual display which is linked to the actions of the performer and/or the generated audio to a greater or lesser degree.²

Because of the lack of a necessary direct cause and effect relationship between performers' actions and the sounds generated by a virtual instrument, the line between composition and instrument can become blurred. A single performer action can trigger multiple musical events - a single button press might trigger playback of many notes for example. Where the user interface is highly constrained and the musical output is complex we might, in fact, consider the 'instrument' to be a composition. To take an extreme example, an instrument with one button which, when pressed,

¹This is very close to the definition proposed by Jordà: "Musical instruments are used to play and produce music, transforming the actions of one or more performers into sound." (Jordà 2004, p. 321)

²Goto includes this visual component in his definition of the term which for him, "...refers to a whole system which contains Gesture, Gestural Interface, Mapping Interface, algorithm, Sound Synthesis, and Interactive Video." (Goto 2000, p.217)

triggers playback of twenty minutes of complex music must really be considered a composition. As we add more buttons and other interaction possibilities which give the performer more influence over the generated music at a greater level of detail, we might start to see the device as a ‘composed instrument’ (Schnell and Battier 2002) and the method of design/performance as ‘interactive composing’ (Chadabe 1984) or a kind of real-time arranging. Here, the composer/instrument designer could be devising processes (generative algorithms, etc) and perhaps pre-composed musical fragments which are triggered and affected by the actions of the performer/composer. As more fine-grained control is given over to the performer and the size of the musical events triggered by individual performer actions is reduced, the virtual instruments begin to more closely resemble traditional acoustic instruments, working within a “one-gesture-to-one-acoustic-event paradigm” (Wessel and Wright 2002, p. 11).

In this thesis, the emphasis is on virtual instruments at the ‘instrumental’ end of this spectrum - instruments which give the performer a degree of control over sounds at a low level - and I will use the term ‘virtual instruments’ to describe these. However, the virtual instruments I will describe in this dissertation are not restricted to ‘merely’ instrumental responses. The term ‘instrument’ implies a particular kind of interaction - one in which the performer uses the instrument as a tool to achieve a pre-planned goal. The musicians who used the software I developed did not necessarily use the software in this way, and nor did they want to. I considered using ‘interactive music systems’ (Rowe 1993) but to me this is an overly broad term which would include digital recording and editing software, sequencers and all manner of other software which I am not concerned with here. So, despite their imperfections, I will use the terms ‘virtual musical instruments’ or ‘virtual instruments’ in this thesis. However, I stress that I use them in the broadest possible sense.

I will use the term ‘interactive composition’ to describe systems where the influence of the performer is comparatively minor and the computer produces responses to performer actions that are longer in duration, perhaps by playing back pre-composed phrases or complex sequences of sounds controlled by generative processes.

1.4 Musical Expression and Exploration

Given that I am taking the view that a virtual musical instrument is designed to facilitate musical expression and/or exploration, it is necessary to consider just what is meant by these terms.

Musical expression is a complex and potentially controversial term. As a starting point, consider this definition from the Oxford English Dictionary:

“The manner of performance (with respect, e.g. to degrees of loudness or softness) suited to bring out the feeling of a musical passage.” (OED Online 1989)

This is perhaps the traditional use of the word in music performance. When performing a piece of notated music a performer subtly alters characteristics of their sound such as dynamic, rhythm and pitch in order to bring the work to life. These variations are called 'expression' and are commonly seen as something added to the bare bones of the music (notes and underlying rhythm). An early goal of some researchers was to develop techniques for enabling computers to play notated music in ways that sounded less mechanistic and more human. They sought to do this by uncovering the 'rules' which human performers applied to create expressive performances (eg. Clarke 1985, Canazza et al. 1997).

However, such a definition of expression is restrictive in many ways and does not accurately reflect contemporary music practice. Of particular concern is the fact that this conception of expression assumes there is a pre-existing score which is interpreted by a performer. Clearly this is not always the case.

Gurevich and Treviño (2007) identify in the literature a common conception of expression which involves performers somehow coding extra-musical elements into the music they produce. These extra-musical elements are usually considered to be 'emotion' and the implication is that a perfect musical interface will enable the error free 'transmission' of emotions from performer to listener. As Gurevich and Treviño point out though, expression in contemporary music practice is a complex and ephemeral concept which resists such simplistic characterisation. Performers, including a number who participated in this research, do not always see the purpose of music performance as expressing a pre-existing idea or state.³ At times the process of music-making is a kind of exploratory dialogue with their instrument, other musicians and the broader environment within which the performance takes place. Just what is being expressed is very difficult to pin down.

Of course, some musicians very explicitly reject the notion of expression in their practice. The composer John Cage, to take a prominent example, made strenuous efforts to try to remove the sense of personal intentionality from his music-making.

"I have found a variety of ways of making music (and I continue to look for others) in which sounds are free of a theory as to their relationships. I do not hear music before making it, my purpose being to hear as beautiful something I have not before heard." John Cage, cited in Holmes (2002, p. 238)

For him, everyday, unpredictable sounds were far more interesting than the artfully constructed compositions traditionally presented in concert halls, and his works aim to allow the listener to simply experience sound without attempting to interpret it, decode it or even relate it to past experience. From this perspective the notion of

³Each of the participants in the study participate in a wide range of musical activities, so their approach to expression is likely to vary depending on context. When playing romantic repertoire in a symphony orchestra they will use a different approach than when they are improvising in a contemporary music ensemble, for example.

personal ‘musical expression’ is something to be avoided, and certainly not designed for.

Given this broad range of approaches to musical expression, in this thesis I take the pragmatic view that a virtual instrument is ‘expressive’ when musicians are able to use it successfully for their particular musical purposes.⁴ The approach I have taken is to work with musicians who are highly trained, experienced and prominent in their field. Thus participants in the user study include principal players from professional symphony orchestras as well as prominent improvisers and contemporary music specialists. As will be described in detail later, musicians actively participated in the design of the virtual instruments and in the user study which examined how well or badly they facilitated expression and exploration.

Because of this, the work described in this thesis could at least partially be seen as an investigation into the nature of musical expression, as the musicians involved approach it. Through this work I have come to the realisation that the interfaces for musical expression which I have designed straddle the boundary between a deterministic approach to music-making (the ‘traditional’ approach for want of a better word) and the non-deterministic approaches of experimental composers such as Cage. Viewed from a traditional perspective they are virtual *instruments* which permit instrumental control. From a non-deterministic perspective they are sonic (as opposed to musical) *environments* which simply produce sounds. How musicians respond to this environment is up to them.

One response to an unfamiliar environment is to explore it. The notion of musical *exploration* has different connotations to that of musical *expression*. Certainly the two concepts are related but when exploring a sonic environment the interaction is less directed towards specific ends. Keeping open the possibility that the virtual instruments might facilitate exploration rather than expression (or may support both at different times) means that musicians whose personal aesthetic might be more aligned with Cage’s may still find them of value.

1.5 Significance of the Research

Increasing computer power has made the development of advanced audio-visual music applications possible even on low-cost hardware. This, coupled with the emergence of programming environments such as Pure Data, Max/MSP and many others has led to a large number of interactive music applications being developed. The continuing growth of the New Interfaces for Musical Expression conference, run annually since 2001, is evidence of the level of interest in this field.

⁴This aligns well with the somewhat recursive definition of expression offered by Fels: “musical expression occurs when a player intentionally expresses herself through the medium of sound”. A well-designed instrument in Fels’ terms, “supports the ability to play music by allowing the user enough control freedom to explore sound space and make music while being sufficiently constrained to allow the user to learn to play the instrument” (Fels 2004, p. 672).

The rapid growth in this area provides many opportunities for creative work and research. As I will discuss in detail later, a key issue for the design of interactive virtual musical instruments has been ‘mapping’ between performer actions and software response. In this thesis I describe a novel technique for mapping between live sound, generated by musicians playing acoustic instruments, and computer generated sounds and visuals. I show that this technique can facilitate the development of virtual instruments which are flexible enough to support a wide range of musical approaches.

The development of virtual instruments should be an artistic process,⁵ but one which is stimulated and inspired by technical developments. Because virtual instruments are intended to facilitate creative work, a purely technical approach is inadequate. Technical issues such as mapping are important only because they impact on the *experiences* of musicians who use virtual instruments and the music they make. Examining the nature of these experiences in a musical context is therefore a critical area.

This thesis presents findings from a study of expert musicians who used the virtual instruments that were developed. The study used qualitative research techniques with their origins in sociology to derive a theory of musician - virtual instrument interaction which accounts for the experiences of the musicians who participated. Thus this thesis does not consider technical issues in isolation, but examines them in the context of musical experience.

In short, this thesis is significant because it addresses two key issues in this area:

- the design of virtual musical instruments, including the mapping between performer actions and virtual instrument response; and
- the experiences of musicians who use these virtual instruments.

1.6 Structure of the Thesis

Chapter 2 (Literature Review) begins by looking at the current state of the art in virtual instrument design. In this chapter I examine the literature and identify the key issues which have arisen. The important area of ‘mapping’, or linking user actions to the sounds and/or visuals produced by the virtual instrument is examined and the use of simulated physical models in music applications is introduced. The comparatively unexplored use of simulated physical models in musical *interfaces*, as opposed to sound synthesis, is introduced. Finally, a number of significant audio-visual works for acoustic instruments and computer are presented in order to provide some context for the virtual instruments I have created.

In chapter 3 (Methodology) the overall approach to this creative/research project is outlined. An approach to interpretive research which draws on action re-

⁵Perry Cook, articulating this position succinctly, advised designers of new musical instruments to, “make a piece, not an instrument or controller.” (Cook 2001, p.1)

search and design science is described. Details of the specific research methods, including participatory design and grounded theory techniques are also presented in this chapter.

In chapter 4 (Virtual Musical Instruments) I describe in detail the virtual musical instruments which were developed and identify the design criteria which guided their development. These criteria are drawn on in chapter 5 (Interactions with the Virtual Instruments), in which the musicians' experiences with the virtual instruments are examined.

Chapter 5 describes the development of a theory which accounts for the way in which the musicians in the user study used the virtual instruments. A key finding of the study was that the musicians interacted with the virtual instruments in three different 'modes of interaction': instrumental, conversational and ornamental. Each of these modes is described and illustrated with quotes from the study. In addition, data gathered during the study is analysed in order to *evaluate* how well or badly the virtual instruments met the design criteria identified in chapter 4.

The findings from chapter 5 led to a re-examination of the initial design criteria. The original design criteria which guided the development of the three virtual instruments were developed with a relatively simplistic view of musicians' interactions in mind. In chapter 6 (Framework for Criteria-Based Design), a more sophisticated set of criteria are presented which explicitly refer to the modes of interaction identified in chapter 5.

Chapter 7 (Implications, Ongoing and Future Work) considers the implications of the findings from the user studies and the development of the revised design criteria.

Finally, the contributions of the thesis are summarised in chapter 8 (Conclusion).

Chapter 2

Literature Review

2.1 Introduction

In this chapter I outline key issues concerning the design of virtual musical instruments. A number of areas of research are relevant here. First, there is the issue of design criteria for virtual instruments: that is, what are the high level design criteria that are likely to lead to successful instruments? Second, there is a body of work looking at the key area of “mapping” - linking the actions of a performer with sounds and/or visuals generated by the virtual instrument. Next, I consider the use of physical models as a mapping layer for virtual instruments. Physical models have been used extensively in sound synthesis but less commonly as a mediating layer between the acoustic and computer-generated sounds. This technique is a very important component of the software developed as part of this thesis as it can help address some common limitations of virtual instruments. Finally, a selection of audio-visual virtual instruments which influenced the design of the virtual instruments I created are discussed.

2.2 Design Criteria

The question of how to design virtual instruments in order to facilitate musical expression is complicated. As discussed in chapter 1 (section 1.4), the term “expression” is itself somewhat nebulous and difficult to separate from musical aesthetics. However, there is a wide ranging literature which considers design criteria for computer-based musical instruments, and, in this section, I present some key ideas which informed the development of the virtual instruments developed during this project.

2.2.1 Ease of Use and Expressive Potential

Wessel and Wright argue that expressive instruments need to provide a “low ‘entry fee’ with no ceiling on virtuosity” (Wessel and Wright 2002, p.12). In other words, they need to provide ease of use without sacrificing the flexibility that will allow a sufficiently motivated performer to develop high-level skills. While this is a superficially straight-forward concept, there is some ambiguity over the need for ease of use and what actually constitutes a low ‘entry fee’ in this context. There are a number of traditional instruments which arguably have a high entry-fee but have nonetheless been very successful. It takes some time to develop the ability to play a single steady note on a violin for example, but the affordances provided by the physical design of the instrument, coupled with a rich history of previous virtuosic performances, provide sufficient motivation for new performers to continue to take up the instrument and work to develop high-level skills. Because there is a history of ‘great performers’ on the violin, we have evidence that great things are possible. Thus, the success of the instrument can not be decoupled from social context.

The concept of virtuosity is also complicated and resists quantification as there is no objective way to measure it. At its simplest, we say that a performance is virtuosic if people agree that it is. Great performers, by definition, redefine the meaning of what it means to be a virtuoso by raising the level of expected performance or making use of extended techniques, simultaneously impacting on instrumental technique and musical language.

What is generally meant by designing for virtuosity is that the instrument should provide the performer with a wide range of expressive possibilities. That is, it should not be overly restrictive or respond in ways that are too simple. However, while the instrument must be capable of producing a wide range of rich, complex sounds (eg. Wessel and Wright 2002, Jordà 2004, Fels 2004, Dobrian and Koppelman 2006), this should not come at the expense of transparency and repeatability. Paine articulates this trade-off neatly:

“The mappings¹ must be such that there is extensive scope for exploration and the discovery of new outcomes, but where the outcomes prove repeatable to the extent that they confirm the cognitive map that the interactor is developing as their relationship with the interactive system deepens.” (Paine 2002, p. 298)

Golan Levin aims to build audio-visual performance instruments which have an “instantly knowable, indefinitely masterable interface” (Levin 2000, p. 56). While building instruments which meet these criteria is difficult, he provides two examples from the physical world - the piano and the pencil - which demonstrate that they are not mutually exclusive.

“Although any four-year-old can discover their basic principles of operation, it is common for an adult to spend many years practicing these tools, and yet still feel that more mastery is possible or that more compositions remain to be expressed through them. Such systems, moreover, have the extraordinary property that an individual may eventually, through their use, discover or reveal a unique and personal voice in that medium.” (Levin 2000, p. 56)

Settel and Lippe (2003) argue that “instrumental expressive range” is critical. By this they mean that the instrument should be able to transmit the performer’s ‘musical message’ and be able to handle a wide range of such messages. Thus, ‘complex’ instruments such as the violin have greater potential for musical expressivity than simple ones such as claves for example. A complex instrument in this context is one that has “high resolution and wide range with respect to dynamics, timbre and pitch” and which has a high level of consistency; it “can consistently render a sonic response specific to a particular instrumental manipulation”. (Settel and Lippe 2003, p. 198)

¹We will consider mapping in detail in section 2.3.

Summary The concepts of ease of use and virtuosity are complex and difficult to separate from cultural context. However, there is a consistent theme in the literature that virtual instruments should be conceptually simple while providing musicians with the ability to create a wide range of complex sounds.

2.2.2 Intimacy and transparency

Fels argues that *intimacy* is key to successful instrument design. Intimacy, in this case, refers to the degree to which the performer feels in control of the instrument. Drawing on the arguments of Moore (1988), Fels argues that this intimacy (“control intimacy” in Moore’s terms) occurs when performers feel a close link between their physical actions and the sound produced by the instrument. If they perceive a close link between their actions and the resulting behaviour of the device then they can feel intimately connected to it. Expanding on this, Fels observes that there are four kinds of relationships that can form between people and objects (Fels 2000, 2004):

1. Response: The object is perceived as a separate entity which responds to control in some way. Interaction in this relationship can be seen as a kind of simple conversation in which the person makes gestures and observes how the system responds. The point here is that gesture and response are perceived as being separate from one another. The person can observe that the system responds to their actions but they are still discovering *how* it responds.
2. Control: The object is perceived as an extension of the body—it becomes ‘embodied’, a part of the person. In this relationship, action and response are coupled and the feeling is of controlling and/or performing with the system rather than responding to it as an independent entity.
3. Reflection: The object is perceived as a separate entity which does not respond to control. A patron of a traditional art gallery might have this kind of relationship with the exhibited paintings. The person may have a complex response to the painting, and communication could be said to occur, but the painting itself is unaffected by the interaction.
4. Belonging: The person perceives themselves as an extension of the object - they are embodied by the object. In other words, their actions are impacted upon, or perhaps even controlled by, the actions of the object. Fels states that, “for this type of relationship, the object must be able to control the person and the person must be in a state to allow the control” (Fels 2004, p.675) and notes that this relationship is more difficult to design for.

It is important to note that these relationships follow a continuum and can co-exist. Studies of the experiences of people with Fels’ work *Iamascope* (Costello et al. 2005), for example, indicate that they frequently move to and fro between response and control states during their interaction. In Fel’s terms, ‘intimacy’ refers to




Figure 2.1: Graph illustrating degree of mapping transparency for player and audience (Fels et al. 2002, p. 116). Expressive instruments will be transparent to both audience and player - positioned in the top right section of the graph.

the degree to which the person is able to embody an object and he argues that this is a critical factor in designing musical interfaces that facilitate musical expression. Further, he claims that building a strong response relationship is a necessary precondition to moving on to the intimate control relationship. Learning how the system works is necessary before one can develop expert skills. A good system will therefore facilitate and motivate the transition from the response to control relationships.

An instrument which facilitates intimacy should also exhibit transparency in operation. That is, the links between player action and the sounds produced must be clearly apparent and understandable for both performer and audience. Fels et al. (2002) use a simple graph to visualise the likely expressivity of individual instruments (fig. 2.1). One axis represents player transparency and the other audience transparency. Different interfaces for musical expression can be placed in the resulting quadrants. Ideal (most expressive) instruments will be in the top right hand corner of the graph indicating they are transparent to both audience and player. They argue that transparency is a predictor of expressivity. ie. A musical interface that is transparent to both audience and musician will be expressive, at least to some extent.

A key characteristic of computer music is that the mechanism that produces the sounds is decoupled from the physical actions of the performers. When playing a trombone, for example, the performer directly manipulates the instrument by blowing air, buzzing the lips and moving the slide. There is a direct physical connection between the performer and the resulting sounds. When the sounds are produced by a computer however, there is no need for the performer to physically excite an object in order to produce sound. Electric power provides the energy required to produce

sonic vibrations via loudspeakers

This provides us with a great deal of freedom. The computer is capable of producing an enormous range of sounds. We are no longer constrained by the physical structure of traditional musical instruments.

One consequence of the movement towards the use of general-purpose laptop computers in live performance is that there is a corresponding breakdown in the link that the audience perceives between performers' actions and the resulting sounds and/or images (Fels et al. 2002, Jordà 2002, Schnell and Battier 2002). Live performances which feature performers sitting at laptops creating and manipulating sounds are common. From the audience's perspective the physical actions of the performers are almost completely decoupled from the sound that is produced. For all they know, they might be checking their email while simply playing a CD or DVD.

The practice of 'live coding' is one technique that has emerged in order to enhance the audience's perception of the links between performer actions and computer generated audio visuals (Collins 2003). In live coding, performers' laptop screens are made visible to the audience, members of which are then able to discern, to some degree, how the actions of the live coders impact on the music.

Ruminating on this issue, Miller Puckette points out that the link between performer actions and computer output is significant not only for audience engagement but also in the building of performance traditions.

"To make [sharing musical ideas] possible, there must be a direct and comprehensible relationship between the controls we use and the sounds we hear. (This would not be a bad thing from the audience's point of view either.) A performer who pushes a button to start a sequence is not showing us how the music was really made; all we learn about the music is what our ears can tell us. But if the performer's actions correspond more closely to the sounds themselves, then we can see something about the music's gestural content, and our own music can be better informed by it. In this way we could evolve a stronger and more meaningful tradition for making computer music." (Puckette 1991)

Fels et al. (2002) argue that using mapping strategies that relate to what they term 'literature', meaning shared knowledge or common understanding, can help build and maintain transparency. Having the virtual instruments respond in physically plausible ways, a technique I have used in the design of the instruments described in this thesis, allows the performer and audience to draw on their practical knowledge of how physical objects respond to forces to understand the workings of the instrument and thus maintain transparency.

Feedback from the instrument to performer is important in maintaining transparency for the performer. Tanaka (2000) argues that feedback gives the performer confirmation that their actions are having the desired effect. The most obvious feedback channel is the audio produced by the instrument, but other channels can help

enhance the feeling of connection between player and instrument. In traditional instruments tactile feedback is provided because of the (usually) direct physical connection between performer and the vibrating sound source. In virtual instruments, this tactile feedback must be explicitly designed for. When it is not possible to provide tangible feedback, Tanaka argues that “the physical gesture of the performer must be apparent in some resulting musical gesture. This is important both to the musician to have a sense of what he is doing, and for the audience to have an understanding of what is taking place” (Tanaka 2000, p. 400). Later I will describe a technique which uses simulated physical models to mediate between performer actions and computer generated sounds and visuals. A benefit of this technique is that it can magnify performer gestures in a way that can enhance both their visibility and audibility.

Visual feedback can also be used to show workings of the system and thus facilitate audience understanding (Arfib et al. 2005, Dobrian and Koppelman 2006). This can occur at different levels. At a low level, it is possible to show a direct representation of the parameters that the performer is manipulating: “For example, measurements of pressure or blowing force² can be displayed in the form of a slider or the degree of illumination of a graphical object.” (Arfib et al. 2005, p. 127) At higher levels, a representation of the systems’ interpretation of these low level parameters and the resulting impact on the sound can be shown. For example, if it has been decided that a combination of increased blowing force and decreasing pressure will result in a darker sound, then the system might display a darker colour on screen.

Moore (1988) argues that “control intimacy” is critical, and that the best traditional musical instruments have a rich palette of sounds that are closely linked to “...the psycho-physiological capabilities of a practiced performer” (Moore 1988, p.21). Instruments that facilitate a high degree of control intimacy allow performers to use subtle physical gestures to manipulate musical sounds in expressive ways. In discussing the limitations of the MIDI protocol, Moore argues that many expressive acoustic instruments facilitate continuous variation of musical parameters. A trombonist, for example, is able to continually make subtle changes to timbre, volume and pitch while playing a single note. Such control comes at the expense of learnability to some extent. Developing proficiency on expressive musical instruments takes time.

Summary In order for musicians to develop an intimate, expressive connection with virtual instruments, the links between their actions and the responses of the virtual instrument need to be clearly apparent and understandable. These links should also be clear to audience members in order to foster engagement with the performer and the work. Using mapping strategies that relate in some way to shared knowledge or common understanding should help in this regard. Likewise, tangible and/or visual feedback can be used to enhance the transparency of a virtual instrument’s operation.

²Clearly, a virtual instrument using a wind controller is being referenced here.

2.2.3 Predictability and Consistency

As well as suggesting that expressive musical instruments should be easy to play initially but also provide motivated performers with the capacity to develop virtuosity, Wessel and Wright (2002) argue for a number of other criteria. They indirectly support the need for transparency and intimacy by arguing that musicians should have control over “the overall volume and density of the sound an instrument is producing, even when a semi-autonomous generative algorithm is handling the details” (Wessel and Wright 2002, p. 14). Further, from their experience, they have found that, “there should be some sort of correspondence between the “size” of a control gesture and the acoustic result” (Wessel and Wright 2002, p. 14). That is, large, energetic gestures by the performer should produce louder, more energetic sounds.

Predictability is another key requirement. Wessel and Wright argue that the performer should always retain high-level control over the virtual instrument, even though automatic processes may take care of many of the details. Predictability becomes complex when the output of the system depends on past actions of the performer as well as current ones. When the state of the instrument changes as a result of performer actions, then the same gesture may not have the same effect at different times. This, of course, does not mean that instruments should not change state; the behaviour of the piano is still predictable even though holding the sustain pedal down will change the effect that pressing the middle C key will have. Rather it suggests that the state of the instrument should be apparent to the performer. Wessel and Wright argue that some kind of visual display indicating the state of the system may be useful in this regard. They, however, prefer to design interfaces that do not accumulate information from multiple gestures to avoid this problem. Paine (2002) argues against this concession, arguing that for an instrument to be truly interactive and not merely responsive, its structure and mappings *must* change over time in response to the actions of the user. Just how much the fundamental structure of a virtual instrument may change before the user loses the sense of interacting with a coherent system is unclear however.³

While the terms ‘predictability’ and ‘consistency’ have subtly different meanings, they are closely coupled. We say that an instrument is predictable when it responds consistently,⁴ when it can, “consistently render a sonic response specific to a particular instrumental manipulation” (Settel and Lippe 2003, p. 198). Traditional acoustic instruments have high levels of predictability in the sense that if the performer executes two identical physical actions upon them, they will produce the same acoustic result. Of course, different instruments have different tolerances in this regard. To produce two identical notes on the violin requires considerably more skill than

³Paine proposes neural networks as a programming technique that might provide the system with sufficient ‘intelligence’ to be able to learn and evolve its own internal structure and hence its response to user actions. The idea is that this approach would allow the structure to change while retaining a degree of transparency for the user. This allows each interaction to be unique and respond to individual actions of the performer as well as the evolving history of the interaction.

⁴I will have more to say about this later (section 6.2.1).

producing two identical notes on the triangle. Instrumentalists develop the skills to produce physical movements that are ‘close enough’ to identical to have minimal perceptual effect on the notes produced.⁵

This points to an inherent trade-off in musical instrument design. We want fine-grained control in order to be able to produce subtle alterations in timbre, but the control should not be so fine-grained that the interaction becomes overly complex and it becomes impossible to repeat a desired effect on demand. Thus instrument designers need to find “a balance between randomness and determinism, between linear and non-linear behaviours” (Jordà 2004, p. 336). Chadabe describes this trade-off in the context of what he terms ‘interactive composing’:

“In interactive composing, the system’s response to the performer must be interesting and informative; that is, it must contain new, unexpected information. Yet the response must also be recognisably related to the performer’s actions, because if the performer did not perceive clearly the effects of his or her actions, the act of performing would cease to have any meaning. The response must be as in a conversation, where the reply to a statement is related to the statement but not a repetition of it.” (Chadabe 1984, p.25)

This concept of conversational interaction with virtual instruments is a key point and will be explored in detail in later chapters.

Summary Virtual instruments should behave consistently so that musicians are able to predict what effect their actions will have. Changing the state of the virtual instrument may lessen predictability. Thus, when instrument state is dynamic, its current state should be made apparent to performers and audience. Balancing controllability and complexity is a key issue and conceiving the interaction between performer and instrument as ‘conversation’ may be useful.

2.2.4 Diversity

Jordà (2004, 2005) argues that an instrument’s *diversity* is a critical characteristic and observes that this diversity can be considered at three levels: micro, mid and macro. *Micro-diversity* is a measure of how effectively an instrument allows the performer to vary two performances of the same piece. An instrument with high micro diversity will allow subtle changes of timbre, tempo, dynamic, etc. in order to allow the performer to make a particular performance of a piece unique. This desirable characteristic of expressive musical instruments is a measure of the extent to which a performer is able to produce nuanced performances (Arfib et al. 2005, Dobrian and Koppelman 2006).

⁵For an overview of the relationship between physical gestures and noticeable differences in generated audio from the instrument designer’s perspective, the reader is referred to Rubine and McAvinney (1990).

Mid-diversity is a measure of how different two pieces played on the same instrument will sound. If two different performers playing two different pieces of music sound essentially the same on a particular instrument, then that instrument has low mid-diversity. A difficulty with attempting to objectively measure mid diversity is the unknown impact of listeners' previous experience with the instrument and/or the genre. To many people, the works of Mozart, Beethoven and Brahms sound the same, but anyone with more than a cursory knowledge of music of that era will instantly be able to recognise and identify the work of each. It is, therefore, likely that an instrument which facilitates the exploration of new musical territory will initially be perceived as having low mid-diversity, but will be re-evaluated more positively as performers and listeners gain expertise with the new style. This is not to say that the concept is without merit, but we should keep in mind the contingent nature of any evaluations we make in regard to the mid-diversity of new instruments.

Finally, there is *macro-diversity* (or *adaptability* (Arfib et al. 2005)), a measure of how adaptable and flexible an instrument is. An instrument with high macro-diversity could be used in many ensembles and different musical styles. Lower macro-diversity indicates a more specialised instrument. Jordà observes that instruments with lower macro-diversity are perhaps more likely to have more personality and therefore have more chance to give birth to new styles, but higher macro-diversity instruments will have more chance of being used, because they will be more able to fit in with existing instruments. Once again, separating intrinsic characteristics of an instrument's design from social context is difficult here, as it is really only by observing how an instrument will, over time, be adapted for use (or not) in existing and new musical styles that we will be able to 'truly' evaluate its macro-diversity.

For this reason, Jorda positions macro-diversity as comparatively less important. Micro and mid-diversity on the other hand are critical in his opinion. A high level of micro-diversity is seen as facilitating the development of virtuosity by allowing the performer to develop their own voice and make each performance their own. High mid-diversity is important because it can motivate users to make the transition from superficial 'toying' with a novelty instrument to more engaged and expressive musical performance.

The importance of diversity will depend on the intentions of the virtual instrument designer. If he or she wishes to develop an instrument to be used in a wide range of performance contexts by many different musicians, then clearly diversity is a critical issue. However, one of the key advantages of software based musical instruments is that they can be easily modified: it is far easier to modify a virtual instrument to use a different synthesis method than it is to physically construct a new violin for example. It is, therefore, perfectly feasible for a composer to create new instruments for every musical work for example, or an improviser to create a new instrument for each performance. Thus, diversity is an area instrument designers need to consider but the degree to which it is prioritised is essentially an aesthetic decision.

Summary Ideally, virtual instruments should be flexible. Jordà calls this flexibility 'diversity' and describes three levels on which it can operate: micro, mid and macro. Micro-diversity refers to the ability of the virtual instrument to allow subtle changes to timbre, tempo, dynamic, etc in order to personalise an individual performance. Mid-diversity is a measure of how different two pieces played on the same instrument will sound, and macro-diversity is a measure of how adaptable an instrument is to different musical styles. The degree to which diversity is important depends on artistic context.

2.2.5 Personality

Virtual instruments should have a different character to traditional computer applications such as word processors, spreadsheets and databases. If musicians are to use them to create compelling music they need to do more than simply obey the commands of performers with maximum efficiency; they need also to motivate and inspire musicians to be creative. In the words of Sergi Jordà:

"Inspiring is an important concept. Music instruments are not only in charge of transmitting human expressiveness like passive channels. They are, with their feedback, responsible for provoking and instigating on the performer new ideas or feelings to express." (Jordà 2005, p.233)

Thus, while virtual instruments need to be controllable and predictable, we must keep in mind that they need to be more than "passive channels". As we will see in chapter 5, musicians in the user study often desired two-way, conversational interaction with virtual instruments. A perfectly controllable, predictable instrument may reduce the tendency for this kind of interaction to develop.

"A musical instruments *raison-d'être*, on the other hand, is not at all utilitarian. It is not meant to carry out a single defined task as a tool is. Instead, a musical instrument often changes context, withstanding changes of musical style played on it while maintaining its identity. A tool gets better as it attains perfection in realizing its tasks. The evolution of an instrument is less driven by practical concerns, and is motivated instead by the quality of sound the instrument produces. In this regard, it is not so necessary for an instrument to be perfect as much as it is important for it to display distinguishing characteristics, or "personality". What might be considered imperfections or limitations from the perspective of tool design often contribute to a "personality" of a musical instrument." (Tanaka 2000)

George Lewis' "interactive musical environment" *Voyager* (Lewis 2000), responds in real-time to live sounds by producing its own sonic response. Designed for use by Lewis himself, a virtuoso improvising trombonist, *Voyager* produces highly complex

and often unpredictable music. Lewis describes the importance of imbuing *Voyager* with its own sense of personality:

“Part of the task of constructing *Voyager* consisted of providing the program with its “own sound.” In *Voyager*, this notion of sound appears in tandem with a kind of technology-mediated animism, expressed as an interactive aesthetic of negotiation and independent computer agency.” (Lewis 2000, p.37)

For Lewis, then, issues of control over *Voyager* are of far less concern than the program exhibiting a kind of identity which fosters highly conversational free improvisation.

Summary Notions of efficiency and control in virtual instrument design are likely to be less important than giving them character and personality which motivates and inspires musicians. This may involve trading controllability for complexity of response.

2.2.6 Notes in time vs. timbre / Micro vs. Macro level operations

As I mentioned in chapter 1 (section 1.3), the lack of a necessary direct cause-and-effect link between performer actions and virtual instrument response means that the line between ‘instrument’ and ‘composition’ can become blurred. In that section I stated that the virtual instruments developed during this project operated at the instrumental end of this spectrum, in that they allow the performer detailed control over generated sounds at a relatively low level.

Of course, virtual instruments can be designed to allow performers to manipulate higher-level sounds such as phrases (groups of notes). Arfib et al describe two levels of manipulation: microscopic (“sound object” level) and macroscopic (phrasing level) (Arfib et al. 2005). As the level of musical material being manipulated increases virtual instruments start to become more like interactive compositions.

Paine (2002) proposes dynamic morphology (Wishart 1996) and spectromorphology (Smalley 1997) as ways to think about sound produced by interactive systems. Traditionally, music theorists have focused on notes and time, with timbre being relegated to orchestration. These approaches argue that timbre ought to be given a more prominent place in music theory. Paine proposes this thinking to be incorporated into computer music approaches as opposed to approaches which tend to focus on the organisation of notes in time at the expense of timbre.

Summary Virtual instruments can operate at different levels, allowing the performer detailed control over sound timbres and/or higher level musical structures such as phrases. The virtual instruments described in this thesis operate at the timbre (or ‘microscopic’) level.

2.3 Mapping

In the preceding section, I have examined the literature on virtual instrument design and highlighted some high-level design criteria and important issues to consider. A key area which impacts upon the character of virtual instruments is 'mapping' - linking users' gestures (of whatever kind) to sound. In this section, I examine mapping in detail.

A key characteristic of digital musical instruments is the decoupling of the gestural actions of the musician from the mechanism which actually produces sound. In traditional instruments there is a direct link between the actions of the performer and the resulting physical vibrations that constitute the instrument's sound. For example, the bowing action of the violin player is directly linked to the vibration of the strings and the resulting sound. In computer music, because the energy to produce sounds is provided electrically, there is not necessarily any direct physical link between the actions of the performer and the sounds produced. This means that a single synthesis technique may have any number of interfaces. A critical aspect of designing computer-based musical instruments is therefore the "mapping" between user interface and synthesis parameters.

Mapping in this context is the process of linking the interface to the sound source. A trivial example of a 'one-to-one' mapping might be to link the vertical movement of a mouse to the pitch of a generated note—the higher the user moves the mouse pointer, the higher the pitch. This is a one-to-one mapping because one parameter generated by the user (the height of the mouse pointer) is directly linked to one parameter of the synthesis engine that makes sounds: pitch. The designer of a very simple (but not necessarily very satisfying) virtual instrument might therefore map pitch to vertical movement of the mouse and volume to horizontal movement.

In addition to the simple one-to-one mapping, more sophisticated mappings are possible and probably desirable. 'One-to-many' mappings map one user-generated parameter to several synthesis parameters (Hunt et al. 2000). (This approach is sometimes described as a "divergent" mapping strategy (Rovan et al. 1997).) If we took our simple one-to-one mapping strategy instrument and extended it so that moving the mouse to the right simultaneously increased volume and decreased vibrato, then we would be using a one-to-many approach. One user-generated parameter (horizontal mouse position) is mapped to two synthesis parameters (volume and vibrato).

The converse is also possible: a many-to-one strategy in which many user generated parameters are mapped to one synthesis parameter. If we again redesigned our mouse controlled instrument so that moving the mouse vertically resulted in large changes in pitch, but moving it horizontally resulted in finer pitch changes then we would be using a many-to-one mapping. Rován et al. (1997) call this "convergent" mapping, and claim that although it is harder to master, it facilitates expressive playing more effectively than one-to-one mappings.

Mapping strategies have a significant impact on users' experiences with an electronic instrument. As Hunt et al. (2003) point out, "by altering the mapping...the entire character of the instrument is changed" (p. 429). Acoustic instruments, the designs of which have been refined over generations, have inbuilt complex, non-linear mappings between the actions of the musician and the resulting sounds. For example, a simple one-to-one mapping approach to the design of a synthesised clarinet which uses a Musical Instrument Digital Interface (MIDI) wind controller to capture the musicians' actions would probably map air speed to volume, with increasing air speed resulting in increased volume. The problem with this approach is that in a real clarinet the mapping is much more complex, with an increase in air flow also affecting the response of the reed which affects the timbre. Clarinetists therefore adjust their embouchure as they increase volume to maintain control over the sound (Rovan et al. 1997). The result is that the timbre of a clarinet played loudly is quite different to one played softly; a loud clarinet sound is not simply an amplification of a soft one.

In a series of three experiments, Hunt et al. (2003) demonstrate that complex mappings such as those built into acoustic instruments result in instruments that, while more difficult to learn initially, are more expressive, controllable (after the initial learning period) and enjoyable to play. Their studies involved creating a number of simple instruments which used various mapping strategies and evaluating them from two perspectives: how well the instruments enabled users to imitate a given sound, and, perhaps more interestingly, how expressive and satisfying they found the experience.

The simplest of these instruments used one-to-one mapping between the movements of basic sliders (both physical and on-screen) and synthesis parameters. Users generally found these uninteresting, limiting and frustrating to use and were less successful at imitating the given sounds with these instruments. In addition, as the sound the users were asked to imitate was made more complex, users did not show as much improvement with more practice as they did with the more sophisticated instruments.

The more complex instruments used many-to-many mappings and, in addition, required what might be termed 'energy injection' to make sound. Instruments that require energy injection need the player to excite the virtual instrument in some way. The theory is that when playing acoustic instruments, "the musician's energy is physically transformed into the sound energy that is heard by an audience" (Hunt and Wanderley 2002, p.97), and that this traditional physical connection between musician and music has important implications for the design of expressive instruments.⁶

In the study, players had to move a slider to excite the virtual instrument. If they stopped moving the slider, then the sound faded. According to the energy injection

⁶A notable exception to this rule is the church organ, which, like many modern electronic and computer instruments, decouples the 'controller' (keyboard) from the sound production system (bellows and pipes), and has nonetheless been popular and successful. This would seem to indicate that while the concept of energy injection and transformation may be important, it is not indispensable and may perhaps be traded off for other benefits such as polyphony and a wide variety of available sounds (Hunt and Wanderley 2002).

tion theory, players should feel more engaged by instruments that link their physical energy to the sounds produced. Results seem to support this, with users showing evidence of greater engagement with these instruments, and spending more time experimenting and exploring their possibilities. These interfaces were more rewarding and “felt more like a traditional instrument and promoted holistic thinking rather than forcing the users to analyse exactly what each of the controls did” (Hunt et al. 2003, p.433).

Summary Mapping between user actions and computer generated sounds is a critical issue for designers of virtual musical instruments. It has been argued that virtual instruments which use cross-coupled, many-to-many mappings and which require energy injection are more engaging and expressive. Some empirical evidence supports this argument.

2.3.1 Limitations of Mapping

Menzies (2002) argues that the concept of mapping implies a kind of stateless interaction, where the output from the sound generator depends only on the gestural input at that specific point in time. He argues for the use of a mapping process where the output depends in some way on the *history* of interactions. This means that the response of the virtual instrument to two identical inputs may change depending on what has happened previously. In other words, the mapping layer has a *state* that the user's actions alter while they play the instrument.

Chadabe (2002) observes that the response of virtual instruments can be placed on a continuum from completely deterministic at one extreme to completely non-deterministic at the other. A deterministic instrument is completely predictable and will always respond exactly the same way to any given input. A non-deterministic instrument on the other hand may use complicated algorithms incorporating chance operations. The output from such an instrument is highly unpredictable. Chadabe argues that instruments that incorporate a significant amount of non-determinism – while still giving the user some influence – stimulate a more conversational style of interaction. In the conversational interaction, the performer shares control of the performance with the virtual instrument. The virtual instrument is influenced, rather than controlled, by the performer, who, in responding to its sometimes surprising behaviour, may in turn alter the musical trajectory of their playing. Chadabe calls instruments such as these ‘interactive instruments’.

One technique for finding the fertile middle ground between deterministic and non-deterministic behaviour in the design of virtual instruments is to use some kind of physical dynamic system in the mapping layer (eg. Menzies 2002, Momeni and Henry 2006). This is discussed in section 2.4.2 below.

Summary It has been argued that conceiving of virtual instrument design purely as a mapping problem implies that two identical gestures will always result in the same sonic response. That is, that it encourages designers to conceive of virtual instruments as stateless ‘black boxes’ which always respond in the same way. Some authors argue that instruments which change state as they are played are more likely to stimulate a more conversational style of interaction.

2.3.2 Instrumental Gestures

In order to discuss the mapping strategies used in the virtual instruments created as part of this research, it will be helpful to clarify the terms used to describe the function of the various gestures⁷ used when playing musical instruments. The terminology introduced here will be used in chapter 4 (Virtual Musical Instruments) when describing the mapping between the live sounds produced on acoustic instruments and the computer generated sounds and visuals.

Players of traditional musical instruments physically manipulate them to produce sounds, using gestures made with various combinations of hand, mouth, feet, etc depending on the instrument. We focus here only on the gestures that are actually required to physically produce sound, the *effective* gestures (Cadoz and Wanderley 2000). There are of course other, possibly unconscious, gestures that musicians often make during performance such as swaying to the rhythm, tapping feet and so on. While these gestures are an integral part of performance for many instrumentalists, and are no doubt coupled with the minimal set of gestures that are strictly necessary, we can consider their direct impact on produced sound relatively trivial and will not consider them here.

Effective instrumental gestures have been grouped into three categories: excitation, modification and selection (Cadoz 1988, Cadoz and Wanderley 2000). Excitation gestures are those that are responsible for providing the necessary energy for the instrument to produce sound. A percussionist striking a drum is executing an excitation gesture for example, as is the violinist drawing their bow across a string and the clarinetist blowing into the mouthpiece.

Excitation gestures may be further grouped into two categories: *instantaneous* and *continuous*. Percussive and plucking style gestures (striking a drum, plucking a guitar string) are instantaneous: “sound begins when the gesture finishes” (Cadoz and Wanderley 2000, p.80). A continuous gesture on the other hand, continues as the instrument is sounding. For example, the blowing gesture of the trombonist continues as the trombone produces sounds – the gesture and sound co-exist.

The second category of instrumental gesture is the *modification* gesture, which modifies the instrument in some way but transfers only minimal energy to the produced sound. Examples include moving a trombone slide, pressing saxophone keys

⁷It should be born in mind that the term ‘gesture’ is used broadly in this thesis. The virtual instruments described in later chapters use acoustic sounds as instrumental gestures, for example.

and shortening a violin string by pressing down on the fingerboard. Modification gestures change the way the instrument responds to excitation gestures in some way.

Modification gestures may also be divided into two categories: *parametric* and *structural*. Parametric modification gestures alter a characteristic of the instrument on a continuum. A violinist, for example, can vary the length of a string from maximum (length of the open string) to minimum (distance from the end of the fingerboard to the bridge) by pressing it down onto the fingerboard. Parametric modification gestures can be continuous or discrete. The violinist uses a *continuous* parametric modification gesture when they use their finger to depress a string onto the fingerboard because the fingerboard is fretless, giving the performer theoretically infinite control over the length of the string between the minimum and maximum points. A guitar, on the other hand, is played with *discrete* parametric gestures because while the performer can press down at any point along the fretboard, the control over string length is limited by the placement of the frets.

Structural modification gestures on the other hand result in “categorical differences” (Cadoz and Wanderley 2000, p.80) in the structure of the instrument. Inserting a trumpet mute for example, changes the response and timbre of the instrument significantly and is an example of a structural modification gesture.

Finally, the instrumentalist uses *selection* gestures to choose between a number of similar elements of an instrument. The pianist, for example, has 88 keys to work with but only ten fingers. Selecting which keys to press allows them to decide which parts of the instrument are acted upon by excitation gestures – depressing keys.

This taxonomy of functional gestures is somewhat clumsy and, as Cadoz and Wanderley (2000) acknowledge, they are “neither exclusive nor independent”. However I use them here because they provide a useful language for discussing the relationship between the acoustic instrumental ‘gestures’ of live performers and the behaviour of the virtual instruments. Table 2.1 summarises these terms.

Functional gesture	Subtype(s)	Example
Excitation	Instantaneous	Striking a percussion instrument
Excitation	Continuous	Blowing a wind or brass instrument; bowing a violin
Modification	Parametric, continuous	Pressing violin string onto fingerboard; moving a trombone slide
Modification	Parametric, discrete	Pressing guitar string onto fretboard; pressing clarinet keys
Modification	Structural	Inserting trumpet mute
Selection	Sequential	Playing a monophonic synthesiser
Selection	Parallel	Playing piano

Table 2.1: Summary of the taxonomy of effective instrumental gestures (Cadoz 1988, Cadoz and Wanderley 2000). These terms will be used in chapter 4 to describe the mapping between performers’ acoustic sounds and the resulting behaviour of the virtual instruments.

Summary In this section I have presented a simple taxonomy of ‘instrumental gestures’, which performers use to create and control sounds. While imperfect, this taxonomy will be useful when describing virtual instruments in later chapters. It should be born in mind that the term ‘gesture’ is used broadly in this thesis to refer to any performer action which impacts upon the behaviour of a virtual instrument.

2.4 Applications of Physical Modelling Techniques in Music

In previous sections, we have discussed the importance of virtual instruments providing rich, complex sounds while retaining transparency and controllability. The issue of mapping between user actions and the behaviour of the virtual instruments has been discussed and an argument for using cross-coupled, many-to-many mappings which require the user to inject energy into the system to produce sounds has been presented. In addition, it has been suggested that allowing virtual instruments to change state during performance will facilitate more complex and interesting musical interactions. In this section, the use of simulated physical models as a mapping layer between user actions and computer generated sounds and visuals will be presented. I will argue that this technique has potential to provide intuitive control over complex musical material, and to effectively link sound and visuals.

2.4.1 Physical models in sound synthesis

Physical models have a long history in sound synthesis. Rather than attempting to directly simulate sound, in physical modelling synthesis the aim is instead to model the *source* of sounds. To put it simply, one might say that rather than building a virtual violin sound, the aim is to build a virtual violin. If we can model enough detail of the physical properties of the violin as an object, we can then ‘play’ it in the virtual environment and obtain realistic sounds. The higher the fidelity of our model, the more life-like the resulting sounds.

Cadoz et al. (1984) argue that traditional synthesis techniques do not lend themselves to satisfying musician-machine interactions, because they focus on the sound as an abstract object and use acoustical terms such as spectrum, modulation, amplitude, etc. to describe it. The job of a musician working within this paradigm is to specify the desired sound in these terms. The problem is that this way of describing sound separates its acoustic structure from its ‘symbolic content’ (Cadoz et al. 1984, p.60). Because of our experiences in the physical world, we have intuitive understanding of the relationships between physical gestures and sound. We know what crinkling cellophane or knocking on wood will sound like for example, because of our practical physical knowledge of how the world works. The parameters of traditional digital synthesis techniques do not map well to this kind of knowledge. A technique

to address this issue is to shift the focus away from the sound as an object by building instead virtual instruments that respond in physically plausible ways. That is, to focus on the cause rather than the effect (Cadoz 2002).

The musician using such an instrument can draw on their knowledge of how things work in the world to interact with the instrument. Rather than having to mentally map between their target sound and abstract synthesis parameters, they may use their intuitive understanding of physical processes and their links with sound to create and experiment with new sounds.

Physical modeling techniques have potential to create and control sounds that provide a higher degree of engagement for both performer and audience. Leman argues that there is evidence that “listening focuses on the moving source of a sound rather than on the sound itself” (Leman 2007, p.236). In other words, when we hear music, we perceive it in terms of physical actions that we associate with such sounds.⁸ These need not necessarily be the physical actions that actually cause the sounds, but actions that we somehow associate with them based on past experiences.

He proposes a model of musical communication based on the encoding and decoding of biomechanical energy in sound. In this model, the performer realises musical goals by physically manipulating an instrument (or, more generally, a ‘mediator’), which translates the performer’s physical energy into sound. The listener, at least partially through a process of associating sounds with physical actions, makes sense of the sound. This is not to say that the listener’s understanding of the music will be identical to that of the performer’s, but rather that the listener will make sense of the sound in their own action-related terms. The implication is that instruments which facilitate a more direct connection between the physical actions of performers and generated sounds are more likely to facilitate musical communication at this gestural level.

Castagne and Cadoz (2003) rationalise their use of physical models by arguing that they produce sounds which are *physically plausible*. That is, because of the simulated underlying physical mechanisms which generate the sounds, we are more able to attribute their source to natural processes, even though the sounds themselves may not sound ‘acoustic’ in the traditional sense. They argue that, “the important feature for a musical sound is not to cause the listener to infer its physical cause, but to present a set of subtle dynamic variations among perceptual parameters that lead the listener to think it was produced in some physical manner” (Castagne and Cadoz 2003, p.2).

Using force-feedback devices, Luciani et al. (1994) have made use of physical models to provide an immersive experience in which on-screen objects are given tangible presence. Using their device, the user can manipulate parts of an on-screen physical model which respond in the same way as actual physical objects. Devices

⁸This is perhaps due to human adaptation to evolutionary pressures associated with the need to identify dangers in our physical environment.

such as this offer exciting possibilities for the design of virtual musical instruments which retain a physical character while also providing freedom to musically explore psuedo-physical structures which would be impossible to actually build in the real world.

Physical models have been used both at the micro-level, to produce physically plausible sounds, and the macro-level, where they are used to initiate sounds and control timbral evolution over longer time scales. Cadoz (2002) describes the creation of a composition entirely realised using physical models. Low frequency (below 20Hz) models are used to articulate sounds that are produced by the high-frequency models (20Hz-20KHz). The piece itself is an elaborate mechanical system that is set in motion and left to run its course. That is, it is a non-real-time work and involves no interaction during the performance.

To date, the use of physical models in computer music has been predominantly aimed at directly synthesising sounds. In this thesis however, I focus instead on the use of physical models as *part of the interface between musician and sound and visual synthesis engine*. That is, using low-frequency models, manipulated by a musician, to set parameters for sound and visual synthesis. As such, a detailed examination of the use of physical models in direct sound synthesis is beyond the scope of this thesis. Interested readers are instead referred to Smith (2004) and Vlimki et al. (2006) for overviews of work in this area.

2.4.2 Physical Models in Musical Interfaces

As we have discussed, physical models offer intriguing possibilities for the creation and control of musical material with strong links to our everyday experience of the world. Putting the direct synthesis of sound aside, there is another way that physical models may be used in the design of virtual instruments - as a mechanism for mediating between user actions and the audio-visual response of the instrument. This is an area that is relatively unexplored.

In a seminal paper presented at the 1992 International Computer Music Conference, Joel Ryan argues that making musical ideas more concrete by using software models⁹ facilitates a more engaging, inventive approach to music making.

“Composers can devise models on the computer which give their ideas a more concrete form. But creating a model or simulation on the computer is more than just a representation in another medium. It has gained in the process the possibility of being touched, played, articulated, and has the power to translate these articulations into the needs of the machine. Thus the narrow logical channels for communication with the computer are greatly expanded.” (Ryan 1992, p. 415)

⁹Here Ryan is talking about a broad range of possible software models or simulations, not specifically physical models.

Using physical models in musical interfaces emphasises the link between physical effort and musical sounds. Ryan argues that the physicality of musical sounds is deeply embedded in our experiences of music.

“The appeal of phantom instruments is in the liberation of imagination from merely physical limitations and electronic music itself was seen as a liberation of sound. But can music thrive with no relation to the phenomenal world? Like poetry, music has a relation to history and lived experience, the emphasis should be on music from somewhere rather than music from nowhere.” (Ryan 1992, p. 416)

Physical Models as Dynamic Independent Mapping Layer

Momeni and Henry (2006) have written about the use of a “dynamic independent visual-mapping layer” (Momeni and Henry 2006, p.50), which involves the insertion of an extra layer into the traditional mapping framework (fig. 2.2). This dynamic layer can be any kind of time-varying generative system. Physical models are one such system, and Momeni and Henry also describe their use of perceptual spaces as another possibility. In this approach, the user’s gestures influence the generative system in some way so that its state changes over time. This dynamic behaviour can be measured and mapped via additional mapping layers to audio and visual synthesis. A similar approach has been described by Choi (Choi 2000) and Mulder et al. (1997). Likewise, scanned synthesis (see chapter 7, section 7.3.3) is based on this paradigm.

To illustrate, a virtual instrument using a physical model as the dynamic layer might work as follows. Gestural input from the user results in force being exerted on the physical model and in response parts of the model move about, bump into each other, etc. Various measurements of the state of the model, such as speed of individual masses, forces being exerted, acceleration and so on, are then separately mapped to parameters for audio and visual synthesis engines. The visual synthesis mapping layer might map the X, Y and Z coordinates of masses to the position of several geometric shapes on screen for example, whereas the audio synthesis mapping layer might map forces exerted on each mass to synthesis parameters such as the individual amplitudes of a set of oscillators.

It can be seen that with this approach we end up with three mapping layers in addition to the dynamic layer. The first maps from user gestures to parameters which change the state of the dynamic layer. The second and third map from measurements of the state of the dynamic layer to audio and visual synthesis parameters.

This approach provides a number of advantages. First, because both audio and visual synthesis parameters have the same source (the dynamic layer), the intimate linkage of sound and vision is greatly simplified. While they may be separated if desired (by treating the outputs from the dynamic layer in dramatically different ways),

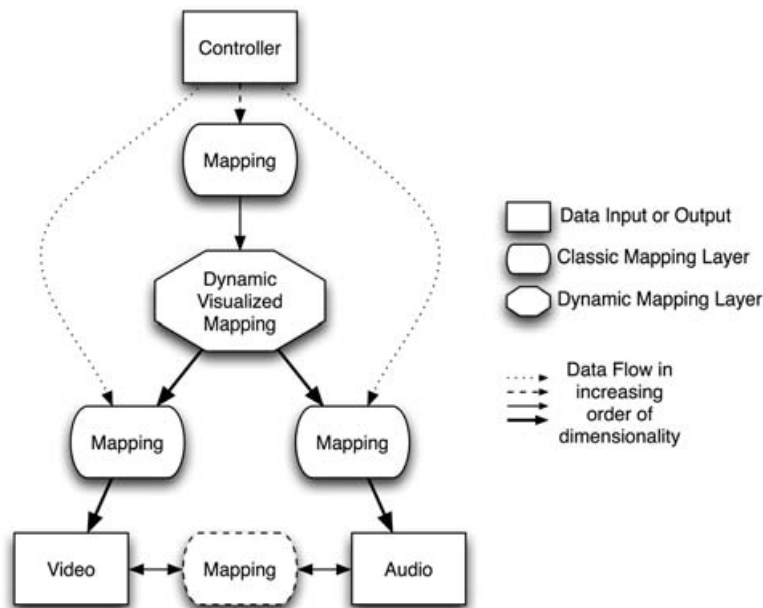


Figure 2.2: Block diagram of a virtual instrument using a dynamic independent visual-mapping layer (Momeni and Henry 2006, p.50).

the ‘default’ condition is likely to lead to clearly perceivable correspondences between sound and vision.

Second, the dynamic layer provides convenient ways to build instruments based on divergent (one-to-many) mappings. A mass-spring physical model which contains a network of say 10 masses linked together with springs can be set in motion by moving only one of the masses. The movement of this single mass results in multiple movements in the overall structure as the force propagates through the network via the links. Each of these movements is predictable at a high level and is a direct result of the initial user action. These derived movements provide extra streams of data which may be mapped to audio/visual synthesis parameters.¹⁰

Third, if the visual display is a representation of the dynamic layer itself (eg. a display of the actual physical model), then the user is more able to understand the state of the system, leading to an improved ability to control the instrument. In addition, such a display can help an audience understand and engage with a live performance as they are more able to perceive what impact the actions of the instrumentalist have on the virtual instrument.

Finally, the movements of the physical model bring a sense of dynamism to the virtual instrument. As the physical model network reacts to energy supplied by the

¹⁰The Web, a physical controller designed by Michel Waisvisz and Bert Bongers (Krefeld and Waisvisz 1990, Bongers 2006) also explores the interconnection of individual controller elements. The web is “an aluminium frame in an octagonal shape with a diameter of 1.20m., and consisting of six radials and two circles made with nylon wire” (Bongers 2006, p. 63). Tension in the strings was measured by custom designed sensors, providing a stream of data for sound synthesis. Because of the interconnected nature of the Web, the output was complex and not always easy to control.

performer it will often oscillate, providing rhythms the player can respond to. By bringing a sense of unpredictability and a kind of simple agency to the interaction, while still retaining high-level controllability, a physical model mapping layer may help stimulate a more conversational style of musical interaction (Chadabe 2002).

Summary In this section I have given an overview of the use of simulated physical models in musical applications. While physical models have a long history in sound synthesis, their use in musical *interfaces* is comparatively under explored. Recently the use of physical models as a dynamic independent mapping layer has been proposed as a way to:

- build virtual instruments based on divergent (many-to-many) mappings;
- link sound and vision; and
- make the behaviour of the virtual instrument more transparent for performer and audience.

The virtual instruments described in chapter 4 use this technique.

2.5 Interactions with Virtual Instruments

Performers can interact with virtual instruments in a number of ways. A key question that arises is: what is the nature of the relationship between the performer and the instrument? Who is in control? Is the virtual instrument merely a tool that responds to commands issued by the performer or does it 'talk back', asserting some degree of autonomy?

While there has been considerable discussion of these issues in the literature, there have been few structured studies of how musicians interact with virtual instruments. An important exception in a closely related area is the study by Andrew Brown into how composers interact with computer-based tools (Brown 2003).

Despite this, a number of taxonomies which group virtual instruments into various categories based on either their technical characteristics or how the interactions between musicians who use them are (or are intended to be) structured have been put forward. These include Robert Rowe's taxonomy of interactive music systems (Rowe 1993), Winkler's performance models (Winkler 1998), the continuum of control described by Joel Chadabe (Chadabe 2002), as well as a taxonomy arising from Brown's study mentioned above.

I will reserve discussion of this work until chapter 7 (section 7.2). The works cited above were kept in mind during the user study described in chapters 3 and 5 but I attempted to avoid having them influence the findings. This is in accordance with the grounded theory technique Glaser and Strauss (1967), particularly as espoused by Glaser (Glaser 1978), which encourages researchers to avoid having preconceptions

about what will be found in data gathered during qualitative studies. Thus, while I flag that such taxonomies exist here, I will present them in detail later when they can be placed in context with the findings from the study described in chapter 5.

2.6 Audio-visual works for acoustic instrument & computer

Having identified a number of design criteria for computer based musical instruments and discussed the use of physical models as an independent mapping layer, I will now discuss a number of significant audio-visual virtual instruments which influenced the design of the virtual instruments discussed in this thesis. Each will be briefly described and discussed in relation to the design criteria outlined in section 2.2.

This summary is not meant to be comprehensive history of virtual instruments, which is beyond the scope of this thesis. Instead this review is limited to representative works which:

- Respond in real-time to live audio;
- Are intended to facilitate musical expression;
- Provide some kind of audio-visual representation of, or response to, the live music;
- Are intended for use in a performance context; and/or
- Make interesting use of physical models.

The reader is referred to Jordà (2005) and Miranda and Wanderley (2006) for broader overviews of the history of virtual instruments.

2.6.1 *Messa di Voce*

The interactive performance work *Messa di Voce* by Golan Levin and Zach Lieberman (Levin and Lieberman 2004) explores very similar territory to the work described in this thesis. Developed in conjunction with singer/composers Joan La Barbara and Jaap Blonk, this large-scale (30-40 minute) work was premiered in 2003. Using speech analysis techniques, the software for *Messa di Voce* transforms the live vocal sounds into dynamic graphical elements. These animated graphics augment and exaggerate various sonic characteristics of the live sound. In addition, the use of movement tracking techniques allows the performers to interact with the generated graphics gesturally. The software component of the work is comprised of 12 modules or “audio-visual vignettes” (Levin et al. 2004, p.7).

In performance, the two vocalists perform in front of a large (approximately 12m x 4m) projection screen showing the visual response to live audio. Both performers




Figure 2.3: A performance of *Messa di Voce* (the *Bounce* module). Here the vocalist Jaap Blonk is making cheek-flapping sounds which cause a stream of bubbles to be shown on screen as if emerging from his head. (*Reproduced with permission.*)

wear microphones and their vocalisations are amplified and also of course analysed by the software in order to generate the live graphics. The software tracks the performers' movements and this means that the animated graphics projected upon the screen can react to the position and movement of their bodies. Thus, at times, steams of bubbles or plasma appear to be emerging from the performers' mouths as they perform.

Messa di Voce is part of a series of works exploring the question: "if we could see our speech, what might it look like?" (Levin and Lieberman 2004, p. 7), and strong emphasis is therefore placed on the link between the live vocals and the computer generated visuals. Few of the software modules produce sound, and those that do tend to replay musical phrases sampled directly from the live performance rather than synthesised or processed sound (Levin 2008). In addition, initiation of playback is not directly linked to the gestures (physical or vocal) of the performers. As such, while Levin and Lieberman designed the software to be "deeply instrumental for [the] vocalists and *commensurately expressive*" (Levin and Lieberman 2004, p. 11), it is predominantly a *visual* instrument controlled by music as opposed to a musical instrument.

This is interesting, because much of Levin's past work has been in developing virtual musical instruments which use a "painterly interface" to dynamically create audio-visual performances (Levin 2000). In his Masters thesis, Levin describes the design and development of five such audio-visual performance instruments which use sophisticated mappings between dynamic drawings (and the gestures that cre-



Figure 2.4: A still from the *Fluid* module of *Messa di Voce* showing a glowing fluid or plasma emerging from the performers' mouths as they vocalise. (*Reproduced with permission.*)

ate them) and real-time audio synthesis. The idea is that the animated drawings created live are perceived as being tightly coupled to the computer generated audio, thereby facilitating an instrumental, expressive interaction. It would be very interesting to see a kind of extension to *Messa di Voce* which used some of these techniques to build a more direct link between the animated graphics which visualise the live sound and computer generated sounds. The virtual instruments developed as part of this thesis could be seen as a kind of preliminary exploration of the possibilities of such a hybrid approach. Seen in this light, they occupy a space between *Messa di Voce* and these earlier Levin works in that the animated graphics simultaneously act as both visualisation of live acoustic sounds, and generator of computer generated sound. The virtual instrument in this case acts as a kind of filter or mediating structure (see chapter 4, section 4.4.2).

2.6.2 The Singing Tree

Part of the large-scale *Brain Opera* by Tod Machover and the Opera of the Future Team at the MIT Media Laboratory, the *Singing Tree*, like *Messa di Voce*, provides a visual response to live singing but was designed for use by novices rather than professionals. Act 1 of the opera involves the audience exploring and creating music with a number of interactive works, one of which is the *Singing Tree*. Physically, it consists of a hood, somewhat resembling those in which public telephones are sometimes mounted, which contains headphones, a microphone and a small, flat screen (figure 2.5). When the user sings into the microphone, software analyses



Figure 2.5: Photograph of the *Singing Tree* at the Haus der Musik, Vienna. (Photograph by the author.)

their singing to determine if they are singing a steady note. If so, they are rewarded with an accompaniment of consonant string, woodwind and vocal harmonies through their headphones. If not, the provided accompaniment becomes more dissonant and makes use of “brassier and more percussive instruments” (Oliver et al. 1997, p.262).

The video display is consistent with (but not directly linked to) the computer generated audio, and is once again intended to reward those who sing a steady pitch:

“The Singing Tree has three unique video streams: the front view of a human face, the side view of a human face, and a human hand holding a flower bulb. Each stream starts in an inactive state; for example, the human face seen from the front has closed eyes. When the participant starts singing a steady pitch, the video wakes and proceeds to an identifiable goal. For example, an eye opens and one zooms inside the eye to find a dancer spinning, or the flower in the hand blossoms. The mapping of the singer’s ability to maintain pitch is direct; the video proceeds toward the goal provided the participant maintains a steady pitch, and reverses otherwise.” (Oliver et al. 1997, p. 262)

As a musical instrument, the *Singing Tree* is obviously and deliberately limited. Its “reward-oriented response” (Oliver et al. 1997, p. 262) gives a form of feedback on vocal performance as opposed to providing an interface for musical expression. The more abstract, evocative nature of this feedback contrasts with software such as *Sing and See* (Thorpe 2002), *Voce Vista* (Miller and Schutte 2002) and *WinSingad*

(Howard et al. 2004, 2007), which provide real-time feedback on singing performance in the form of graphs and other visualisations of a more quantitative nature.

Despite the fact that both audio and visual streams share a common control source (the voice), they share only one control parameter: steadiness. When the computer determines that the user is singing a steady tone, the video will play forwards and the computer generated audio will have a consonant quality. If not, playback direction is reversed and the audio becomes dissonant. This means that the richness of the link between computer generated audio and video is limited. Micro-temporal events in the video – the opening of the flower for example – have no impact on the audio and vice versa, providing limited scope for using the *Singing Tree* as an expressive instrument.

Like the *Singing Tree*, the virtual instruments I have developed provide both audio and visual responses to live sound. However, as discussed, the audio and visual elements of the *Singing Tree* are simple and also distinct from one another. That is, while the generated visuals are intended to be consistent with the audio, the link between them is not fundamental. If, for example, the software developed a fault which caused video playback to freeze, the audio would not necessarily be impacted. The audio and visuals for my virtual instruments, on the other hand, are integrated: the physical model acts as a source of control data for both. If a bug were to stop the simulation, both audio and visuals would halt.

2.6.3 The Metasaxophone

The Metasaxophone is an acoustic Selmer tenor saxophone that has been extensively augmented with sensors in order to convert various instrumental gestures into a continuous data stream (output as Musical Instrument Digital Interface (MIDI) messages) to a computer running Max/MSP. A mixture of Force Sensing Resistors (FSRs), triggers and an accelerometer are used. The FSRs are used to continuously measure pressure applied by the performer to the keys and thumb rests of the instrument. Depending on the configuration of the software, the triggers may be used by the performer to initiate actions or alter the state of the software in some way. Data from the accelerometer indicates the position of the saxophone in two dimensions (left/right and up/down) (Burtner 2002, Burtner and Serafin 2002). Additionally, the Metasaxophone is equipped with three small microphones which may be positioned at desired locations inside or outside the instrument.

This combination of acoustic instrument, continuous and discrete sensors and microphones provides a complex and rich hardware interface for musical expression. Matthew Burtner, designer and player of the Metasaxophone, has composed a number of works for the instrument. Of particular relevance to the work described in this thesis is the use of the Metasaxophone instrument/controller to explore what Burtner describes as “instrument controller substitution” (Burtner 2002, p.208) in the work *S-Trance-S*:

“Instrument controller substitution experiments with the recombination of an instrument controller interface and a physically modelled instrument of an entirely different type. In an ongoing project with Stafania Serafin, the Metasaxophone has been used as a controller for bowed string physical models. By bowing the string from within the gestural space of a wind instrument, new expressive potentialities of the model are opened. The disembodied nature of physical models becomes a means of recombining it with other interfaces, creating extended techniques for physical models that would not be possible for the real instrument.” (Burtner 2002, p.208)

In the work *S-Trance-S*, the Metasaxophone’s pressure sensors are mapped to various parameters of a bowed string physical model called the ‘ExBow’.¹¹ As the piece progresses, the mapping becomes more complex. Initially, data from the Metasaxophone’s continuous controllers is mapped to control parameters of a single bowed string physical model (‘ExBow 1’ in figure 2.6). For example, the B key pressure sensor (at the topfront of the saxophone in figure 2.6) is mapped to the ‘inharmonicity 1’ parameter of the ExBow 1 bowed string physical model. At a certain point in the piece, a second bowed string (ExBow 2) is introduced, but with different mappings. In addition to the inharmonicity 1 parameter of ExBow 1, the B key now maps to the ‘noise’ parameter of ExBow 2. That is, pressing on the B key will increase the value of the inharmonicity 1 parameter for ExBow 1 and simultaneously increase the value of the noise parameter for ExBow 2.

It can be seen that as more strings are added throughout the piece, the complexity of the virtual instrument increases rapidly, leading, presumably, to a corresponding decrease in controllability and therefore transparency for both performer and audience. This, of course, is not a criticism of the piece, as the overall artistic goals of the work did not include designing an all-purpose expressive musical instrument. Rather this increasing complexity was a deliberate compositional strategy to gradually increase the overall timbral richness as the piece progressed.

Another interesting aspect of *S-Trance-S* is the use of “instrumental cross synthesis” (Burtner 2002, p. 210), in which convolution techniques are used to combine the sounds of saxophone and strings. In this way, the dynamic, physically modelled string sounds are in a sense merged with the live saxophone to create hybrid sounds (see figure 2.7). This use of the live instrumental sound coloured or mediated by the state of a dynamic physical model is a similar technique to that used in *Partial Reflections I* (see chapter 4).

S-Trance-S, in effect closely couples two separate instruments: an acoustic saxophone, and a virtual instrument which allows the performer to manipulate four physically modelled strings using pressure sensors and switches. The instruments are coupled because the pressure sensors are mounted on the saxophone’s keys. This

¹¹ Note that the complete instrument, comprising the Metasaxophone controller and software including the mapping layer and physical models, was developed for this particular work; the intention was not to create a general purpose musical instrument.

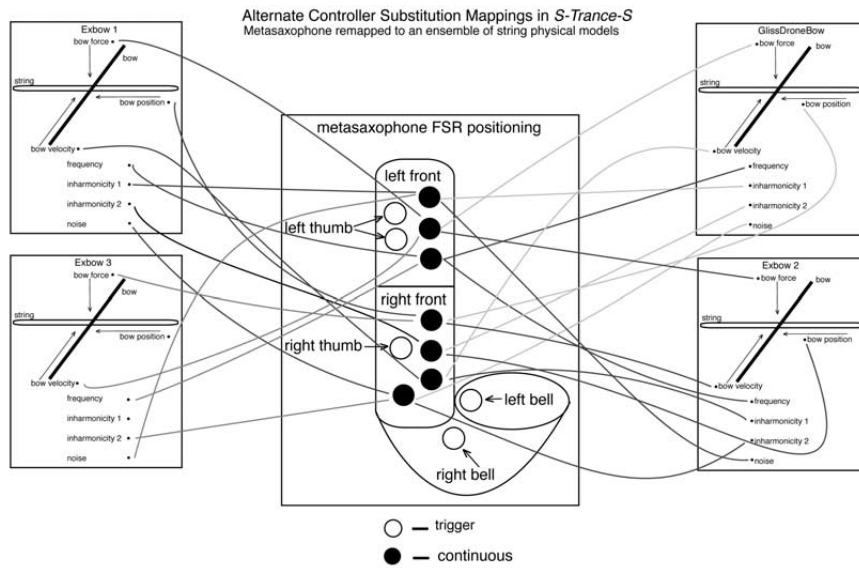


Figure 2.6: Diagram showing mappings from the Metasaxophone sensors to parameters of ExBow physical models (Burtner and Serafin 2002, p. 138)

means that while the pressure sensors (and through them the physical modelled strings) may be manipulated on their own, without blowing the acoustic instrument, the musician can't play the saxophone in the traditional sense without simultaneously playing the virtual instrument. As soon as the keys of the saxophone are pressed to change pitch, the mapping from the associated pressure sensors to the physical models will cause the state of the virtual instrument, and thus the generated sounds, to change.

This work explores similar territory to the virtual instruments I have developed. As in *S-Trance-S*, a traditional acoustic instrument is used to manipulate a simulated physical model. However, while I have not seen this work live, from descriptions of the work and videos available online, my observation is that I have placed greater emphasis on transparency and control than Burtner. The Metasaxophone interface is very complex, which opens up a large number of possibilities for gestural control over the computer generated sounds using finger, hand and arm movements, but would also tend to reduce transparency. My virtual instruments use acoustic sounds alone as a source of gestural control. This reduces scope for non-traditional interaction, but increases the likelihood the instruments' behaviour will be transparent to audience and performer (see section 4.4.2).

In addition, the virtual instruments described in chapter 4 provide a visual (as well as aural) representation of the state of the virtual instrument, whereas *S-Trance-S* does not. It therefore seems likely that Burtner has prioritised richness and complexity of response over controllability and transparency.

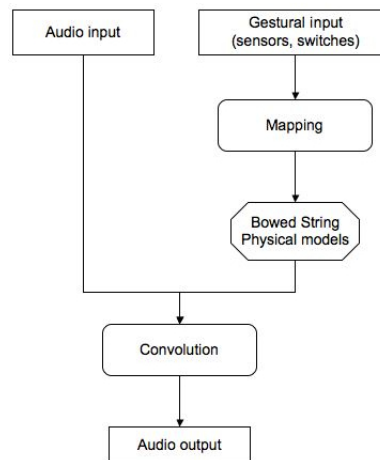


Figure 2.7: Diagram showing the relationship between live audio and acoustic and sensor data streams from the Metasaxophone in the piece *S-TranceS*. Note that although not shown here, the bowed string sounds are also convolved with one another, which adds further to the sonic complexity of the piece.

2.6.4 Hyperinstruments

From 1987-1991, a group of researchers at the MIT Media Lab led by Tod Machover developed a series of 'Hyperinstruments'. A Hyperinstrument in this sense is essentially a traditional acoustic instrument augmented by software. Various performance gestures are measured and fed into the computer system as parameters for various audio synthesis techniques. The performance gestures can be measurements of the physical actions of the performer, such as bowing speed or wrist angle of a cellist for example, or the actual acoustic sounds produced by the instrument (which after all are encoded physical gestures themselves).

The HyperInstruments, "range from the highly deterministic (one physical event is tied to a particular musical result), to the computationally flexible (musical result chosen by particular context, and modified according to the specific musical gesture received), to the truly intelligent (analysis by rule of performed music and the machine-choice of an appropriate response)" (Machover 1992, p. 4). By these definitions, the *Partial Reflections* virtual instruments described in this thesis fall into the 'computationally flexible' category.

Of particular relevance is the 'Hypercello' a virtual instrument developed in collaboration with renowned cellist Yo-Yo Ma in parallel with the composition of the work *Begin Again Again...* by Machover. Like the Metasaxophone (see section 2.6.3) a number of sensors gather information on the physical gestures of the cellist while playing the instrument. This stream of information is used by software to trigger and control various synthesis processes and parameters of effects applied to the acoustic sounds made by the cello. Several techniques are used to map between physical

and musical gestures and the computer generated audio. These include:

- using different articulation styles on the instrument to trigger a change of state in the software;
- using pitch recognition techniques to enable specific accompaniments for particular phrases played on the cello to be generated by the computer;
- tracking the position of the bow on the string to select playback of particular timbres by the computer;
- using a combination of bow position, bow-finger pressure and tremolo speed to control volume, timbre and transformation of the live cello sound;
- using onset detection to control playback of computer generated tone clusters; and
- using the cello to control additive synthesis (Machover 1992).

Because of its similarity to the control and synthesis techniques used in *Partial Reflections I*, the techniques used to map between instrumental gestures and additive synthesis output are worthy of further consideration. There are two systems that were devised during the design and composition process (Machover 1992, p. 77). For technical and aesthetic reasons the performance version of the hypercello used the simpler of these two. Here, however, we concentrate on the more complex system, as this is closer to that used in *Partial Reflections*. The mapping strategy is described by Machover below:

“The timbre instrument was designed to be controlled from the cello. The interface to the cello assigned a different function to each string. The C string selected a fundamental frequency note to be played. A down-bow started a note, and an upbow stopped it. While bowing down, the loudness of the note was set from the amplitude of the bowing. The note continued until the next upbow. The G string selected a timbre from a pre-set library. One timbre as assigned to each chromatic pitch. If a note was being played when a new timbre was selected then a timbre interpolating mutator was started. The D string selected a mutator, in the same way that the G string selected a timbre, and the A string was used to supply additional parameters to the mutator.” (Machover 1992, p. 77)

The notes played by the cello can be seen as selection gestures (see section 2.3.2). Notes played on the C string select the pitch played by the additive synthesis engine, G string notes set the timbre and D string chooses a ‘interpolating mutator’ which provides transitions between timbres. While this use of live sound to control additive synthesis processes seems a good fit, there is a lack of available information on musicians’ experiences with this instrument. How transparent did they find the

interface? Did they find it an expressive instrument? Which aspects did they find most satisfying? What frustrated them?

Another characteristic of the hypercello is that the response of the software changed significantly as *Begin Again Again...* progressed. In fact, the entire work is an exploration of the balance of power between acoustic cello and computer. "In the beginning of the piece, the solo, independent quality of the cello is emphasized" (Machover 1992, p.79). Gradually, the computer asserts itself: "By the *Jaggedly Aggressive* section, there is clearly a sort of struggle going on. The solo cello is felt to be struggling, both to keep up with the incessantly rapid tempo of the computer part, and to make itself heard over the many layers that are beginning to accumulate" (Machover 1992, p.79). Finally, acoustic instrument and computer are reconciled:

"There is no longer a conflict between solo cello and accompaniment; the two are inextricably linked, with the tremblings of the one being reflected in the reactions of the other, and vice versa. The hypercello and its extended sonic world—through the struggle and reconciliation of the first part of the work—have finally become one, a new kind of instrument." (Machover 1992, p.80)

Some of the techniques used by the Hypercello, as used for *Begin Again Again...*, were used for the virtual instruments I will shortly describe in chapter 4. Like Machover, I have used acoustic instruments as a controller for additive synthesis. Like him, I used notes played on the acoustic instruments as 'selection' gestures. However, in other respects the Hypercello and my virtual instruments are very different. Firstly, the Hypercello exists physically - it is not a purely virtual instrument. Secondly, (at least for the piece *Begin Again Again...*), the software response changes significantly over time, whereas in my instruments the response is essentially consistent. Finally, the Hypercello is not used to trigger or control computer generated visuals in *Begin Again Again...*, whereas my instruments integrate audio and visuals.

2.7 Conclusion

In this chapter I have given an overview of key issues in virtual musical instrument design. While there are no recipes which are guaranteed to produce successful virtual instruments, there are a number of commonly accepted design criteria. These include:

- Virtual instruments should be conceptually simple, while providing musicians with the ability to create a wide range of complex sounds.
- In order for musicians to develop intimate, expressive connections with virtual instruments, the links between their actions and the responses of the virtual instrument need to be clearly apparent and understandable.

- Musicians should be able to predict what effect their actions will have on the virtual instrument. Thus when instrument state is dynamic, its current state should be made apparent to performers and audience.
- Virtual instruments should be adaptable to new musical contexts. However, the degree to which diversity is important depends on artistic context.
- Notions of efficiency and control in virtual instrument design are likely to be less important than giving them character and personality which motivates and inspires musicians. This may involve trading controllability for complexity of response.
- Virtual instruments can operate at different levels, allowing the performer detailed control over sound timbres and/or higher level musical structures such as phrases. The virtual instruments described in this thesis operate at the timbre (or 'microscopic') level.
- Virtual instruments which use cross-coupled, many-to-many mappings and which require energy injection are likely to be more engaging and expressive.
- Some authors argue that instruments which change state as they are played are more likely to stimulate a conversational style of interaction.

I have proposed the use of physical models as an intermediate mapping layer between performer actions and computer generated sounds and visuals and outlined how they might be used to provide musical interactions which balance control and complexity. In chapter 4 I will describe how this technique was used to create a series of three virtual instruments.

Finally I have presented some representative audio-visual works for acoustic instruments and computer in order to provide further context for the virtual instruments I have created. Before I describe these, however, I will spend some time outlining how I have combined practice (the design of virtual instruments in an artistic context) and research (the study of musicians who used these instruments). This is the subject of chapter 3.

Chapter 3

Methodology

3.1 Introduction

The aim of this practice-based research project is to design virtual musical instruments and examine their impact on the practice of musicians who use them.

The main goals and objectives are:

1. To develop a set of virtual instruments which use simulated physical models to mediate between live sound and computer generated sound and visuals;
2. To investigate musicians' experiences with these virtual instruments;
3. To devise a set of design criteria for virtual instruments of this type.

The overarching research questions are:

- What are the design criteria for virtual instruments which use physical models as a mapping layer between live audio and computer generated audio-visuals in order to facilitate musical expression and exploration?
- How do expert musicians interact with virtual instruments such as these?

The questions I will consider in this chapter are:

- How should we go about designing the virtual instruments?
- How should design criteria which guided the development of the instruments be identified?
- How should the experiences of musicians using the virtual instruments be studied?
- How should the data gathered from the studies be analysed in order to build an understanding of what is going on and the relationship between the design criteria and the musicians' experiences?

At a high level, the proposed structure of the project is as follows:

1. Examine the literature to identify initial design criteria for virtual instruments.
2. Develop a series of virtual instruments using iterative, participatory design methods.
3. Conduct a series of user studies to examine how musicians interact with the virtual instruments.
4. Derive a set of design criteria for virtual instruments of this type informed by practice (step 2) and research (steps 1 and 3).

A diagram of this process can be seen figure 3.1.

In the following sections I will describe and justify the specific research methods that were used.

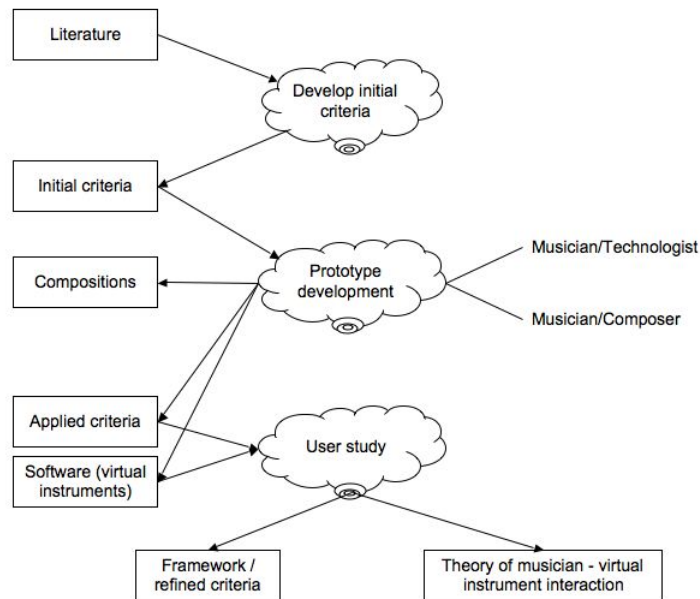


Figure 3.1: Overview of the research project. Processes are shown as clouds, concrete outcomes as boxes.

3.2 Research Paradigm

Before describing the research methods in detail, it is necessary to spend some time outlining the overall research paradigm applied in this research. To do this will require some consideration of the ontological and epistemological positions which underpin the research methods that have been applied.

Crotty (1998) argues that in developing a research plan, researchers need to consider two questions:

- Which methods should be used?
- How can these methods be justified?

These questions require the researcher to confront their underlying assumptions and beliefs about the nature of knowledge. Crotty suggests that the basic elements of the research process - methods, methodology, theoretical perspective and epistemology - should be carefully considered and the researcher's position made explicit. Crotty defines these terms as follows:

- *Methods*: the techniques or procedures used to gather and analyse data related to some research question or hypothesis.
- *Methodology*: the strategy, plan of action, process or design lying behind the choice and use of particular methods and linking the choice and use of methods to the desired outcomes.

Epistemology	Theoretical perspective	Methodology	Methods
Objectivism Constructionism Subjectivism (and their variants)	Positivism (and post-positivism) Interpretivism • Symbolic interactionism • Phenomenology • Hermeneutics Critical inquiry Feminism Postmodernism etc.	Experimental research Survey research Ethnography Phenomenological research Grounded theory Heuristic inquiry Action research Discourse analysis Feminist standpoint research etc.	Sampling Measurement and scaling Questionnaire Observation • participant • non-participant Interview Focus group Case study Life history Narrative Visual ethnographic methods Statistical analysis Data reduction Theme identification Comparative analysis Cognitive mapping Interpretative methods Document analysis Content analysis Conversation analysis etc.

Figure 3.2: Crotty's four elements of research (Crotty 1998, p. 5).

- *Theoretical perspective*: the philosophical stance informing the methodology and thus providing a context for the process and grounding its logic and criteria.
- *Epistemology*: the theory of knowledge embedded in the theoretical perspective and thereby in the methodology." (Crotty 1998, p. 3)

Crotty's taxonomy of research paradigms will be used here to help clarify the approach taken in this research. Somewhat unusually, Crotty blends ontology (the nature of reality) and epistemology (the nature of knowledge) together, arguing that, "ontological and epistemological issues tend to emerge together" (Crotty 1998, p. 10). However, he does not ignore ontological issues but rather considers them in conjunction with epistemological issues as they arise.

Starting in the leftmost column of figure 3.2, we consider epistemology, the nature of knowledge. The epistemological position taken in this research is *constructionism*. According to Crotty, constructionism is, "the view that all knowledge, and therefore all meaningful reality as such, is contingent upon human practices, being constructed in and out of interaction between human beings and their world, and developed and transmitted within an essentially social context." (Crotty 1998, p. 42)

The constructionist epistemological position stands in contrast to the objectivist position, which holds that objects in the world exist and have meaning independent of human beings.

Objectivism is the epistemological view that things exist as *meaningful* entities independently of consciousness and experience, that they have

truth and meaning residing in them as objects ('objective' truth and meaning, therefore), and that careful (scientific?) research can attain that objective truth and meaning." (Crotty 1998, p.5–6)

Objectivist epistemology is most commonly linked with the positivist theoretical perspective and the natural sciences. Researchers applying this paradigm make strenuous efforts to maintain separation between the observer and the object of the research. Typically this involves making 'objective', quantitative measurements of phenomena which are considered relevant to the research. In the natural sciences, this has been extremely effective. However, at the very least in creative domains such as music this is problematic. The position taken in this research has been that music is a social construction (Gergen 1999). That is, what constitutes 'good' music (and by extension a 'good' musical instrument) is completely dependent on the social context in which it is created and experienced. From this perspective then, it is not possible to objectively measure whether any particular music is 'good' or 'bad'. Goodness or badness only has meaning in a social setting.

Thus the theoretical perspective that has been applied here is *interpretivism*.

A positivist approach would follow the methods of the natural sciences and, by way of allegedly value-free, detached observation, seek to identify universal features of humanhood, society and history that offer explanation and hence control and predictability. The interpretivist approach, to the contrary, *looks for culturally derived and historically situated interpretations of the social life-world*. (Crotty 1998, p. 67, italics in original)

The choice of a constructionist epistemology and interpretive theoretical perspective informed the choice of methodology and methods which will be justified and described in the following sections.

3.3 Research Approach

The overall structure of this research was informed by Design Science and Action Research. These methods share a common emphasis on intervening in some situation in an attempt to build understanding and improve the lives of the people involved.

3.3.1 Action Research

The term 'Action Research' was coined by Kurt Lewin (Lewin 1946) and this approach has since been widely adopted, particularly in education. At a high level, action research is concerned with improving a situation through theoretically-informed action and reflection. Research progresses through an iterative cycle of planning, action, observation and reflection (figure 3.3). Kemmis and McTaggart (1988) describe the action research cycle as comprising four steps in which participants undertake:

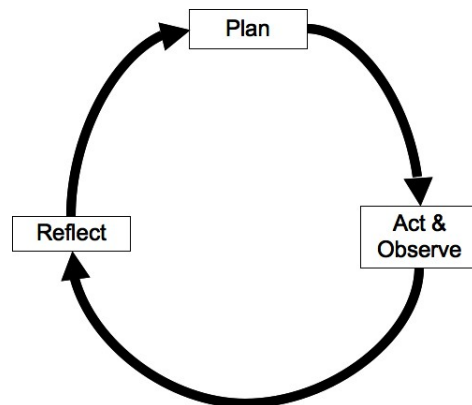


Figure 3.3: Action research cycle.

- “to develop a plan of critically informed action to improve what is already happening;
- to act to implement the plan;
- to observe the effects of the critically informed action in the context in which it occurs; and
- to reflect on these effects as a basis for further planning, subsequent critically informed action and so on, through a succession of cycles.”
(Kemmis and McTaggart 1988, p. 10)

It could, of course, be argued that all human activity follows these steps. What differentiates action research from normal, everyday good practice is an emphasis on rigour and the explicit building and refining of theoretical understanding. Thorough gathering of evidence by careful observation enables critical reflection and evaluation of the impact of action on the problem setting.

The definition of Rapoport (1970) emphasises the interplay between theory (of social science in Rapoport’s case) and practice:

“Action research aims to contribute both to the practical concerns of people in an immediate problematic situation and to the goals of social science by joint collaboration within a mutually acceptable ethical framework.” (Rapoport 1970, p. 499)

The literature on action research emphasises that it is a collaborative, rather than individual activity, arguing that human activities take place in a social context and should not, therefore, be analysed in isolation. Thus, the phases of observation, reflection and evaluation described above are group activities with the aim of generating positive cultural, as well as individual, change. Unlike the natural sciences, action research is fundamentally participatory in nature, in that action researchers are actively engaged in understanding and improving their own practices. Understanding

and improvement of others' practice is a secondary goal (Kemmis and McTaggart 1988).

According to Kemmis and McTaggart, the general form of an action research plan is: "We intend to do X with the view of improving Y" (Kemmis and McTaggart 1988, p. 19). An important point is that what is done and the reasons for so doing are likely to change as the research progresses. As the action researchers develop understanding of the situation and build experience by attempting to realise their planned improvements, it is natural that goals and methods will be adjusted.

A problem with using action research in the context of musical expression is that it is generally used by larger groups in domains such as classroom teaching where the emphasis is on the group working together to identify problems and address them collectively. This is problematic in the context of musical expression because composers and musicians tend to have diverse artistic goals and highly individual methods. A major part of composers' artistic development is the creation of their own 'voice', so developing a single virtual instrument for several different musician/composers may result in a bland, overly generic instrument which attempts to cater to all styles. One of the advantages of computer based interfaces for musical expression is their capacity for individualisation. If, in order to involve greater numbers of musicians in the project, an attempt is made to create virtual instruments which are suitable for use by a large group of musicians in diverse musical contexts then the risk is that more radical design alternatives will be excluded, and simple virtual versions of traditional instruments produced. The decision in essence is whether to privilege specific relevance over general relevance. If an instrument is designed for an individual musician and/or composer, we can be somewhat confident that it will meet their needs even if it is not used by anyone else. A general purpose instrument on the other hand may not offend most musicians, but may not be satisfying for any of them.

A second issue with action research is the relative importance given to the development of theory. Rapoport (1970) argues that there is a tension in action research projects between the desire of researchers to develop theories with strong practical foundation and practitioners' desires to address more immediately practical concerns. Rapoport claims there are many examples where action research results are "purely service oriented with little or no yield for social science" (Rapoport 1970, p. 506). Of course the other extreme is equally problematic:

"On the opposite, 'ivory tower', extreme is the researcher who has used the industrial environment to gather information for purposes unrelated to any concerns of the organisation or those in similar situations and therefore of little or no relevance to their requirements. These researchers may be seen as parasitical by those involved in the practical affairs, or as highbrow 'smash-and-grab' types." (Rapoport 1970, p. 506)

There is no definitive solution to this problem; it is inherent in the nature of action research with its dual emphases on both practical intervention and contribution to theoretical understanding. The critical point is to ensure that participants understand and agree that the needs of researcher and 'client' are important and are to receive adequate attention.

A final issue with using action research in this project is that the action research literature emphasises that the researcher intervenes in some situation, usually at the request of a 'client' in order to address a 'problem'. For this research, the focus is on exploring an opportunity - the potential use of simulated physical models in live music applications - rather than addressing a particular problem. While it is hoped that this technology may indeed address some issues that are problematic with interfaces for live music performance, the focus is on exploring possibilities and not on solving specific problems.

So, from action research, the following principles are taken:

- Active involvement of both researcher and composer/musician in the design/research process.
- Identification of successful and unsuccessful strategies through practice.
- Contribution to practice and theory.
- Use of an iterative approach of action and reflection.

Recognising the tension between the researcher's need to develop theory and the composer/musician's desire to focus on artistic outcomes, some aspects of the research have been conducted outside the artistic collaboration which produced the virtual instruments. This is contrary to action research principles which would suggest that research and practice should be intimately bound together in order to maintain relevance to practitioner participants. However, by examining how musicians unfamiliar with the project interact with the virtual instruments, some objectivity and the chance of fresh perspectives is gained. It is recognised though, that by conducting an external evaluation, issues are being examined that are of limited interest to the musician who is more interested in practical applications than developing theory. In other words, some rigour and general relevance is gained at the expense of immediate relevance (Schön 1983).

3.3.2 Design Science

Design Science is a research method in Information Systems that emphasises the relationship between the design process and theoretical knowledge. A fundamental principle of design science research is that "knowledge and understanding of a design problem and its solution are acquired in the building and application of an artefact" (Hevner et al. 2004, p.82). The idea is that relevance to a particular domain is maintained by building systems in close collaboration with domain experts

– building systems that attempt to address their specific issues. Academic rigour is maintained by thoroughly evaluating the systems that have been developed in use and examining their impact. By drawing on the experience of development and findings from the evaluation process, a theoretical understanding of the problem domain is developed and refined.

Design science is positioned as a natural compliment to the natural sciences approach to research (Simon 1981, March and Smith 1995). In the natural sciences, the researcher attempts to carefully observe phenomena with the aim of uncovering the underlying laws which explain why things are as they are and allow us to make predictions. A key principle of natural science research is that of minimal (or at least highly controlled) intervention; the process of observation should not disturb the phenomena being studied. Design, on the other hand, “is concerned with how things ought to be, with devising artifacts to attain goals” (Simon 1981, p.133). This implies intervention in the area under investigation. The designer-researcher creates an artefact which in some way disturbs the status quo and examines and evaluates its impact.

It is important to note that it is not necessary to choose *either* a design science or natural science approach. March and Smith argue that:

“...an appropriate framework for IT research lies in the interaction of design and natural sciences. IT research should be concerned both with utility, as a design science, and with theory, as a natural science. The theories must explain how and why IT systems work within their operating environments.” (March and Smith 1995, p.255)

This approach implicitly draws on some of the ideas of action research, in that by creating an artefact for use in a situation, the researcher actively participates in the area under study and is not a passive, objective observer. There are many similarities between action research and design science but surprisingly little dialogue between these two streams of research. (Järvinen (2007) has recently pointed out the similarities between the two research paradigms.) A strength of both approaches is that they do not mandate particular approaches to the evaluation/reflection phases of the research, although it might be observed that past work in action research is largely qualitative, whereas design science has a quantitative flavour.

Figure 3.4 shows a high level view of the design science research process. By drawing on the knowledge base of the designers (on the right of the diagram) and working with people in the domain under investigation (on the left), artefacts are developed which address problems or provide new opportunities. Because existing theories are used to inform its design, the new artefact is in a sense an embodiment of those theories. By observing the impact of the artefact on the people who use it and the work they do, the theories can be tested and refined and/or new theories proposed. The ‘testing’ of theory in design science has a more utilitarian focus than in the natural sciences:

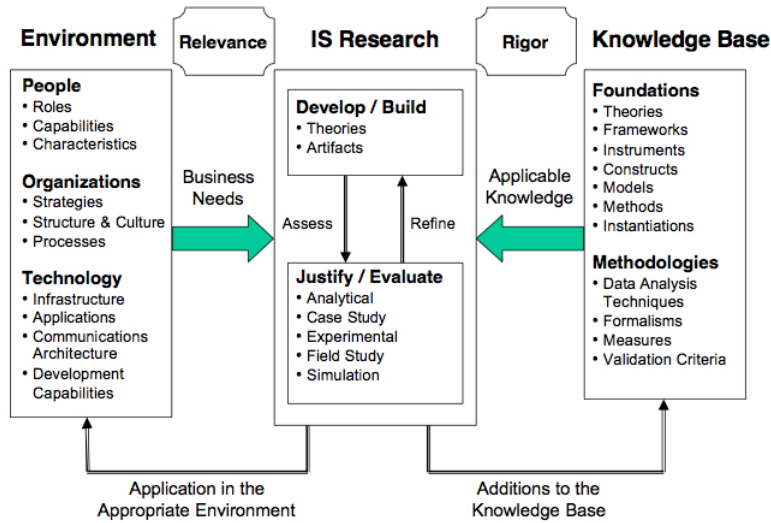


Figure 3.4: Overview of the design science approach to Information Systems research (Hevner et al. 2004, p.80). While the emphasis is on ‘business needs’, the paradigm is of course also applicable in other domains.

“Whereas natural science tries to understand reality, design science attempts to create things that serve human purposes. It is technology-oriented. Its products are assessed against criteria of value or utility - does it work? is it an improvement?” (March and Smith 1995, p.253)

For this research the following characteristics of design science are drawn:

- Building understanding of the domain through the design and use of artefacts for use in specific situations.
- Drawing on relevant theories to inform design.
- Contributing to theoretical understanding by conducting a thorough evaluation of the impact of the artefact on the practices of expert users in a specialist domain.

3.3.3 Design Approaches

Having decided on an overall approach drawing on action research and design science, it is necessary to actually design software. The question addressed in this section is: What methods should be used to design new virtual musical instruments based on physical models?

Prototyping, participatory design and agile methods

Prototyping has a long history in software design. In computing, a number of design methods and approaches have used the fundamental idea of iterative development coupled with formative evaluations. The fundamental idea is that prototype software is constructed in order to explore design possibilities and the impact of the software on user practice. The construction of 'rough and ready' prototypes allows the designers to:

- assess the technical feasibility of various designs and develop appropriate technical skills;
- gather feedback from potential users on how well the software addresses their needs;
- refine the design through successive cycles of building and evaluating the prototypes; and
- examine the impact of the prototypes on the work practices of users.

The active involvement of users in developing requirements for the prototypes and providing feedback as their design is refined is very important. The Participatory Design approach which emerged in Scandinavia in the early 1970s has been influential in emphasising the benefits of actively involving users in software design. Fundamental to this approach is the idea that introducing new technologies into users' practice alters that practice. That is, in using virtual instruments, musicians' approach to music-making will be changed, at least to some degree. By working with musicians to develop the virtual instruments in the first place, and involving them in exploring what impact they have on their musical practice, we increase the probability that the impact will be positive. By carefully documenting and examining how musicians actually use the virtual instruments to make music, it is possible to build understanding of both design techniques and music practice as the instruments' design and musicians' practice co-evolve.

Another important guiding principle for participatory design is the importance of respecting and accounting for current practice in the design process (Bødker et al. 1992, Suchman and Trigg 1991, Suchman 1987). Active user participation in design helps ensure that software designers take account of how users approach their practice currently, and where they see potential for exploration and/or improvement. In this research, composer/musician Ben Marks composed music for two of the virtual instruments. However, it should be noted that the process of design and composition proceeded in parallel, often in the same room. It was *not* the case that the virtual instruments were designed and then handed over to Ben with the expectation that he would compose music for them, or that Ben provided a completed composition for which computer generated sounds and visuals were to be created.

Broadly speaking there are two approaches to prototyping. One, named ‘horizontal prototypes’ by Floyd (1984) involves the rapid construction of non-functional prototypes which allow potential users to see what the software might look like. The user interface has the appearance of a finished product, but no (or perhaps minimal) functionality. ‘Vertical prototypes’ on the other hand have a reduced feature set, but features that are present are fully implemented. It seems likely that both approaches have benefits (Grønbæk 1989) depending on the difficulty in constructing functional prototypes.¹

Vertical prototyping was the approach used in this research. The use of the Pure Data programming environment supported the construction of simple virtual instruments with minimal effort, meaning that the musician could experiment with the software, give feedback and develop compositional ideas in a rapid cycle. In addition, Pure Data’s ability to remain running while being modified allowed many changes to be made to the software in the presence of the musician.

The technique of allowing vertical prototypes to ‘evolve’ or grow as requirements surface in dynamic or complex situations has long been considered good practice (eg. Bally et al. 1977, Boehm 1986, Brooks 1987). Prior approaches to software design (eg. Royce 1970) are generally construed (perhaps somewhat unfairly) as requiring the discovery of user requirements, usually in a series of meetings or interviews, before any software is actually written. Once requirements have been decided, the technical specifications are identified and the software is developed. An obvious flaw in this approach, and the main reason why it has fallen out of favour, is that it is predicated on the notion that users know what they want the software to do, or, even more dangerously, that a software designer can identify what users *should* want the software to do. In situations where novel uses of technology are proposed, or where requirements are emergent, this is not the case and this up-front design approach is not suitable.

A potential drawback of the prototyping approach is the tendency for ‘tunnel vision’ to develop (Sol 1984). This occurs when design proceeds as elaboration on the current prototype, and more radical approaches are not considered. One recommended technique to help avoid this problem is the deliberate exploration of a range of alternatives by developing of a series of prototypes (eg. Kyng 1988). Thus, while three virtual instruments were developed to a relatively advanced stage during this research, a number of other prototypes were developed, explored and discarded on the way.

Proponents of ‘Agile’ development methodologies (Beck 2000, Cockburn 2002, Abrahamsson et al. 2002), advocate an iterative process incorporating short phases of development and frequent feedback from users. The idea is that by getting nearly continuous feedback from the intended users, software requirements can be adjusted ‘on the fly’. Beck and Fowler (2001) draw parallels with driving a car: when

¹In circumstances where proposed systems are very complex or use new and expensive technologies, cardboard ‘mock-ups’ can also be valuable in exploring design possibilities (Ehn and Kyng 1992).

driving we don't simply point the car in the right direction, start the engine and then take a nap. Rather, we continually monitor our direction based on our current circumstances and make the necessary adjustments to our steering as we go.

“Software development is a process. It can go well, or it can go badly. To keep it going well we must continually direct it. To direct it we must frequently assess the direction it is going, compare this to the direction we want it to go, and then make careful adjustments.” (Beck and Fowler 2001, p. 13)

Perry Cook recommends that designers of digital music controllers should “make a piece, not an instrument or controller” (Cook 2001, p.1). Attempting to build a generic ‘super instrument’ suitable for all musical purposes is difficult because of the difficulty of gathering requirements and getting feedback. The question is: who should assess the direction and decide which adjustments to make?

The strategy that was applied in this research was to design virtual instruments for use by one composer/musician for one piece of music. This meant it was possible to get feedback rapidly and make design decisions without undue consideration of what impact they would have on other musical contexts. The expectation was that this would result in an instrument that was suitable for at least one musical application (the piece that was composed in parallel with the design process), and that its suitability for broader use could be explored afterwards. In other words, rather than attempting to design an instrument based on speculative understanding of what musicians would want it to do, it was decided instead to work on developing an effective instrument based on the specific requirements of a particular musician, and gather feedback from other potential users later.

The advantages of taking this approach are:

- Feedback cycles are shorter. Gathering feedback from one committed composer is easier than gathering feedback from a large number
- The compositions and controllers can co-evolve. The composer makes changes to the piece which provides an opportunity for the interaction designer to alter the behaviour of the instrument, which in turn inspires the composer to try new musical ideas, and so on.
- At a minimum, one piece of music will be written for the instrument.
- The process helps ensure that artistic goals remain at the core of the project and are not swamped by technical considerations.

There are however some disadvantages. Because they are designed for a specific work, the expressive range of the instruments will be limited. It should be remembered that these prototype instruments are intended to enable the careful examination of musicians' current practice and how that practice is impacted by instruments of this type. In this sense, the approach draws on the 'prototyping' approach

of Mogensen (1992). As the name implies, prototypes in the provotyping approach are used to provoke current practices by using them “as a concrete medium for calling forth experiences in current practice, instead of focusing on how the prototype could be improved.” (Mogensen 1992, p. 47).

A second disadvantage of prototyping from the researcher’s perspective is that it is harder to distil general rules for instrument design from the process. Just as the grounded theory approach (Glaser and Strauss 1967) builds theory by attempting to explain what is going on in specific cases rather than attempting to find general explanations for a large number of cases, this approach attempts to find design criteria which are effective in a specific situation. It may be that these criteria are more broadly applicable, but this is something that is explored *after* the design process has concluded.

3.4 Research Methods

In previous sections, the overall shape of the research has been described. An approach drawing on action research, design science, prototyping, participatory design and agile methods has been outlined. In this section, the specific methods which have been applied are presented.

3.4.1 Identifying design criteria

Following the design principles outlined, a number of virtual musical instruments were developed, and these have a number of characteristics. There are two key questions which now need to be considered:

1. What are the characteristics of these virtual instruments?
2. What impact does using the virtual instruments have on the music-making of musicians who use them?

When considering the characteristics of the virtual instruments, it is necessary to consider which perspective to take. It would be possible to describe the instruments in a purely functional terms - to describe exactly what the software does but it is equally important to consider *why* they work in this way. Of course, the software behaves in the way it does because the designers made it that way, with the intention that it would facilitate the achievement of some goals. The question is: what were those goals? Identifying the design criteria - the goals of the designers - can help in two ways. Firstly, it makes the intentions of designers explicit and secondly it facilitates the evaluation of the resulting artifacts.

Portillo and Dohr (1994) define a design criterion as, “a measure of value used by the designer to conceptualize, test and evaluate the project purpose in the design process” (Portillo and Dohr 1994, p. 405). An important point to note here is that

design criteria are used, perhaps tacitly, during the design process as the designers develop their ideas and the designed artefacts in parallel.

In this research, the design criteria applied by the software designer and the musician were identified drawing on several sources:

1. Reflective online diaries (blogs) kept by both the software designer and musician.
2. An interview with the musician after the development of the virtual instruments.
3. Reflection by the software designer.
4. Examination of software version control logs.²

Having identified the criteria applied by both participants in the design process, it was then possible to evaluate the software by asking:

- How well does the software meet the design criteria which emerged during design?
- What is the relationship between the criteria manifest in the virtual instruments and the experiences of musicians who use them?
- How successful is the software overall?

To answer these questions it is necessary to examine the experiences of musicians who use the virtual instruments. The following sections detail the methods used to do this.

3.4.2 User Studies

In order to examine how well the virtual musical instruments meet their design criteria and how successful they are overall, it is necessary to study the experiences of musicians who use them. Just how to approach this is a difficult issue. 'Traditional' Human Computer Interaction (HCI) approaches have focused on measuring user performance when carrying out various well-defined tasks such as navigating a website or entering figures into a spreadsheet. Software designed to facilitate musical expression presents a problem in this context as it is difficult to formulate tasks to assign to users which are measurable but also meaningful (Wanderley and Orio 2002). If the aim was to produce a general-purpose musical instrument for performing traditional music, then evaluation would be simpler. Tasks such as playing a scale, trilling, etc. could be assigned and measurements made to ascertain how successfully users were able to execute them. The benefit of this approach is that it would be possible to somewhat objectively compare two different musical instruments in terms

²The version control software Concurrent Versioning System (CVS) was used during software development to track changes at every stage of the design process.

of this restricted definition of playability. However, where the instrument is intended to create new and unusual sounds – to explore new languages of composition and performance – this approach is problematic. Part of the rationale for creating these instruments is that they disrupt musicians' ways of thinking about music so that they are stimulated to try new ways of playing and composing. Attempting to determine how effective they are at facilitating performance of current styles of music might be interesting, but it would not facilitate learning about how to design instruments which encourage divergent thinking.

For this reason evaluation methods in this research have been qualitative. In 2005, the National Science Foundation sponsored a workshop examining creativity support tools which brought together a number of researchers in the field. Of particular relevance for this work was a breakout session on methods of evaluation which arrived at the conclusion that a complete understanding of creativity support tools requires the use of several evaluation techniques. The workshop report details suggests that:

“There is a FAMILY of evaluation techniques which must be brought to bear in order to CONVERGE on the key issues involved in the study of Creativity Support Tools.” (Hewett et al. 2005, p. 14, capitalisation in original)

In other words, researchers need to draw on a range of methodological tools and techniques in order to build a rich understanding of the impact of tools for the support of creative work on users' work. Because the use of simulated physical models as a mediating mechanism in virtual instruments is a new technique, gathering quantitative data on how successful they are is problematic. Getting numbers would be easy, but understanding what they meant would not. Thus, at this stage of their development qualitative methods were used with the expectation that the findings may go some way to facilitating quantitative measurements as practical knowledge of this type of virtual instrument develop.

A series of evaluation sessions was therefore conducted in which a number of expert musicians used the virtual instruments and discussed their experiences. Before outlining the procedures in detail however, the nature of the researcher's past experience in the domain under study will be discussed.

If the behaviour of participants is going to be analysed without adopting an *a priori* theoretical framework (as is the case when using the Glaserian approach to grounded theory development (section 3.4.2)), then a skilled practitioner in the area under study will be in a position to identify and understand what is going on. Crabtree (2003), discussing ethnomethodological studies, argues that researchers need to have (or develop) a degree of competence in the field under study. The researcher who has what Crabtree calls 'adequate mastery' of the domain, “can recognize as members recognize *what* is going on in the phenomenal field of practical action under study and *how* it is getting done.” (Crabtree 2003, p. 81, italics in original)

This was very important in this research. It was very helpful to be able to interact with the musicians who participated in their own 'language'. Music is a highly specialised field with a developed vocabulary and historically informed practices. The fact that the principle researcher had a degree in music performance and several years of professional experience was very significant, as it meant that musicians did not have to explain the fundamentals of their art and craft to a researcher who had no experience of it. Conversations with and feedback from the musicians could therefore be at a high level, and musicians could speak freely without feeling the need to simplify their expressions or translate them into lay-persons' terminology.

Participants

All participants were professional musicians based in Sydney who all had a minimum of five years professional experience. Because the software is designed to work with the trombone, initial participants were trombonists, in order that any technical issues could be ironed out without introducing undue complexity. Players of other instruments were subsequently recruited for the study; however, the majority were wind or brass instrumentalists as the software was designed for monophonic, pitched instruments.

Because the virtual instruments are intended to encourage new techniques for musical expression, an important consideration when recruiting musicians was that they had an interest in new music and were open to the potential use of new technologies in their music-making.

Some participants had greater experience with improvised music and others with notated music. Numbers were approximately balanced to ensure we took account of these different perspectives.

In line with UTS ethics procedures, it was made clear that participation was entirely voluntary and that video and interview data would be kept secure and not supplied to third parties. It was also made clear that participants did not have to answer any of the interview questions if they did not wish to.

Each evaluation session took around two hours.

There were seven participants in total. Due to the degree of expertise of the participants and the in-depth nature of the evaluation, this was a sufficient number to provide detailed insight into the experiences of expert musicians with the virtual instruments. Note that in qualitative research the emphasis is on generating, rather than validating, theory (Glaser and Strauss 1967). It is not claimed that this study proves or disproves the existence of particular phenomena. Rather, this research was exploratory and intended to provide detailed insight into the experiences of the specific musicians who participated in the evaluation, and to generate theories consistent with what was observed. The research was rigorous, in that elements of the theory were clearly linked to data gathered during the evaluation, but no claims are made that the theory is applicable to the broader musical community or even to other

musical instruments. It is hoped, however, that this research will provide a sound basis for future research which may attempt to rigourously and more broadly *validate* the concepts and relationships carefully uncovered in this study. Such validation would be likely to involve larger numbers of musicians using virtual instruments in a simplified and more controlled context.

Data collection

A pilot study investigating the audience experience of a sound-controlled interactive instrument was conducted in the Beta_Space exhibition space at the Powerhouse Museum Sydney in September 2006. This study differed from the evaluation by experts in that the participants were not selected on the basis of their level of musical expertise but were randomly selected Powerhouse patrons. By coincidence two of the participants did have a degree of musical training but the majority were musical novices. The intention was to examine the patrons experiences with the virtual instrument and to trial evaluation techniques that would later be used with the expert performers.

As the focus was on how the participants used the virtual instruments, they were simply asked to play with them in any way they liked and to stop playing when they felt they had had enough.

The video-cued recall technique (Omodei and McLennan 1994) was used to get a rich understanding of how these users experienced the virtual instrument. This was a two-step process. First, each participant's session playing with the virtual instrument was video recorded. They were then taken to another room and asked to watch the video that had just been recorded and, while watching the playback, to verbally report what they perceived or thought during the interaction. That is, the video was used as a cue to prompt the participants to recall what was going through their mind as they were playing with the instrument. These video-cued recall reports were also video recorded. Finally, users were asked a series of open questions about their experiences with the software.

For each participant then, the study provided video recordings of the participant:

- playing with the virtual instrument;
- retrospectively describing what they perceived and thought during their session with the virtual instrument;
- answering a series of questions about their interaction with the instrument.

While it was successfully used in the pilot study, there are several reasons why it was decided that the video-cued recall technique would not be suitable for the expert evaluations. First, it was felt that the time required of each participant could quickly become excessive. For example, suppose the musician played with the virtual instrument for 75 minutes. When finished, they would then have to watch a video of the

entire 75 minute session again, while recollecting in detail what was going through their mind at every moment. Thus, the evaluation and reflection would consume two and a half hours. Further follow-up interview questions would add extra time. While musicians may have been inclined to be generous with their time, it was necessary to be mindful of making unreasonable demands.

Second, because of the duration of the evaluation sessions, the time between the musicians' use of the virtual instruments and subsequent retrospective report would have been significant, leading to a reduction in accuracy of recalled thoughts and perceptions.³

In order to avoid these problems it was instead decided to gather a mix of concurrent and retrospective reports from the musicians (Ericsson and Simon 1993). When using the concurrent approach, the idea is that the musicians continuously verbalise what is going through their mind as they play the virtual instrument, reducing the time between thought and verbal expression to a minimum. However, asking musicians to generate fully concurrent think-aloud reports presented practical problems because brass musicians are unable to speak (intelligibly) and play their instrument at the same time. A sensible compromise was to ask the musicians to verbally report what they were thinking and perceiving as frequently as they were able during their time using the software. This means that they were effectively providing a large number of smaller retrospective reports as they played for a time, commented on what was happening, played some more, made further comments and so on.

A disadvantage of this approach is that having to continually stop to comment on what they were doing may have prevented the musicians from becoming fully immersed in the compositional task. That is, they may have not attained the 'flow state' which, according to Csikzentmihalyi (1996), is so important in creative work. However, there is significant evidence that the think-aloud procedure has minimal impact on the cognitive processes of research subjects engaged in problem-solving tasks (Ericsson and Simon 1993). It seems reasonable to assume that this extends to creative tasks also. In addition, professional musicians are likely to be adept at moving in and out of their optimal 'performance' frame of mind as this is part of their daily work. In any event, there appeared to be no less-disruptive alternative for gathering detailed information about the musicians' experiences and thoughts on using the virtual instruments.

It is important to note that in addition to gathering information about what the musician was thinking and experiencing as they used the virtual instruments, the musician's opinions on the software and suggestions for how it could be improved were actively solicited. As experts in their field, it was hoped that the musicians would be able to provide insight into the nature of the virtual instruments, their potential uses, limitations and areas for improvement. The intention was that the musicians would become engaged with the design process and in a sense become co-designers.

³The gap between action and retrospective report is the duration of the evaluation. That is, a 75 minute evaluation means a 75 minute time gap between action and report.

Where possible, suggestions by the musicians were implemented immediately in the software in order to help clarify and confirm the suggestion and gather immediate feedback. As such, the format of the evaluation was flexible. There was a standard procedure (described in section 3.4.2), but when interesting issues arose this was varied. Because the emphasis of this study is on theory generation rather than verification, rich data gathering was prioritised over consistency of procedure. The process was more akin to a user dialogue than usability testing (Buur and Bagger 1999).

The participants' interaction with the prototypes and the interview was recorded using a digital video camera. Video data has been stored in accordance with Creativity and Cognition Studios' ethics procedures.

In order to provide additional perspective, an observer (not the author), experienced with interactive artworks, was present during evaluations. The observer noted the actions and comments of the musician. At times the observer participated in discussions. Following the evaluation session, the observer and researcher discussed the session and compared initial impressions of the themes that emerged. The notes provided additional data for analysis and a degree of triangulation.

Study procedure

Set-up The investigation took place in the Creativity and Cognition Studios' main studio. The computer output was displayed on a large rear-projection video screen and sound output was through Genelec 1029A studio monitor speakers coupled with a 7050A sub-woofer. Sound input was via a Rode NT1-A microphone connected to the computer through a small mixing desk.

Introduction & user instruction Participants were asked to play with the virtual instruments one at a time. Some musicians preferred to work with the software without receiving a detailed explanation of how it works. However, on most occasions the musicians were told:

- That the virtual instrument responds to their sound via the microphone;
- That the screen will display a visual representation of the virtual instrument; and
- That the virtual instrument will produce sounds.

If desired they were given further details, including:

- The fact that their sound exerts force on the on-screen virtual instrument;
- What effect changing pitch, volume and timbre will have on the instrument; and
- How the movements of the virtual instrument create sounds.

Investigation of the users' experience Participants were asked to use the virtual instruments to make music in any way they liked. The focus of the investigation is on what they are able to do with the instruments, what impact using the instruments has on their music making and any suggestions for improvements, so the musicians were not asked to perform specific musical tasks. They were, however, asked to verbally reflect on their experience with the instruments using a variation of the 'think aloud' approach (Lewis 1982, Ericsson and Simon 1993).

While playing with the instruments, the musicians were asked to report and reflect on their experience.

The instruction to participants was:

"We would like you to report what you perceive or think while you are interacting with the software. We would like to get as complete a report of what is going through your mind as possible, so please don't worry if what you say is inconsistent or incomplete - just report what is going through your mind at the time. Don't feel that you must break your concentration to make a report. Please try to give reports at a time that feels natural and appropriate to you."

Following the familiarisation process, the musicians were asked to prepare and perform a short piece using the virtual instrument. The instruction was:

"We would like you to prepare and perform a piece of music with the software. You can use manuscript paper if you wish to jot down any ideas but please don't feel that you have to. There are no constraints on the style or duration of the piece; we're just interested in how you use the software and your experiences in so doing. Please remember that we are evaluating the software not your performance.

If you would like any aspect of the software adjusted let me know and I will do what I can to accommodate your request."

Interview Participants were next asked a series of questions relating to their experience with the virtual instrument. The questions were:

1. Tell me about the piece of music you played.
 - (a) Why did it have these characteristics?
2. Do you have some comments about how easy or hard it was to write for or perform with the software?
3. Do you have some comments about the sound produced by the software?
4. Do you have some comments about the visual display?

5. While you were interacting with the work, did you become aware of any particular characteristics of your playing?
6. Did you play differently today than you normally would?
 - (a) In what ways?
 - (b) Can you say why?
7. Do you have any suggestions or proposals for how we might improve or extend this software?
8. Can you think of any uses for this software?
9. Is there anything I should have asked you?
10. Do you have some comments or questions about what we've done today? Is there anything we should have done differently?

Questionnaire Finally, participants were asked to complete a questionnaire, ranking their level of concurrence with the following statements on a five-point Likert scale from strongly disagree to strongly agree.

- The instrument responded to the music in a way that I would describe as 'natural'.
- The instrument's behaviour was consistent. (ie. It responded to identical notes in the same way.)
- The instrument was conceptually simple.
- The instrument allowed me to create complex musical and visual effects.
- I found the instrument interesting and engaging.
- I felt in control of the instrument.

Demographic details Participants were asked the following in order to gather basic demographic information:

- Do you have any formal qualifications? If so, what are they?
- Can you summarise your professional experience in a few sentences?
- Would you be prepared to tell us your age or age-range?

Data analysis

The video-recordings of the musicians playing the virtual instruments and talking about their experiences were a very rich source of data. A challenge was to identify consistent themes and patterns in order to make sense of this information. Techniques from grounded theory method (Glaser and Strauss 1967, Glaser 1978) were therefore used to code and analyse the data gathered. This method was a good fit for this purpose because it facilitated the generation of theory closely tied to the evidence from rich qualitative data. At a high level, the basic steps of the grounded theory analysis process as applied in this study were:

1. Transcribing the evaluation sessions.
2. Open coding: that is, identifying and labelling incidents in the data (including non-verbal data). During open coding, the researcher continually asks a series of questions, such as, "What is this data a study of?", "What category does this incident indicate?" and "What is actually happening in the data?" (Glaser 1978, p. 57). This is done line by line, coding each sentence. As coding progresses, incidents are constantly compared with one another to identify similarities and differences.
3. Memoing: as ideas emerge regarding the codes and their relationships during coding, the researcher stops to make a note. Memoing aids the process of linking the descriptive codes into theory.
4. Sorting: memos are sorted and arranged in order to identify core issues and their relationships with one another and thus build theory which is grounded in the gathered qualitative data.

The software Transana (Woods and Fassnacht 2007) was used to facilitate this process. Transana is open source software for conducting qualitative analysis of video and audio data. In particular, it facilitates the annotation of video with keywords and comments as well as tools for grouping related video clips together based on these keywords. This enables the researcher to build hypothesis and gather evidence through detailed examination of both the verbal comments and responses of the musicians and their behaviour while using the software.

3.4.3 Grounded theory method

The term 'grounded theory' refers to theory inductively generated by analysis of data. The method for generating such theories proposed by sociologists Glaser and Strauss (1967) is also often colloquially referred to as 'grounded theory', so to avoid confusion, in this thesis 'grounded theory' refers to the theory and 'grounded theory method' refers to the method. While the method is agnostic with regards to the type of data analysed, in general it is applied to qualitative rather than quantitative data.

The grounded theory method is a slightly more formal version of methods long used in ethnographic studies to make sense of the large tracts of complex qualitative data gathered in the field.

The grounded theory method is aimed at generating, rather than testing, theory. In positivist or post-positivist scientific research the emphasis is on testing hypothesis which are derived from existing theories. Glaser and Strauss (1967) argued that theory generation was comparatively neglected in scientific research, and proposed techniques to facilitate and enhance this process. The grounded theory methods aim to provide a rigorous approach to generating theory intimately connected with the data from which it is drawn. By focusing on maintaining the connection between theory and data, Glaser and Strauss argue the method helps ensure that generated theories retain relevance to the situation under study and applicability in the real world.

The key aspects of the grounded theory method are as follows:

- Theory generation and data gathering proceed in parallel. Sampling is driven by the need to further develop the emerging theory, not by the desire to enhance the generalisability of the findings.
- Use of coding techniques to facilitate analysis of qualitative data.
- The technique of memoing is used to aid the generation of theory.

Data gathering and sampling

In the grounded theory method, the data gathering and theory generation processes inform one another. It should be obvious that whatever data is gathered will influence the emerging theory; in fact the grounded theory method's *raison d'être* is to maintain the relationship between theory and data. It might not be so obvious that the emerging theory will also inform data gathering strategies. As the theory develops, questions arise which suggest further avenues of exploration. For example, a theme that emerged during analysis of earlier interviews conducted during this research was that of control. The degree to which musicians were able to control the virtual instruments seemed to be important. Up to that point, the musicians that had participated in the study were predominantly from an orchestral background. The desire to understand how much their emphasis on the issue of control was due to this background led to the recruitment of musicians with a background in free jazz, who brought an additional perspective to the investigation and helped to flesh out the control issue.

This approach is in contrast with sampling approaches in research in the positivist tradition, in which attempts are made to ensure that research samples are representative so that research results are generalisable. Grounded theory aims to produce theory which accounts for the specific situations under study. The level of general applicability of a grounded theory is something that is discovered as the research

progresses, possibly over several research projects. The intention is that the theory in its current state will always account for the situations (the individual studies) from which it was derived. General applicability is a secondary, but nonetheless important, goal.

Another key point here is that analysis commences as soon as the first data is gathered. It is not the case that a full set of evaluations are conducted before analysis begins. By analysing as data gathering progresses, selective sampling can ensure that data relevant to the emerging theory is gathered.

Coding

In grounded theory method, theory is generated from the data by the analytic processes of coding and memoing. 'Coding' in the grounded theory method refers to the detailed examination of the gathered data and the identification and labelling of incidents. When dealing with verbal data from interviews, a common approach is to transcribe the interview session and mark up the transcript with codes during analysis. A number of software packages facilitate this process.

The process of coding is very simple, but there are a number of guidelines can help guide researchers, particularly those in the early stages of a grounded theory study. One of the co-originators of the grounded theory method, Barney Glaser, offers four 'rules' for coders (Glaser 1978). These are:

1. Ask questions of the data.
2. Analyse the data line by line.
3. Analyse your own data.
4. Interrupt coding to memo (make a record of) ideas.

Ask questions of the data Rule one refers to a series of questions that Glaser recommends analysts keep in mind as they examine the data. These questions are intended to prompt the researcher to keep an open mind and avoid imposing their pre-conceived notions on their analysis of the data.

The first question is, "What is this data a study of?" (Glaser 1978, p. 57). In asking this high-level question, the researcher is prompted to continually consider whether the initial focus of the study is in fact appropriate to the data that has actually been gathered. It is possible, or even likely, that the research will change direction as data gathering and analysis proceed. The intention is that the researcher is open to this and does not attempt to 'force' the data to align with his or her preconceptions.

Glaser's second question is, "What category does this incident indicate?" (Glaser 1978, p. 57). Placing incidents in the data into categories or groups with similar incidents is at the core of grounded theory analysis. In so doing, the analyst is constantly comparing one incident with another to see if they can be grouped together in some way or if a new category needs to be created.

Practically, this grouping is usually achieved by labelling incidents in the data with keywords or short phrases. If a musician in this study were to say, “I felt uneasy because the software was continually judging my playing” then the researcher might label this sentence, ‘judging playing’. (In fact, such a sentence would be likely to have several labels attached.) As analysis proceeds, it is possible that other incidents in the data will appear to the researcher to also relate to the software judging the musician’s playing. By comparing these incidents with the previous one the researcher will decide whether to label them in the same way, or, if they are sufficiently different, to create a new label for them.

The third and final question Glaser recommends researchers ask themselves as they analyse is, “What is actually happening in the data?” (Glaser 1978, p. 57). That is, what are the underlying processes which can account for what is happening in the specific situation reflected in the data? By continually considering this question, researchers build theories that explain or otherwise account for what is happening in the data.

Analyse the data line by line The second of Glaser’s rules for coding requires that researchers look at data (interview transcripts in most cases) in great detail: line by line. The intention is that by examining the data at this micro level, the researcher is forced to fully take account of what is actually in the data. It provides the added benefit of allowing the researcher to clearly link the resulting theory to specific elements of the data itself.

A common theme in the grounded theory method is the need to code at a detailed level in the initial stages of analysis.

“Thus grounded theory coding does not go with wholistic reading of an interview or field note for the “overall” conceptual impression. This impressionism neglects the details, and their comparative meanings for concept generation. The resulting concept is not rigourously generated and verified as to pattern and fit. It is what it is, just a juicy overall impression, which seems to fit but has no systematic relevance to a theory and no systematic generation.” (Glaser 1998, p. 140)

The intention is that theory generated by line-by-line coding is clearly linked to incidents in the data to a very specific degree. The danger of a holistic, impressionistic approach is that the analyst will unintentionally ‘force’ the data to fit their preconceived notions, ignoring incidents in the data that do not fit.

Analyse your own data The process of coding in the grounded theory method does not involve taking a pre-existing set of codes and applying them to data gathered during an investigation. Rather, the codes are created during analysis as the analyst compares and labels incidents in the data. For this reason, it is not desirable

(or even possible) for grounded theory analysis to be 'outsourced' (Glaser 1978). In this research, all analysis was conducted by the author.

Interrupt coding to memo (make a record of) ideas. While coding is a key part of grounded theory analysis, it is not sufficient in itself to produce grounded theory. Coding - the labelling of incidents in the data - facilitates analysis by grouping together incidents which are similar in some way. However, merely labelling these incidents does not produce theory. In order to produce theory which accounts for or explains the data in some way, it is necessary for the analyst to consider relationships between labelled incidents in the data. This process is ongoing. That is, during analysis the analyst continually considers how labelled incidents relate to each other and documents their evolving understanding of the situation under study by recording their thoughts in 'memos'. The format of memos is entirely up to the individual researcher. The intention is to capture the researchers thoughts quickly during analysis so that insights are retained. As such, memos are not expected to be grammatically correct or understandable to anyone other than the researcher involved. The memoing process is in effect a first step towards writing up findings from the study (Charmaz 2006).

"Memos vary from being a "jot" of a few words or a jot to write a memo later all the way through theory bits to a 5 to 10 page paper. In short they are anything that captures the meaning of conceptualized ideas, where they go theoretically and for theoretical sampling. There is no formalization of them. They just flow out with total freedom." (Glaser 1998, p. 178)

Saturation

A final issue for those using the grounded theory method is when to stop gathering and analysing data. Unlike quantitative techniques, the grounded theory method does not have rules which specify required sample sizes in order to produce results of a certain measurable significance. However, there are guidelines that help the researcher determine the point at which additional sampling and analysis is unlikely to lead to further insight. The non-numerical measure of the degree to which the researcher has reached this point is called 'saturation'. By Glaser and Strauss's definition:

"*Saturation* means that no additional data are being found whereby the sociologist can develop properties of the category. As he sees similar instances over and over again the researcher becomes empirically confident that a category is saturated." (Glaser and Strauss 1967, p. 61)

In other words, when the researcher finds that new data is not yielding new categories, properties of those categories or insight into the relationships between them,

then saturation has been reached.

3.5 Conclusion

In this section, I have described and justified the research approach and specific methods used in this project. I have stated that the research takes a *constructionist* position on the nature of knowledge and an *interpretivist* theoretical perspective (section 3.2). The research approach draws on both Design Science and Action Research in that:

- The researcher actively participates in the area under investigation.
- Successful and unsuccessful virtual instrument design strategies are identified through practice.
- The project aims to contribute to both practice and theory.
- A iterative approach of action and reflection is applied.
- Virtual instruments are designed and used to both provoke current practice and help understand it.
- The virtual instruments will be evaluated in terms of how well they address musicians' needs.
- The virtual instruments will be designed using an iterative design methods.

I have described a software development method based on participatory design techniques, prototyping and agile methods and outlined how this will help ensure the virtual instruments retain practical relevance to the musicians involved.

In addition, I have described the importance of making the design criteria which guides development of the virtual instruments explicit so that it is possible to examine:

- how well they met these criteria;
- the relationship between the criteria manifest in the virtual instruments and the experiences of musicians who use them;
- how successful or unsuccessful the virtual instruments were overall.

Finally, I have described the structure of a user study which will examine the experiences of musicians who use the virtual instruments and outlined how the grounded theory method will be used to analyse the qualitative data gathered.

In the following chapter I will describe the virtual instruments that were designed and identify the design criteria which guided their development.

Chapter 4

Virtual Musical Instruments Based on Simulated Physical Models

4.1 Introduction

This chapter describes the virtual instruments developed as part of this project in detail. In chapter 2 a number of criteria for building virtual musical instruments were identified, along with a technique for using physical models as a mapping layer between gesture and computer generated sound and visuals. In chapter 3, an approach to research in this creative domain drawing on action research, design science, participatory design and human-computer interaction was described. In this chapter, I present the virtual instruments that were designed along with the design criteria which guided their development. The descriptions of the virtual instruments in this chapter are predominantly technical, but I also briefly describe the ‘character’ of each virtual instrument in order to illustrate how the technical characteristics relate to aesthetic concerns.

So that readers of this thesis are more readily able to understand how the interactions between musician and virtual instruments are structured and the creative context within which they were created, a DVD is provided. The DVD includes a short video in which I demonstrate each of the virtual instruments. In addition, the disc contains video recordings of performances featuring the virtual instruments. The contents of the DVD are listed in appendix I.

4.2 Design process

As discussed in chapter 3, a participatory design approach was used. As such, the virtual instruments were developed in close collaboration with a musician and composer, Ben Marks. Ben is a trombonist and composer based in Brisbane, Australia who has extensive experience in a wide range of musical genres, but particularly in contemporary music. He has played with the contemporary music ensemble ELISION and also with ensembles throughout Australia and overseas including Ensemble Offspring, Libra Ensemble, Southern Cross Soloists and the Singapore, Adelaide, Western Australian, Melbourne and Queensland symphony orchestras. He has toured with ELISION to Poland, Korea, England, Norway, Germany, Austria and New Zealand. Ben regularly gives solo recitals in Brisbane and throughout Australia and has received a number of scholarships and awards including grants from the Australia Council, Arts Queensland, and Ian Potter Foundation along with academic awards from The University of Melbourne and Griffith University. In addition, Ben teaches trombone and runs various ensembles at the Queensland Conservatorium of Music, Griffith University and the Queensland University of Technology.

The collaboration involved the simultaneous design of virtual instruments and composition of music for those instruments. Ben was primarily responsible for composing the music and I was primarily responsible for developing software. However, as discussed in section 3.3.3, the process was truly collaborative and both of us contributed to all facets of the artistic outcome to some degree.

This work resulted in the development of two virtual instruments and associated compositions entitled *Partial Reflections I* and *Partial Reflections II*. In addition, a third virtual instrument with a simplified interaction scheme, *Spheres of Influence*, was developed (see section 4.7). Each of these instruments will be described in detail in the following sections.

The compositions and virtual instruments were premiered in concert at the Sydney Opera House Studio in August 2006. Since then, they have been featured in several concerts in Sydney and Brisbane. *Spheres of Influence* was exhibited at Beta_Space in the Powerhouse Museum, Sydney in September-October 2006.

4.3 Design criteria

A number of design criteria guided the development of the virtual instruments. In this section a set of design criteria which guided development will be introduced and discussed with reference to specific characteristics of the three virtual instruments described previously. Reflective diaries (web logs or 'blogs') were used during the creative process to document these criteria as they developed. As the reader will see, several of these criteria were derived from the literature reviewed in chapter 2. As outlined in chapter 3, making the criteria which guided development explicit enables us to examine:

- how well the virtual instruments embody the criteria;
- the relationship between the criteria manifest in the virtual instruments and the experiences of musicians who use them; and
- how successful or unsuccessful the virtual instruments were overall.

The criteria are:

1. The virtual instruments should be able to interact with a wide variety of acoustic instruments.
2. The virtual instruments should respond in a way that seems natural.
3. Virtual instrument response should be consistent.
4. Virtual instruments should be simple but allow skilled musicians to create complex effects.
5. The virtual instruments should be interesting, engaging and motivate the musician.
6. The musician should feel in control of the virtual instrument, but it should retain the ability to surprise.

7. The virtual instrument should encourage a playful, exploratory approach, especially in new users.
8. The relationship between live sound, the behaviour of the virtual instrument and the resulting sounds and visuals should be apparent to observers (eg. audience members).

In this section each of these criteria will be introduced and discussed.

Criterion 1: The virtual instruments should be able to interact with a wide variety of acoustic instruments. The fact that the virtual instruments are 'sound controlled' means that they will respond to any sound. It was felt that this flexibility was an important characteristic which meant that the virtual instruments could be used in a variety of musical contexts. The compositions by Ben Marks include a wide range of instrumental and vocal sounds and the intention was that the virtual instruments would respond in interesting and meaningful ways to this diverse range of sounds. This characteristic also meant that the virtual instruments are able to be used by many other instruments.

Criterion 2: The virtual instruments should respond in a way that seems natural. It was felt that an advantage of using simulated physical models to map between live sounds and computer generated sounds and visuals is that the resulting virtual instrument has a 'naturalness' about it. This term 'natural' is a little ambiguous. What is meant is that the instruments responded in ways that were consistent with our expectations based on our experience of objects in daily life. Because the instruments moved in ways that we would expect physical objects to move their response was not unexpected or outlandish. The intention was that musicians would feel an intuitive connection between the acoustic sounds they produced and the sounds and visuals generated by the virtual instrument.

Criterion 3: Virtual instrument response should be consistent. This criterion is related to criterion 2 above (instruments should respond in a way that seems natural). Just as it would be expected that an everyday physical object would respond in the same way to two identical physical actions, so should the virtual instruments. That is, two perceptually identical notes should appear to have the same effect on the virtual instrument. In order that Ben could use the instruments in performance, it was necessary for him to have confidence that the virtual instruments would behave in broadly consistent ways. An instrument that was completely inconsistent would not be controllable (see criterion 6 below).

Criterion 4: Virtual instruments should be simple but allow skilled musicians to create complex effects. A guiding principle for the design of the virtual instruments was that they be as simple as possible (but no simpler). With computer music

(and software in general) there is often a temptation to add functionality. This temptation was resisted for a number of reasons. Firstly adding extra functionality tends to decrease reliability, and as the software was used for live performance reliability was a high priority. Secondly, it was felt that building virtual instruments of greater complexity tended to decrease the transparency of the interaction, particularly to the audience (see criterion 8). Finally, it seemed that it was not necessary to make the virtual instruments more complicated. The simplest instruments seemed capable of creating complex visual and sonic results. The structure of the virtual instruments seemed more important in this regard than their complexity.

Criterion 5: The virtual instruments should be interesting, engaging and motivate the musician. It should be obvious that any instrument designer would want their instruments to engage and motivate musicians. This criterion provides the context for all the other criteria.

Criterion 6: The musician should feel in control of the virtual instrument, but the it should retain the ability to surprise. While it was intended that musicians using the virtual instrument should feel that they had a degree of control over it, it was also felt that for the interaction to be interesting and meaningful the virtual instrument should be able to surprise the musician to some extent. That is, while its behaviour should be predictable at a high level, its structure should be such that it may be able to get into states that generate music or visuals that the musician would not necessarily have expected.

In addition, it was hoped that by using the virtual instrument the musician might discover something about their technique or gain insight into some aspect of their music making. A sensitive and responsive virtual instrument has the potential to make detailed aspects of the musicians' technique more apparent. While this could be annoying, it may also stimulate the musician - perhaps to work on improving technique, or more interestingly, to explore the limits of their instrument in ways that they had not previously considered.

Criterion 7: The virtual instrument should encourage a playful, exploratory approach, especially in new users. It was hoped that the virtual instruments would encourage musicians to explore new musical ideas in a non-judgemental, playful way and that musicians would be stimulated to consider questions such as, 'What does it do if I play...?' and 'How can I make it...?'

Criterion 8: The relationship between live sound, the behaviour of the virtual instrument and the resulting sounds and visuals should be apparent to observers. Both participants in the design process (composer/musician and technologist) had had previous experience, as performers and audience members, with

music software which was nominally responsive to live audio but which from the audience's (and often the musicians') perspective seemed largely disconnected from it. A very important criterion for the virtual instruments was therefore that the connection between the musician's actions and the behaviour of the virtual instrument should be clearly apparent to audience members and, of course, the musician. This is consistent with the position articulated by Fels et al. (2002) (see section 2.2.2), that an instrument's expressivity is largely determined by its degree of transparency to both audience and musician.

4.4 Overview of the virtual instruments

Having outlined the criteria which guided the design process, the virtual instruments will now be described in detail. There are three virtual instruments: *Partial Reflections I*, *Partial Reflections II* and *Spheres of Influence*. While they differ in important ways, they also share a number of common characteristics:

- The virtual instruments are controlled with acoustic sounds via a microphone.
- The visual output is a representation of the state of the virtual instrument. From the musician and audience's perspective, this visualisation *is* the virtual instrument.
- Acoustic sounds are mapped to forces which act on a virtual physical model.
- The physically plausible movements of the physical model provide parameters for sound synthesis.

These common characteristics will be described in more detail in the following sections.

4.4.1 Controlling the virtual instruments

An important requirement for this project was that the virtual instruments that were created should be usable by a wide range of instruments. A number of 'hyperinstruments' such as the hypercello (Machover 1992) and the metasaxophone (Burtner 2002, Burtner and Serafin 2002) (see chapter 2), enhance acoustic instruments by adding sensors, buttons or switches which allow physical gestures to trigger and/or control computational processes. Myatt (2002) has gone to considerable lengths to develop a wireless system which enables a saxophonist to generate control gestures without disturbing their instrumental technique. This opens up considerable possibilities, but for this research the downside is that great effort needs to go into designing an interaction strategy for each type of instrument (or perhaps even for individual musicians). Thus, in order that the software and music for this work would be as available to the wider musical community as possible, it was decided that no hardware other than a computer, microphone and speakers should be required.

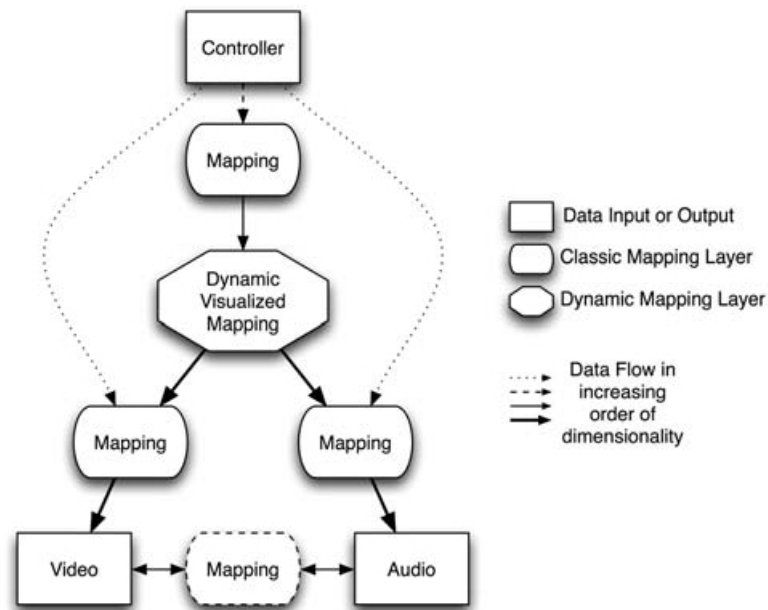


Figure 4.1: Block diagram of a virtual instrument using a dynamic independent visual-mapping layer (Momeni and Henry 2006, p.50).

4.4.2 Physical model as mapping layer

In chapter 2, we discussed the use of physical models in a mapping layer between gesture and computer generated audio and video. In this approach, a simple physical system is modelled in software and manipulated with gestures of some kind. The gestures are mapped to virtual forces which act on the physical model which moves in response in ways that are physically realistic. These movements can then be used as parameters which control both video and audio synthesis. The diagram in figure 4.1 is a high-level view of this process.

In the works developed as part of this project, the 'controller' in figure 4.1 is the live audio from the performer which is passed through an audio analysis module — the *fiddle~* object bundled with the Pure Data programming environment (Puckette et al. 1998). This module examines the incoming audio signal and estimates the following information:

- The pitch of the live sound.
- The amplitude (volume) of the live sound.
- Peaks in the harmonic spectrum. That is, the strength and frequency of the most prominent harmonic components of the sound.
- Identification of note articulations or 'attacks'.

This dynamic stream of data, derived from the acoustic sounds captured by the microphone, is the means by which the musician manipulates the virtual instrument.

In other words, the musician uses their sound to influence or control the virtual instrument. Note that none of the virtual instruments uses all of the above parameters, but all make use of at least the pitch and amplitude of the live sound.

The above parameters are used in different ways. As discussed in chapter 2, the gestural function of these parameters can be classified as *excitation*, *modification* or *selection* gestures (Cadoz 1988, Cadoz and Wanderley 2000). In all of the virtual instruments, the fundamental pitch acts as a selection gesture and amplitude as excitation. As the mechanism of selection and excitation is similar across all three instruments, the basic mapping strategy will be described here.

For all the virtual instruments the overall interaction is structured as follows:

- The musician exerts ‘force’ on part of the virtual model (a mass-spring physical model) by playing into a microphone. Because the model incorporates the laws of physics, it responds in a physically plausible way to this musically generated force. That is, it moves in response in a way that seems natural.
- A representation of the physical model is shown on-screen during performance and is visible to both performer and audience. From the performer and audience’s point of view, the physical model *is* the virtual instrument.
- As the physical model moves, sounds are produced. This means that audience and performer hear computer-generated sounds (controlled by the physical model) as well as acoustic sounds from the traditional instrument.
- The musician only interacts with the software via the microphone. There are no additional buttons, mice or sensors attached to the acoustic instrument. This means that the software will work with any musical instrument or vocals.

Each of the instruments has a mass-spring model at its core. The open-source module ‘Physical Modelling for Pure Data’ or ‘pmpd’ (Henry 2004a) provided this functionality. Using pmpd, the programmer creates a model by creating mass objects and linking them with link objects. The mass objects have properties such as weight and initial position in 3D space. Link objects, which have properties such as rigidity and damping, link two masses. Once the model has been constructed the simulation can be started. Applying force to a mass object will cause forces to propagate through the network of masses and links as pmpd does the necessary calculations to apply the appropriate physical laws.

There are two kinds of masses in the simulation: fixed and mobile. Fixed masses have a set position from which they do not move. The simulation is essentially unbounded in space, and there is a fixed ‘camera position’ from which the mass-spring model is viewed. Thus to ensure that the model is always visible it is necessary to use at least one fixed point to anchor the rest of the model. The mobile masses can move around in three dimensional space. They are constrained in their movements by the links which bind them to other mobile and/or fixed masses. The movement of

the mobile masses generates the data that is used to control the audio and visual synthesis.

In the *Partial Reflections* and *Spheres of Influence* instruments the amplitude of the live, acoustic sound is mapped to the quantity of force exerted on a mass in the model (figure 4.2). The louder the live sound the greater the force. Because the structure and properties of the model are different in each instrument, the response will be different but the link between volume and force remains. Because the resulting movements of the physical model result in sounds being produced, the volume of sound acts as an 'excitation' gesture.

Volume was chosen as an excitation gesture because we felt it gave a superior sense of control to the musician. It was preferred over pitch or harmonic structure as an excitation gesture because there seemed to be an intuitive connection between the force required to increase volume on the trombone and increasing force upon the virtual instrument.¹ The linking of the players physical energy to simulated energy in the physical model is an example of using 'energy injection' to make sound. Experiments by Hunt and Wanderley (2002) indicate that players are more engaged by virtual instruments which use this technique (see section 2.3).

In acoustic instruments, the volume of sound produced is directly proportional to the amount of energy supplied by an excitation gesture. That is, the more energy the musician supplies to the instrument by blowing it, plucking a string, etc. the louder the instrument will sound. Thus, the volume of the acoustic sound is a reflection of the amount energy supplied by the musician. By linking this to force acting on the simulated physical model, in effect we visualise the excitation gesture used on the acoustic instrument and, in addition, give the musician the feeling they are using the same excitation gestures on both acoustic and virtual instruments. This technique has two benefits. Firstly, it aids the performer in understanding and controlling the process. Secondly it helps the audience understand and appreciate what's going on; it makes operation more transparent. Two of the virtual instruments (*Partial Reflections I* and *Spheres of Influence*) used volume as a 'continuous' excitation gesture and one used it as an 'instantaneous' excitation gesture.

In all the virtual instruments pitch is used as a 'selection' gesture. Recall that a selection gesture is used to choose between a number of similar elements of an instrument (Cadoz 1988, Cadoz and Wanderley 2000) (section 2.3.2). A pianist depressing a key on a piano for example is using a selection gesture to choose which note to sound. The virtual instruments described here all used pitch to select which individual masses of the mass-spring physical models should receive force. Each of the physical models had twelve masses, one for each pitch class. The pitch class sounded by the musician determines which of these twelve masses will be acted upon by force. The amount of force, as we have discussed, is proportional to the volume of the sounded note.

¹When blowing a wind instrument, force is required to increase the amount of air passing through the embouchure in order to increase the volume of the sound produced.




Figure 4.2: Mapping from the volume of the source sound to force acting on the simulated physical model.

Pitch was well suited to use as a selection gesture, mainly because trained musicians have highly developed control over it. A professional trombonist with a four and a half octave range can accurately produce over 50 individual notes and can apply subtle microtonal variations to each note if desired. Their control over dynamic levels is also highly developed, but is more relative. Dynamic levels such as pianissimo, forte and fortissimo do not map to specific sound pressure levels for example. Rather, musicians judge how loudly to play based on musical context. For this reason, they are less able to break their dynamic range up into discrete chunks and so using volume as a selection gesture for individual masses would be likely to result in a virtual instrument which was less intuitively controllable.

The mapping of physical model parameters to visual display was simple. The intention was that the visuals would behave as the ‘body’ of the virtual instrument which the musician manipulated. Thus, the position of each mobile mass in the simulation was simply mapped to the position of a sphere in 3D space. It is certainly the case that more complex and potentially interesting visual mappings could be employed, but in order to contain complexity of development and maintain transparency of operation, this simple mapping was retained for each virtual instrument. However, this is certainly an area for future exploration.

The overall effect is that the performer and audience see a projection of a simulated physical model which moves and makes sounds in response to the sounds played by the performer. To the audience and musician, the physical model *is* the virtual instrument. In an editorial in the *Journal of New Music Research*, Paradiso and O’Modhrain argue that such an approach is an emerging technique for improving

transparency:

“One of the main challenges facing the designer of musical interfaces used in performance is to produce mappings that, at least periodically, strip away layers of abstraction, allowing the audience to smell the digital sweat as the artist pushes their instrument to the edge. One technique that relates to this issue involves projecting captivating and causal real-time graphical representations that tightly correlate with the audio component (not necessarily a light show, with visuals independently slaved to the sound, but instead generating both audio and video streams at a higher level, perhaps through the same algorithm or via a video-game metaphor). In this way the performer nudges and stokes a musical process that produces captivating visuals that, in some sense, replace the presence of the physical instrument. The graphics produced by the audio-visual environment ground the audience and the performer, effectively becoming the instrument’s “body.” (Paradiso and O’Modhrain 2003, p.347)

Given that we tend to perceive sounds as encoded physical actions (Leman 2007), the virtual instruments can be seen as recontextualising the physical forces embodied in the live instrumental sound, making them visible, and therefore also more readily perceivable.

4.5 Partial Reflections I

Having described the common characteristics of all three virtual instruments, I will now describe each of them in turn. I will begin by describing their technical characteristics and follow this with an ‘aesthetic perspective’ in which I describe the ‘character’ of the virtual instruments in more evocative terms.

4.5.1 Physical model

The first virtual instrument created as part of this project, *Partial Reflections I* was based on a very simple physical model, in fact an almost unchanged example model provided with the ‘Physical Modeling for Pure Data’ set of objects (Henry 2004a). Twelve masses are linked together to form a string, fixed at one end (figure 4.3). Each of the masses, with the exception of the mass at the bottom of the string is linked to two neighbour masses by a short spring.

The parameters of the simulation are summarised in table 4.1. These figures are provided for the reference of those who may wish to further explore the use of physical modelling and our particular models. However, details of the simulation techniques and equations will not be detailed here, as they are beyond the scope of this thesis. No claim for contributing to either physical modelling computation techniques or physics are made here. Rather, this is an investigation of the use of existing



Figure 4.3: The mass-spring model for *Partial Reflections I*. Circles represent masses and connecting lines represent springs. The smaller, shaded sphere's position is fixed, all others are mobile masses.

techniques in virtual musical instrument design. For details on the methods used by the physical modelling objects in the 'Physical Modelling for Pure Data' objects the reader is referred to Henry (2004a), Henry (2004b) and Henry (2006).

Simulation parameter	Value
Mass of mobile masses	200
Link length	50
Link rigidity	5
Link damping	1
Mass speed damping	1

Table 4.1: Mass-spring model simulation parameters for *Partial Reflections I*. Note that the values shown here are unit-less, as the choice of units has no impact on the simulation. Simulation parameters are proportional in that a force value of 10 applied to a mass of 100 will have the same effect as a force of 1 applied to a mass of 10, provided the simulation rate is unchanged (Henry 2006). Thus while it is possible to conceive the mass as being in kilograms and the link length in meters if desired, this has no intrinsic meaning in the simulation.

4.5.2 Mappings

The mappings for *Partial Reflections I* are summarised in table 4.2. The fundamental idea is that the musician can exert force on the physical model by playing into the microphone. Each individual mass of the physical model is associated with a particular pitch class (see figure 4.3). The pitch class of the notes played by the musician determine which mass has force applied to it. If the musician plays an A for example, then the mass second from the top has force applied; if F is sounded then the force

is instead directed at the mass third from the bottom and so on. In response, the masses will move around and, because they are linked to other masses in the string, this will cause all the other masses to move also.

Data from audio analysis	Mapped to	Continuous or discrete	Gestural function
Overall volume; volume of partials	Quantity and direction of force	Continuous	Excitation
Fundamental pitch class; partial pitch classes	Targets of force	Discrete	Selection
Fundamental frequency; partial frequencies	Frequency of oscillators	Continuous	Modification

Table 4.2: Summary of mappings from live audio to physical model in *Partial Reflections I*.

Force is applied to the mobile masses in three directions. The overall volume is mapped to force applied in the negative 'Z' direction. That is, away from the musician, in toward the screen. The volume of the strongest partial applies force in the negative X direction, pushing masses towards the left of the screen. The volume of the second strongest partial applies force in the positive X direction, towards the right of the screen. The intention is that when the musician plays a note, the associated mass will be pushed backwards away from the musician but also towards the left or right depending on the timbre of the note. By changing the timbre of a sustained note, the musician can alter the direction of force to the left or right.

Additive synthesis is used to generate sounds controlled by the movements of the physical model. In additive synthesis, complex sounds are produced by combining a number of simple waveforms - typically sine waves (Roads 1996). The pitch of the note played by the musician (ie. the frequency in Hertz) is mapped to the frequency of an oscillator associated with each mass. If the musician plays an A with a frequency of 440Hz (A above middle C) then the 'A' mass oscillator is set to oscillate at that frequency. If they subsequently play an A an octave lower, with a frequency of 220Hz, then the 'A' mass oscillator will then be set at 220Hz rather than 440Hz.

The frequencies of the three strongest partials in the live sound are mapped similarly. If the 'A' played by the musician has strong partials at frequencies with pitch classes of E, G and C#, then the oscillators associated with those pitch classes will store the frequencies of those partials.

Data from the physical model is used to control the output of the oscillators (table 4.3). The speed of each individual mass is mapped to the volume of its associated oscillator. The faster the mass moves, the louder the output from its oscillator.

The overall process is a variation on the analysis/resynthesis technique. Roads (1996) describes analysis/resynthesis as a three step process:

Data from physical model	Mapped to
Speed of masses	Volume of oscillators
X, Y, Z co-ordinates of mobile masses	X, Y, Z co-ordinates of on-screen spheres
Force exerted on mobile masses	Brightness of associated on-screen spheres

Table 4.3: Summary of mappings from physical model behaviour to audio synthesis in *Partial Reflections I*.

1. "A recorded sound is analysed
 2. The musician modifies the analysis data
 3. The modified data are used in resynthesising the altered sound."
- (Roads 1996, p. 144)

At a high level, the synthesis used in *Partial Reflections I* (and *Spheres of Influence*) follows this sequence although the modification stage (step 2) is a little more complex. In step 1, the live sound is analysed. In step 2, the audio analysis data is used to manipulate a simulated physical model which in turn produces the data used in step 3. In effect, the sound produced by the virtual instrument is a resynthesis of the original sound mediated by the physical model.

An additional simple visual effect was used to help make the operation of the instrument more transparent to the musician and the audience. The mass that was associated with the pitch class of the note currently being played by the musician was made to glow. The louder the sound, the brighter the glow. In addition, the masses associated with the pitch class of the three strongest partials in the live sound also glow, but with less intensity. The idea was that this would help reinforce the fact that the virtual instrument was responding to the changing harmonic structure of the live sound and changing the timbre of the sounds it produced. When the musician stops playing or changes note or timbre, the glow of the previously illuminated masses gradually fades.

4.5.3 Aesthetic Perspective

Partial Reflections I responds well to long notes with changing timbres. The 'physical' response of the virtual instrument emphasises changing timbres by making both them visible and more audible. The sounds it produces are clearly reminiscent of the acoustic sounds played to it, but remain distinct from them. The sounds are in a sense a reflection or echo of the live sound mediated by the (psuedo) physical structure at the core of the virtual instrument.

The movements of the virtual instrument give the impression of a strange, aquatic creature which is energised by sounds. The live sounds cause it to glow and writhe in a strange kind of slow motion dance, as it 'sings' back to the performer. Its song is



Figure 4.4: Ben Marks plays the virtual instrument for *Partial Reflections I*. Note the bright glow of some of the spheres indicating that they are currently having force exerted upon them. (Photo: Ros Hodgekiss)



Figure 4.5: Ben Marks plays the virtual instrument for *Partial Reflections I*. (Photo: Ros Hodgekiss)

always gentle and comprised of shifting, pulsing timbres; it does not play sequences of notes but generates a continually shifting, humming soundscape as it resonates to the sounds of its environment.

4.6 Partial Reflections II

Partial Reflections II was designed to have a very different character to *Partial Reflections I*. Overall, the software and music for *Partial Reflections I* is slow and smooth and emphasises changing timbres. In contrast, *Partial Reflections II* provides a strong rhythmic drive and is far more energetic. Unlike the first movement with its emphasis on longer notes with changing timbres, the software for second movement ignores all acoustic sounds other than the very beginnings of notes - the moment of articulation or 'attack'. Because the trombone has a very large range of potential styles of articulation, this still leaves the musician with considerable room for exploration.

4.6.1 Physical model

The underlying physical model for the software for this movement is again quite simple, and again is comprised of 12 masses - each one associated with a pitch-class. A link joins each individual mass to a fixed central point. Figure 4.6 shows the arrangement. The model is configured so that as soon as the software is started the 12 spheres begin oscillating very rapidly about the fixed point. In this initial state the length of the springs is so short that it is not really possible to see how many spheres there are (see figure 4.7).

The parameters for the physical model simulation were chosen so that the masses would be in perpetual motion, continually spinning (see table 4.4). The intention was to produce a model which had a very 'spiky' and unstable character that responded quickly and aggressively to the live sound.

Simulation parameter	Value
Mass of mobile masses	100
Link length	0
Link rigidity	100
Link damping	4
Mass speed damping	2

Table 4.4: Mass-spring model simulation parameters for *Partial Reflections II*.

4.6.2 Mappings

The mappings for *Partial Reflections II* are summarised in table 4.5. When the performer plays, the computer detects the onset of the note (ie. the moment of artic-




Figure 4.6: The mass-spring model for *Partial Reflections 2*. Circles represent masses and connecting lines represent springs. The smaller, shaded sphere's position is fixed, all others are mobile masses.



Figure 4.7: Photograph of performer and *Partial Reflections II* virtual instrument just before playing commences. (Photo: Ros Hodgekiss)

ulation) and determines its pitch class and volume. These are used to direct force onto the physical model, with the pitch-class selecting which sphere will have force applied to it and the volume determining how much force is applied. The louder the detected note, the greater the force.

Data from audio analysis	Mapped to	Continuous or discrete	Gestural function
Volume of note at onset	Quantity of force	Discrete	Excitation
Pitch class of note at onset	Targets of force	Discrete	Selection
Sample of attack (100ms)	Stored sample associated with each mass	Continuous (?)	Modification

Table 4.5: Summary of mappings from live audio to physical model in *Partial Reflections I*.

Force is always directed in the direction that will cause the masses to spin anti-clockwise. For example, if the mass is currently positioned at “3 o’clock” then the force will be directed in an upwards direction. If it is positioned at “2 o’clock” then force will be directed both an upwards and to the left. The effect is that by playing notes the musician can spin the masses anti-clockwise around the fixed point. The louder they play each note, the faster the associated sphere will spin.

As a result of the masses spinning more rapidly, they are pushed out from the central point. To illustrate, if the performer were to play a series of short, loud Cs, then the C mass would rapidly accelerate and move out into a higher orbit around the central point. When the performer stopped playing Cs, it would gradually spin back down to the central point. (In our model the friction is very low, so the masses never stop spinning completely. When force is stopped they instead orbit rapidly at a very low altitude around the central point.)

Rather than the additive synthesis technique used to produce sounds in *Partial Reflections I*, the *Partial Reflections II* virtual instrument plays back samples of the live acoustic sound. The physical model is used to trigger sample playback and manipulate the playback speed (and therefore also the pitch). When a note onset is detected, the software samples 100 milliseconds (ms) of the live audio from the microphone. That is, the moment of articulation is recorded and stored. The software has twelve buffers, each associated with a pitch-class, which can store 100ms of audio. The pitch of the note onset is estimated and this determines which buffer will store the sample. To extend our previous example, if the musician plays a C, then the first 100ms of that C will be sampled and stored in the buffer associated with that note. Any previous C onsets stored in that buffer are replaced by the new sample. This is the mechanism by which the performer alters the sounds produced by the virtual instrument.

Table 4.6 shows a summary of the mappings between the data generated by the

behaviour of the physical model and sound output. Playback of the sampled onsets is triggered by the movement of masses in the physical model. Each time a mass completes a half turn, the software plays back the recorded sound linked to that mass with one additional modification: the higher the orbit, the slower the playback. The effect of slowing the playback is to lower the pitch of the played-back note. So, if the mass has a very high orbit (because it has had a lot of force exerted on it) then the note that plays back every half rotation will be pitched quite low.

Data from physical model	Mapped to
Position of masses in orbit	Initiation of sample playback
X, Y, Z co-ordinates of mobile masses	X, Y, Z co-ordinates of on-screen spheres
Force exerted on mobile masses	Brightness of associated on-screen spheres

Table 4.6: Summary of mappings from physical model behaviour to audio synthesis in *Partial Reflections II*.

As with *Partial Reflections I*, masses which are having force applied glow in order to visually reinforce the workings of the virtual instrument.

A more complete example should help to summarise and illustrate the behaviour of *Partial Reflections II*. When the software starts the spheres are spinning rapidly around a central point at a very low altitude (see figure 4.7). If the performer plays a Bb several things happen:

1. The Bb mass has force exerted on it in proportion to the volume of the Bb.
2. The first 100ms of the attack are sampled and associated with the Bb mass.
3. In response to the force, the Bb mass is pushed out into a higher orbit.
4. Every half turn, the 100ms of recorded Bb is played back, but with playback speed (and therefore pitch) reduced by an amount proportional to the distance of the sphere from the central point.
5. When the performer stops playing Bbs the Bb mass will gradually spin back into the central point and as it does so the playback speed returns to the speed at which it was sampled.

4.6.3 Aesthetic Perspective

Partial Reflections II has a more energetic, driving character than *Partial Reflections I*. Perpetually in motion, the virtual instrument begins as a compressed ball of kinetic energy. In response to rapid-fire acoustic sounds it explodes about the screen like a mad atom, barely able to retain its shape as it is buffeted by sonic forces. Aurally, *Partial Reflections II* is relentless, spitting back short segments of recorded acoustic



Figure 4.8: Photograph of Ben Marks performing with *Partial Reflections II* virtual instrument. Note that due to the slow camera shutter speed, the photograph shows six masses on screen. Actually, to the naked eye, one mass was orbiting around the central point when this photo was taken. (Photo: Ros Hodgekiss)



Figure 4.9: Photograph of Ben Marks performing with *Partial Reflections II* virtual instrument. Here, two masses are orbiting the fixed central point. (Photo: Ros Hodgekiss)

sounds in a continually driving rhythm. When the performer is silent *Partial Reflections II* gradually returns to its tightly coiled ‘resting’ state, but does not become still or silent. Unlike *Partial Reflections I*, which resonates in sympathy with live sounds, *Partial Reflections II* is aggravated and energised by them.

4.7 Spheres of Influence

The virtual instruments for *Partial Reflections* were designed to be used in performances of music specifically composed for them. As such they had some characteristics which made them somewhat less controllable as musical instruments. For example, because the physical model for *Partial Reflections I* was a string fixed at one end, playing notes that were associated with the masses at the non-fixed end caused a lot more movement in the model than playing notes associated with masses at the fixed end. This suited the music that was written for that instrument, but in order to investigate the relationship between control and musical expression, the design of the instrument was modified to make it more like a traditional instrument, with a greater ability to exert control over individual notes.²

A simplified version of the virtual instrument from *Partial Reflections I* was created. The design of this virtual instrument, which we called *Spheres of Influence* (confusingly referred to as *Charmed Circle* in some papers), altered the interaction in several ways. Firstly, the musical features extracted from the live audio were reduced to two: pitch and volume. Secondly, to ensure that all musical pitches had equal influence on the model, the physical structure was redesigned to arrange the masses in a circle (figure 4.10). Finally, the links between the masses were removed so that forces exerted on one mass would have no impact on the others.

4.7.1 Physical model

The underlying physical model is slightly more complicated than the screenshot might indicate. The structure is in fact a cone, viewed from the base looking up to the point. Each of the masses which is visible on screen is attached via a spring, to the point of the cone, which is located in the centre of the screen. An additional set of springs links the visible masses to fixed masses which are positioned at points around the base of the cone (see figure 4.11). None of the fixed masses or the springs are shown on-screen. The parameters of the physical model simulation (table 4.7) are similar to those used in *Partial Reflections I* and were chosen so that the model was quite responsive to live audio, but was more stable than *Partial Reflections II*.

²In addition, a number of musicians who had seen and played with *Partial Reflections* had indicated that a simpler, more generic instrument might also be useful in teaching.




Figure 4.10: Screenshot of *Spheres of Influence* in its resting state.




Figure 4.11: Diagram of the physical model for the *Spheres of Influence* virtual instrument. Shaded circles are fixed masses, non-shaded circles are mobile masses and lines represent springs. The fixed masses around the outside are closest to the viewer, the fixed mass in the centre is in the far distance. The mobile masses (non-fixed) are positioned in between the outer fixed masses and the fixed central mass. None of the fixed masses are visible on-screen.

Simulation parameter	Value
Mass of mobile masses	200
Link length	400
Link rigidity	4
Link damping	0.5
Mass speed damping	0.5

Table 4.7: Mass-spring model simulation parameters for *Spheres of Influence*.

4.7.2 Mappings

When the musician plays a note, the software determines which pitch is being sounded. As with *Partial Reflections I* each mobile mass is linked to a particular pitch-class so that whenever a particular note is played, the sphere associated with that pitch-class has force exerted on it. The force is always exerted in an outward direction, so the sphere is effectively pushed outwards by the musician's note. The force is proportional to the volume of the note, so loud notes have a greater impact on the string of spheres than soft notes.

When the musician stops playing, the sphere that was being pushed outwards will spring back and oscillate for a time as it gradually returns to its resting position. In effect, the musician can 'pluck' the string on which the sphere sits by playing notes into the microphone. How hard the string is plucked is determined by the volume of the sounded note.

Data from audio analysis	Mapped to	Continuous or discrete	Gestural function
Overall volume	Quantity of force	Continuous	Excitation
Fundamental pitch class	Targets of force	Discrete	Selection
Fundamental frequency	Frequency of oscillators	Continuous	Modification
Amplitude of peaks in harmonic spectrum	Colour of on-screen spheres	Continuous	Modification

Table 4.8: Summary of mappings from live audio to physical model in *Spheres of Influence*.

The simple additive synthesis technique used in *Partial Reflections I* was also used in *Spheres of Influence*. That is, for each pitch class the frequency of the sounded note sets the frequency of an oscillator associated with that pitch class. The oscillator then produces a sine wave at that frequency at a volume that is proportional to the velocity of the associated mass.

Data from physical model	Mapped to
Speed of masses	Volume of oscillators
X, Y, Z co-ordinates of mobile masses	X, Y, Z co-ordinates of on-screen spheres

Table 4.9: Summary of mappings from physical model behaviour to audio synthesis in *Spheres of Influence*.

An additional visual element was incorporated into *Spheres of Influence*. The amplitude of the peaks in the harmonic spectrum were mapped to the the colour of the on-screen masses. Thus changing the timbre of the acoustic sound played by the musician would result in changing colours of the sphere associated with the pitch



Figure 4.12: Photograph of Ben Marks performing with the *Spheres of Influence* virtual instrument. (Note: This photograph was taken before the mapping between the harmonic structure of the acoustic sound and the colour of on-screen spheres was enabled.) (Photo: Ros Hodgekiss)

class of the note being played. This is something of a contrast with *Partial Reflections I* which would respond to changing timbres in the live sound by changing the timbre of the computer-generated sound and by moving in different directions, but did not incorporate colour. The timbre of sounds produced by *Spheres of Influence*, on the other hand, will not change in response to the timbre of the live sound, but the colour of the on-screen masses will.

4.7.3 Aesthetic Perspective

Spheres of Influence has a simpler, more direct character than *Partial Reflections I*, to which it is clearly related. Like *Partial Reflections I*, it responds well to shifting timbres rather than rapid streams of notes. Because each mass is independent (i.e. not linked to other masses) and partials in the live sound are ignored, the sounds it produces are simpler and have a more uniform character. While shifting timbres in the live sound do not change the sounds produced by *Spheres of Influence*, they do change its colour. A whistle will make a mass glow bright red, whereas a more complex acoustic sound will generate rich colours which react to and highlight the dynamic timbres of acoustic instruments. In contrast to *Partial Reflections II*, *Spheres of Influence* exhibits more placid, staid behaviour. Aggressive live sounds have an effect, but they quickly pass as *Spheres of Influence* bounces back to rest.




Figure 4.13: Screenshot from the *Spheres of Influence* virtual instrument, illustrating how different visual colours are produced in response to acoustic sounds with different harmonic structures.

4.8 Discussion

While each of the virtual instruments uses a simulated physical model to map between live sound and computer generated sound and video, each has a unique character. In this section key characteristics and similarities and differences between the virtual instruments will be outlined and discussed. In chapter 5 these will be considered in more detail in the context of the musicians' experiences with the virtual instruments.

4.8.1 Coupling

As discussed, the physical models at the core of the virtual instruments make use of fixed and mobile masses. An important difference between *Partial Reflections I* and the other virtual instruments is that the mobile masses in *Partial Reflections I* are linked to one another. In effect the physical model is one long string with masses placed along it (see figure 4.3). This means that when one mass moves the movement impacts upon all the other masses to a greater or lesser degree depending on the position of the mass on the string. For example, if the G mass (bottom of the string) has force exerted on it, causing it to move backwards, then because the G mass is linked to the F# mass which is in turn linked to F, E, and so on, the entire string of masses will be impacted.

This has the effect of simultaneously increasing the complexity and decreasing

controllability of the virtual instrument. While it is easy to understand the effect of playing individual notes on the physical model, playing several notes in succession can result in the string of masses moving in ways that are complex enough to be unpredictable. As we will discuss in chapter 5, this balance between complexity and controllability is a critical one for virtual instrument design. Coupling mobile masses in this way is one way to make the response richer without compromising transparency to a significant degree.

4.8.2 Continuous versus discrete excitation gestures

Both *Partial Reflections I* and *Spheres of Influence* use the volume of the live sound as a continuous excitation gesture. That is, there is no threshold volume below which the virtual instrument would not respond; the continuous stream of volume data is mapped to a continuous stream of force applied to the physical model. When the input volume was low, so is the force. Garth Paine argues that using input data in a continuous stream in this way, “more accurately reflects an ongoing interrelationship between the user and the system, and as such provides a more flexible, responsive and artistically rewarding outcome than a system based on the triggering of pre-made finite content (audio samples, for example).” (Paine 2002, p. 295).

Partial Reflections II on the other hand, used discrete excitation gestures, in that the audio analysis module was used to detect note onsets (attacks). Sounds below a certain threshold volume were therefore ignored by this virtual instrument. Once a note onset was detected, only 100ms of the live sound was mapped to force exerted on the physical model. The connection between live sound and the physical model is not continuous as it is with *Partial Reflections I* and *Spheres of Influence*, meaning that the behaviour of *Partial Reflections II* is less directly connected with the live sound of the musician. As we will discuss in detail in chapter 5, this tended to encourage the musicians to consider it an ornament or effect applied to their sound rather than an instrument or conversation partner with which they had an intimate connection.

4.8.3 Sound

Both *Partial Reflections I* and *Spheres of Influence* produced sound using additive synthesis whereas *Partial Reflections II* used sampling and playback of the musician’s live sound. Sampling was used in *Partial Reflections II* because it was felt that it would provide the musician with the ability to produce a wider range of sounds with the virtual instrument. The composition for this instrument made extensive use of breath sounds and percussive attacks as well as more conventional trombone techniques to arguably explore a wider soundscape than the music for *Partial Reflections I*.

In theory additive synthesis is capable of producing arbitrarily complex sounds.

The implementation of additive synthesis used in *Partial Reflections I* can produce complex and interesting sounds, but the diversity of timbres is limited by the fact that only twelve oscillators are used. The emphasis is on the *transition* from one timbre to another as opposed to individual sounds at any single point in time. Because the physical model is always moving in response to live sounds (at least while the microphone is on), the timbre of the synthesised sound is always changing with varying degrees of subtlety.

The simpler mappings used for *Spheres of Influence* mean that the timbres produced tend to be less complex. The additive synthesis oscillators store only the frequencies of the detected fundamental pitches of the acoustic sound. Partials in the live sound have no impact on either the movement of the physical model or the timbre of the synthesised sounds. In some ways, *Spheres of Influence* is like a simple keyboard instrument that allows the player to set the pitch of the notes produced by each key. This makes it highly controllable, but also means the sounds it makes are more like a simplified echo of the live sound. There are still 12 oscillators, as there are for *Partial Reflections I*, but because they store only the fundamental frequencies of the live sound the range of pitches they produce tends to be reduced.

The timbre of the sounds produced by *Partial Reflections II* does not change in the same way as those produced by *Partial Reflections I*. Because it simply samples acoustic sounds, it can produce a far wider range of sounds than the additive synthesis used in *Partial Reflections I* and *Spheres of Influence*. However, the connection between the changing state of the physical model and the timbre of sounds produced is limited. Really, the movement of the physical model changes only the pitch of the samples as they are played back³ and how frequently they are played. As *Partial Reflections II* simply mimics the musician's sound, they are able to change the timbre directly by altering the timbre of the attacks they play on their acoustic instrument, but they cannot 'get inside' the sound to alter its nature in more subtle ways.

4.8.4 Effect of timbre change

By changing the timbre of the sound they produce on their acoustic instrument the musicians are able to influence each of the virtual instruments in different ways. The physical model of *Partial Reflections I* responds to changing timbre in the input sound by changing the way it moves. As discussed in section 4.5.2, by changing the timbre of the input sound, the musician can push the mass associated with the current note to the left or right. Because the movement of the masses are linked to the additive synthesis engine, this will also have an audible effect on the timbre of the generated sound. Also, as the frequencies of partials set the frequency of the additive synthesis oscillators there is an additional audible effect.

³In *Partial Reflections II*, the distance of the mass from the central point changes the rate of sample playback. The samples are recorded at 44.1KHz. Playback speed ranges from 44.1KHz (ie. unchanged) down to 11KHz. Playback speed affects pitch, so notes played back at 11KHz will sound two octaves lower than recorded.

Changing the timbre of the live sound has an impact on the sounds produced by *Partial Reflections II*, because it simply samples and replays the live sound. Thus, the sound that is played into the virtual instrument is essentially echoed back, albeit with altered pitch. Changing timbre does not have any effect on the forces applied to the physical model, so there is no change in visual effect either.

Spheres of Influence responds to changing timbre by changing the colour of the on-screen spheres, but there is no effect on the timbre of the generated sounds. That is, there is a visual effect but no sonic effect.

In summary, it can be seen that changing timbre has the greatest effect on *Partial Reflections I* - changing the state of the physical model and thus the visuals and also the generated sound. While changing timbre has a direct effect on the sounds produced by *Partial Reflections II* there is no effect on the model and visuals. Finally, *Spheres of Influence* responds to changing timbre visually, but there is no impact on generated sounds or on the physical model.

4.9 Technical details

The virtual instruments were constructed using the open source visual programming environment *Pure Data* (Puckette 1997). The *Physical Modelling for Pure Data* set of objects by Cyrille Henry (Henry 2004a) provided the building blocks for the physical models and the *Graphical Environment for Multimedia* objects by Mark Danks and IOhannes M. Zmölzig were used to display the physical models in 3D.

4.10 Conclusion

In this chapter I have described three virtual instruments that were developed in collaboration with the composer and trombonist Ben Marks. The design criteria which guided their development were identified and explained, and details of the links between live sounds played by a musician on an acoustic instrument, forces exerted on a simulated physical model and the generated sounds and visuals were provided. Finally, characteristics of the virtual instruments were compared and contrasted.

In the next chapter I will present findings from a study which examined the experiences of a group of highly experienced musicians who used the virtual instruments.

Chapter 5

User Studies: Interactions with the Virtual Instruments

5.1 Introduction

In this chapter the user studies will be presented. Previously I have outlined a program of development and research and described a set of virtual musical instruments which have been designed to facilitate musical expression. Here I will address the following questions:

1. How do expert musicians interact with the virtual instruments?
2. How well do the virtual instruments meet the design criteria identified in chapter 4?

To address these questions data was gathered from a number of sources. These included:

- A user study in which seven professional musicians were videoed interacting with the virtual instruments, commenting on their experiences and responding to interview questions. This was the primary source of data.
- Notes made by observers who attended the musicians' sessions with the virtual instruments. These provided additional perspective on the musicians' experiences and helped identify whether the virtual instruments met the design criteria.
- A questionnaire administered during the musicians' sessions with the virtual instruments which attempted to directly elicit their opinion on whether the virtual instruments met the design criteria.
- Audio recordings of interviews with audience members attending a concert at which music for two of the virtual instruments was presented. These were important because one of the design criteria relates to *audiences'* experiences of the virtual instruments.
- A short interview with Ben Marks, the composer for whom the virtual instruments were created. This provided additional insight into musicians' experience with the virtual instruments from the perspective of the musician most familiar with them.

To address question 1 (How do expert musicians interact with the virtual instruments?), the video recordings of the musicians' sessions with the virtual instruments were transcribed and the grounded theory method was used to generate a theory of musician–virtual instrument interaction. The observers' notes provided additional perspective on this data and informed the development of the theory. The theory is presented in section 5.3 (Findings: Interactions with the Virtual Instruments).

The design criteria described in chapter 4 (section 4.3) reflect the goals I had for the virtual instrument software as I developed it during the artistic collaboration

with composer Ben Marks. In other words, they reflect the characteristics that Ben and I wanted the software to have. To address question 2 (How well do the virtual instruments meet the design criteria identified in chapter 4?), the gathered data—in particular the questionnaires and observers' notes—were examined for evidence that the virtual instruments did or did not meet these criteria. These findings are presented in section 5.4 (Findings: Evaluating the Virtual Instruments).

However, before jumping into the findings, the next section will briefly summarise the research methods and provide some examples of the specific techniques used to analyse the qualitative data gathered during the user studies.

5.2 Method Summary

The overall methodology of this research has been outlined in detail in chapter 3. The first stage of the project, designing the virtual instruments and identifying design criteria, has been described in chapter 4. The second stage, examining musicians' experience with the virtual instruments in order to refine the design criteria, involved a series of user studies in which professional musicians' interactions with the virtual instruments were closely examined. While the process of analysis has been described at a high level in chapter 3 (section 3.4.2), I will briefly summarise the process and provide a specific example of how the grounded theory method was used in this particular study here.

5.2.1 Participants

Seven musicians participated in the study. Each had a high degree of expertise and significant professional experience in their field. Because the virtual instruments were designed to work with monophonic instruments, the musicians were wind or brass players. Figure 5.1 details the demographic details that were collected from the participants.

	Musician	Summary of professional experience	Qualifications	Age / Range	Instrument
G		"Professional player for 23 years as an orchestral player, soloist and teacher."	ASCM	41	Trombone
M		"17 years professional orchestral performer. 20 years teaching experience. Arranger and composer Dip. Music (Performance) Sydney Conservatorium of Music using computers."	Dip. Music (Performance) Sydney Conservatorium of Music	37	Trombone
J1		"Sydney Symphony, Western Australia Symphony etc. Principal trumpet Australian Opera and Ballet Bmus (Sydney Conservatorium), Orchestra, teacher of music for twelve years."	Bmus (Sydney Conservatorium), GradDip (ANU)	31	Trumpet
J2		"Huge range of musical playing, composed and improvised. New music experimental music last 24 years."	Associate Diploma of Jazz Studies	44	Trombone
P1		"Orchestral musician for 30 years. 14 years with Opera Australia. His chamber music experience and improvised music experience spans last 25 years."	Bmus (Hons)	45-55	Clarinet
D		"* Classical freelance playing. * Contemporary 'Art Music' ensembles - Offspring, Halcyon, etc. * Teaching clarinet students."	LTCL, B.Mus BA (Hons-Philosophy)	35	Clarinet
P2		"15 years playing experience in jazz and improv/new music ensembles. Member of many groups including Australian Art Orchestra, Band of Five Names, DIG, Missy Higgins, etc."	BCA (Hons), MMus, GradDip I.M (Lib)	38	Trumpet

Figure 5.1: Demographic information collected from the musicians who participated in the study.

5.2.2 Data Collection

The user studies took place at the Creativity and Cognition Studios at the University of Technology Sydney, Broadway campus. The virtual instrument software was running on an Apple Macintosh computer and the visuals were projected onto a large rear-projection video screen. Audio output was via a set of Genelec 1029A studio monitor speakers, coupled with a 7050A sub-woofer. Sound input was via a Rode NT1-A condenser microphone connected through a small mixing desk to the audio input of the computer.

The steps involved in the user studies were:

1. The basic workings of the virtual instruments, including how they responded to the pitch and volume of the musicians' sound, were explained.¹
2. The musicians were given free reign to use the virtual instruments as they pleased. They were asked to verbally report and reflect on their experience as they did so.
3. The musicians were encouraged to prepare and perform a short piece of music using the virtual instrument.²
4. When the musician had finished with each virtual instrument, a semi-structured interview was conducted and a questionnaire administered.

The questions asked during the interview were:

1. Tell me about the piece of music you played.
 - (a) Why did it have these characteristics?
2. Do you have some comments about how easy or hard it was to write for or perform with the software?
3. Do you have some comments about the sound produced by the software?
4. Do you have some comments about the visual display?
5. While you were interacting with the work, did you become aware of any particular characteristics of your playing?
6. Did you play differently today than you normally would?
 - (a) In what ways?
 - (b) Can you say why?

¹Some musicians preferred to use the virtual instruments without prior instruction, in which case this step was skipped.

²Most musicians did not need to be asked to do this as their method of becoming familiar with the virtual instruments was to make music with them.

7. Do you have any suggestions or proposals for how we might improve or extend this software?
8. Can you think of any uses for this software?
9. Is there anything I should have asked you?
10. Do you have some comments or questions about what we've done today? Is there anything we should have done differently?

The questionnaire asked the musicians to rank their level of agreement with the following statements on a five-point Likert scale from strongly disagree to strongly agree.

- The instrument responded to the music in a way that I would describe as 'natural'.
- The instrument's behaviour was consistent. (ie. It responded to identical notes in the same way.)
- The instrument was conceptually simple.
- The instrument allowed me to create complex musical and visual effects.
- I found the instrument interesting and engaging.
- I felt in control of the instrument.

Responses to the questionnaire are analysed in section 5.4.2. The questionnaire data was not analysed using grounded theory methods.

The musicians interactions with the virtual instruments and the interviews were video recorded using a digital video camera. In addition, notes were taken during the session by an observer. For six of the seven sessions the observer was the co-supervisor of this research, Dr. Linda Candy and for the final session, interactive art curator and PhD student Lizzie Muller took this role.

In addition to the user studies, an interview with Ben Marks, the composer with whom the performance works using the virtual instruments were created, was conducted on January 9, 2007 via iChat (a text based internet chat tool). This interview complemented the user studies by providing insight into the experiences of the musician who was most familiar with the virtual instruments. As such, it provided valuable additional perspective on issues relating to musicians' interactions with the virtual instruments and was used as an additional source of data for the user study.³

³A full transcript of this interview can be seen in appendix D.

5.2.3 Qualitative Data Analysis

The user studies provided a very rich source of data. As discussed in detail in chapter 3 (section 3.4.3), techniques from the grounded theory method were used to facilitate qualitative analysis of the videos. The first step was to transcribe and annotate the videos. The open-source software Transana (Woods and Fassnacht 2007) was used to facilitate this process. Transana is a software package designed to aid qualitative analysis of video data. Transana enables the creation of transcripts which contain time codes. This maintains the link between the transcript and video and enables the researcher to quickly view the section of video associated with any section of the transcript and vice versa. It should be emphasised that Transana is only a tool which simplifies the management and examination of large amounts of video data. It does not automate analysis.

With Transana, clips of interesting video data can be created and labelled with codes (known as 'keywords' in Transana) which are specified by the researcher. For this research, Transana 'keywords' equate to 'codes' in the coding scheme. Once coding is complete, searches can be made which find all clips assigned particular codes. For example, a search could be made which found all video clips from all participants which were assigned the code 'control'. Each of these clips could be examined in detail to find key points of similarity and difference. These features were invaluable when dealing with the more than fourteen hours of video gathered during the studies.

During open coding, the video and transcript were examined in depth. The labelling of incidents in the data was achieved by creating video clips and labelling them with codes. If an incident arose which seemed to require a code that had not previously arisen, a new code was created. That is, codes were created on demand.

In grounded theory studies it is common to approach coding without any pre-defined codes (Miles and Huberman 1994, Glaser 1992). Beginning coding with a pre-existing coding scheme is also possible (Miles and Huberman 1994) if a suitable scheme exists or if the purpose of the research is to verify (rather than generate) theory. For this study, no pre-existing coding scheme was used.⁴

Of course, merely labelling incidents in the data does not create theory. In the grounded theory method, *memoing* is the process by which the analyst reflects upon and documents their evolving understanding of the situation under study. Memoing also helps the analyst to link the codes together into a theoretical framework. Memos are simply notes written by the researcher. They do not have a required format, the intention being simply that insights are captured quickly so that they are retained.⁵ Memoing in this study made use of a feature of Transana which allows the researcher to attach 'notes' to transcripts or collections of clips.⁶

⁴A fuller description of coding can be found in chapter 3, section 3.4.3.

⁵See Chapter 3, section 3.4.3 for a fuller description of memoing.

⁶See Appendix A for a list of memos from this study.

5.2.4 Analysis Example

The Grounded Theory method is a technique intended to help facilitate qualitative analysis. It is *not* an automatic procedure that takes transcripts as an input and produces theory as an output by following rigorous and repeatable processes. Two people using grounded theory methods on the same data may well produce quite different theories depending on their prior experience, interests and areas of expertise. Thus a problem which confronts researchers writing up findings from grounded theory studies is how to clearly demonstrate to readers the link between the data and the generated theory. In this thesis the following strategies are used:

- The specific analysis processes applied to this research are detailed in the following sections with reference to a transcript of one of the user study sessions.
- The codes which were created during open coding are included in appendix B.
- A sample report generated by the Transana software package which details the video clips and codes (known as 'keywords' in Transana) assigned to those clips is included in appendix C.

Transcription and Open Coding

In previous sections the process of analysis used for this research has been described. In order to clarify how the grounded theory techniques of transcribing, open coding, memoing and sorting (see chapter 3, section 3.4.2) were applied, a detailed example of the analysis process as it was applied to a short excerpt from one of the sessions will now be provided.

The first step in analysis of each of the user studies was to transcribe the video recorded verbal data. In fact, with one exception, transcription and analysis proceeded in parallel, as the principal researcher (myself) was also the transcriber. Obviously, in order to accurately transcribe the sessions, I had to attend to the words of the musician in detail. Transcription time was therefore a good opportunity to also use line-by-line open coding techniques. I did not transcribe the final user study session, involving musician P2,⁷ as it was felt that by this stage the key themes had emerged and that there was less to be gained from this. However, especially in the early stages, the process of transcription was invaluable in becoming completely familiar with the data.

The following excerpt from the session with musician J1 will be used to illustrate the process of open coding. J1 was the third musician to participate in the study, so a basic coding scheme had already been created during analysis of the previous interviews. This particular excerpt is useful here because it illustrates how previously created codes are applied to the transcript and how new codes are created when required.

⁷It was instead transcribed by my father!

A⁸ While you were interacting with the work did you become aware of any particular characteristics of your playing?

J1 Um, I think that I would like to make my chromatic scales a lot more even, so the lights were much... You know they were pretty honest. If you're not playing 100% rhythmically accurate it doesn't give you that effect.

A Yeah ok, I should be honest- I'm not sure how honest the software really is because it's got quirks. I mean sure there might be imperfections in your rhythm, but there's also imperfections in (the software). It's kind of like limits in how accurate it can be - it's no where near as accurate as the human ear. I mean it's much more computery - it's much more literal. So if you play a totally even chromatic scale but there's just a slight change in pitch of one of those notes, it can...[exaggerate the defect]

J1 I liked the idea of getting the thing to move away from me. So crescendoing to push it away and then as it came back adding in a lot of disjointed like ((sings)) din din din din. So it sort of as it flowed back it got lots of colours in.

A So pushing it back...

J1 And then playing with it. I liked doing that, I thought that was pretty... Cause then you not only have movement like the movement of it coming back and you don't have to worry about it moving any longer, you can worry about making it colourful.

A Yeah, kind of like setting up the base harmonies or something and then improvising on top.

J1 Yeah, yeah.

Transana facilitates the creation of video clips which contain sections of transcript (and associated video footage) that are relevant to the area under investigation. Thus, when I identified that this section of the session contained relevant material, the first step was to create a clip containing that section. Obviously, not everything that occurs during a two hour user study ends up in a clip as there are times when the chat is informal or off-topic, breaks are taken, etc. However, on average approximately 75% of the video material from each session ended up in at least one clip.

Having created a clip which contained the transcript and video, I examined the transcript carefully to try to identify 'what was going on' and mark up the clip with codes. The first code that was associated with this clip was quite utilitarian. Because the musician was responding to a pre-planned question ("While you were interacting with the work did you become aware of any particular characteristics of your playing?"), the code for this question was added to the clip. This aided later analysis by

⁸In the transcripts, 'A' is Andrew Johnston, the author.

making it possible to search for the clips from each of the musicians' sessions where they responded to this particular question.

Next, I considered the line, "Um, I think that I would like to make my chromatic scales a lot more even, so the lights were much... You know they were pretty honest." It seemed that the musician was talking about how the software had brought an aspect of his technique (evenness of chromatic scales) to his attention. Musicians had previously talked about how the software had allowed them to discover something about their playing that they had previously been less aware of, so a code, 'discovery', had previously been created for this. I therefore added the 'discovery' code to the clip.

In this sentence, the musician mentioned the 'lights', a visual aspect of the virtual instrument and how he would like to play more evenly so that they would do something, but he does not finish his thought. However, mentioning the visuals in this way still seemed to indicate a visually-driven approach, something which had arisen before. I therefore added the 'visual exploration' code to the clip.

Next, referring to the visual feedback of the 'lights', the musician says, "You know they were pretty honest". I interpreted this as a positive comment on the accuracy of the virtual instrument's response to his sound. Previous musicians had commented on the consistency of the virtual instruments' response and how it could be trusted to respond accurately. Therefore, I added the existing codes 'consistency' and 'trust' to the clip.

In the next fragment from the transcript ("If you're not playing 100% rhythmically accurate it doesn't give you that effect"), the musician seemed to be discussing how the software was particularly sensitive to rhythmic aspects of his playing. Once again, this theme had arisen before and a code had been created. In this case, the code was 'exaggerates/sensitivity', so this was added.

In the next fragment, I comment on the limitations on the accuracy of the software. Analytically this was not considered particularly significant, so this was not coded. However, in the next section J1 brings up an interesting point:

"I liked the idea of getting the thing to move away from me. So crescendo-ing to push it away and then as it came back adding in a lot of disjointed like ((sings)) din din din din. So it sort of as it flowed back it got lots of colours in."

Here the musician was making explicit something that previous musicians had not expressed, a kind of physical play in which he is deliberately 'pushing' the virtual instrument (*Partial Reflections I* in this case) back and then letting it fall back to its resting position as he 'coloured' it with his sound. Thus, a new code, 'physical play', was created and added to the clip. When creating codes in Transana it is possible to enter a description also, so I entered a description that seemed to fit the type of interaction J1 was describing:

Physical play Musician describes pushing, prodding or otherwise playing in a physical way with the virtual model.

I also created a 'memo' which captured my thoughts at that moment:

"From J1's session I got the idea of physical play. eg. When he describes pushing string spheres back and then filling them in with colours as they swing back. (1:22:00 approx.) This seems to me to tie in with the idea of counterpoint. Setting up the model - getting it into an interesting state - and then playing over the top of it. This leads to suggest an important characteristic of the virtual instruments:

CHARACTERISTIC: Instruments should have 'momentum', which allows the musician to get the instrument into an interesting/useful state and then play in counterpoint with it. Simple example would be setting up a drone which can be improvised over. Thinking about it, maybe this isn't a characteristic as such but rather a way of playing or compositional/improvisational technique which the virtual instrument can help or hinder."

J1 goes on to say, "I liked doing that, I thought that was pretty...", which was coded with the existing code 'like/dislike'. In addition, he comments that, "...you not only have movement like the movement of it coming back and you don't have to worry about it moving any longer, you can worry about making it colourful." In the interview I describe this as, "kind of like setting up the base harmonies or something and then improvising on top", and the musician agrees. As indicated in the memo above, I felt that this was a kind of 'compositional/improvisational technique'. As other musicians had discussed these kinds of techniques in previous sessions, the existing code 'musical strategies' was added to the clip.

So, it can be seen that the process of coding involves examining the data in detail and labelling 'incidents' in the data with codes. Similar incidents are labelled with the same codes. Once this is done, Transana allows the researcher to search for clips which have been associated with particular codes. For example, it would be possible to search for all clips which had had the code 'physical play' associated with them. This would retrieve the clip referenced above and others from each of the musicians' sessions. These kinds of searches helped enable comparisons between incidents in the data.

Building Theory

Obviously, merely labelling incidents in the data does not create theory, but building up a coding scheme in this way facilitates what Glaser and Strauss (1967) describe as the 'constant comparison' technique. Constant comparison simply involves comparing incidents in the data with one another, identifying similarities, differences



Figure 5.2: Whiteboard sketch showing the modes of interaction. Note that 'ornamental' mode was still characterised as an 'effect' at this stage of the analysis.

and relationships which are recorded in memos as the researcher identifies them. Through this process the researcher builds a theory which helps to make sense of the situation under examination. Memos help facilitate and, to some extent, document the researcher's evolving understanding of the links between these incidents.

While the coding scheme and memos are included as appendices to this thesis in order to help demonstrate that grounded theory analysis was indeed carried out, it is important to note that these are not a complete record of the analysis process. In this research coding and memo-writing were undertaken primarily to *facilitate* analysis rather than document it.⁹ A number of other, less easily documented procedures were also used to help build theory. For example, whiteboard sketches were used to help make sense of the different modes of interaction that were observed (figure 5.2). Also, old computer punch-cards, which had various aspects of the musicians' interactions and virtual instrument characteristics written on them, were arranged in several different ways on the floor in order to help clarify the various relationships that seemed to exist between them (figures 5.3 and 5.4).

These procedures are examples of techniques that were used in moving from coding the data to linking the codes together into a theory that can help explain or illuminate the situation under study. The grounded theory literature has a number of suggestions for how to go about this. The fundamental idea is that the researcher

⁹Thus it is not claimed that the coding scheme created during this analysis is a contribution of this research; it should instead be considered a by-product of the analysis process which generated the theory. The theory, of course, *is* a contribution.




Figure 5.3: An arrangement of computer punch-cards labelled with modes of interaction and virtual instrument characteristics. These were used to help clarify and make coherent the emerging theory.




Figure 5.4: Another arrangement of punch-cards, showing a number of desirable and undesirable virtual instrument characteristics as described by the musicians in the study.

examines the codes that have been created during open coding and attempts to identify higher-level concepts that make apparent patterns in the codes, and relationships between them. The approach described above draws primarily on the suggestions of Glaser (1978) and Miles and Huberman (1994).

5.3 Findings: Interactions with the Virtual Instruments

This section is primarily concerned with the first of the two questions this chapter addresses: How do expert musicians interact with the virtual instruments? The question of whether the virtual instruments met the design criteria will be considered in section 5.4 (Findings: Evaluating the Virtual Instruments).

In this section, the ways in which the musicians interacted with the virtual instruments will be described in depth. In addition, characteristics of the virtual instruments and their impact on the musicians' experiences will be considered.

5.3.1 Modes of Interaction

During grounded theory analysis it is expected that open coding will lead to the discovery of a 'core' category, a key issue which appears to have particular relevance to the situation under study (Glaser and Strauss 1967). The core category emerges during analysis as the researcher continually compares incidents in the data, noting relationships between incidents in memos.

During analysis of the data gathered during this study, it became apparent that the musicians did not always approach the virtual instruments in the same way. Sometimes a musician would express frustration because they felt they did not have enough control over the virtual instruments, but then at other times the same musician would complain that the virtual instruments were not autonomous enough, and that they wanted their behaviour to be *less* predictable. It seemed that the qualities the musicians sought in a virtual instrument would change during their interactions - that they interacted with the virtual instruments in different *modes*. Thus the core issue which emerged during analysis was that of *modes of interaction*.

A particularly striking example of a musician using different modes of interaction arose during the session with musician J2 in which he was comparing the experience of interacting with the three virtual instruments:

J2 This [*Partial Reflections I*] seems to be a more complex response to the previous one [*Partial Reflections II*]. And, particularly from the point of view of... The other one [*Partial Reflections II*] seemed not to create other harmonies. It was just like a delay. So that was an effect. So it tended to keep me diatonic. Just me, other players would do different things. When I hear this [*Partial Reflections I*], I don't know whether it's working off the overtone series or what and coming in with...you know...not the specific

root kind of idea, then it feeds me into other harmonic ideas. Cause I'm thinking of it as a player.

A Yep, that's interesting.

J2 Yeah, well I have to otherwise I don't sound very good.

A I'm just curious, did you think of the first one [*Partial Reflections I*] as a player as well or not so much.

J2 I thought of it more as an effect. As if I had...((stamps foot)).

A A delay pedal.

J2 Yeah yeah. It felt more like that. This [*Partial Reflections I*] feels like... I can imagine a person if you want.

A Ok, so you mentioned about the more complicated harmonic response. Was there anything else about it that made you think of it more as a person?

J2 It's not obviously a pulse. Um...The nature of a swell kind of thing. It's not linear. Maybe I'm not using the right terms. [A I know what you mean though.] But it swelled at times. [A Yes.] So it gives you a feeling of conversation. Whereas the other one felt specifically like a direct response to what I just played, where this feels more like a conversation.

It was apparent that J2 was describing two quite distinct approaches to the virtual instruments. He described his interaction with *Partial Reflections I* as a "conversation", because the sounds it was producing were complex and gave him "other harmonic ideas". In fact, he was deliberately conceiving of the virtual instrument "as a player", because if he didn't he didn't "sound very good". On the other hand, he approached *Partial Reflections II* quite differently, characterising it somewhat dismissively as "an effect", which worked a little bit like a simple delay pedal, which had a more "direct response" to what he played. At least at times then, J2 was not conceiving of the virtual instruments as 'instruments' in the traditional sense. That is, he was not particularly motivated to directly control them in the same way that he had learnt to control his acoustic instrument. Rather he was interested in having them respond to his playing in ways that would challenge him and lead his improvising in new directions.

The excerpt above illustrates the concepts of *conversational* and *ornamental* interaction. Ornamental interaction occurred when the musician thought of the virtual instrument as a kind of audio and/or visual effect which was added to their sound. In ornamental interactions, the musician was less concerned with controlling the virtual instrument and was less aware of it 'talking back' or providing new musical ideas. During the iChat interview Ben contrasted *Partial Reflections I* and *Partial Reflections II* in similar terms to J2:

A so maybe- you use the trombone more as a medium and [*Partial Reflections I*] as a kind of.. well I was going to say muse but that's probably not the right word..

A stimulant? partner?

B beast

...

A ok, this is an important point i think. You don't really conceive of [*Partial Reflections I*] as an instrument in the same way as the trombone then?

B no. It is more another musician I'm jamming with, although predictable and only really responding to me. The sound and look of this other musician though is enough to want to jam though

A i see. Do you conceive of [*Partial Reflections II*] in the same way?

B the composed piece is kind of like an agreement

B [*Partial Reflections II*] is a bit different I think

A it's more of an instrument?

B I feel more like it supports and empahsizes what I'm doing - it aids my communication with the audience

B I never thought of that before...shit.....

A can you think of any analogies from the trombone world? like a mute maybe? or some kind of electronic effect?

B a kind of light show. the sound was interactive though but more like a sampler that manipulates than it's own intelliegence

B hard one to think of. Not a mute but a periphary add-on device of sorts

A but not something you 'converse' with a la spheres?

B there was still a type of conversation but the sculpture was more fixed in its ways

Thus, Ben saw *Partial Reflections I* as a kind of "beast" or "another musician" that he was "jamming with". However, *Partial Reflections II* in contrast was something that supported and emphasised what he was doing and therefore aided his "communication with the audience". He describes *Partial Reflections II* as a "periphary [sic] add-on device of sorts", something that "manipulates" his sound as opposed to something with "it's own intelligence". In this sense, Ben also describes ornamental interaction - the virtual instrument adds something to his sound but does not talk back.

Contrast the styles of interaction described by J2 and Ben above with that described by musician G:

G What I like about it is the sensitivity; the ability to change the sound via my sound. What I'm finding challenging - I'm not saying it's wrong - but what I'm finding challenging is if I have an effect that I like, to be able to guarantee that same effect on demand.

A Yeah that's the consistency...

G Yeah that's my challenge. Ok I like that effect, but for example, I want to start the piece that way or I want to start a section that way, and that I know if I'm going to play X that I'm going to get X. That's the challenge for me. And I think that that's probably a combination of getting used to the software and B. maybe fine tuning the software for a particular work...

Clearly the interaction ideal that G is talking about here has a more instrumental character. He talks of "getting used to the software" so that he is able to "guarantee" that he is able to create particular pre-conceived effects during performance. Consistency is thus very important to G here in order that he can develop virtuosity - full control over all aspects of both acoustic and virtual instruments. Later, G goes on to describe how if the software supports this instrumental mode well then he can trust it to work without the visuals providing confirmation that the virtual instrument is responding as he expected:

G So it doesn't need me to really look at it. I can get an effect - or I think I can get an effect - by having pre-prepared myself and the software to get the effect that I want.

Initially, this was the kind of interaction that it was anticipated would dominate. While the design criteria that guided development had suggested that the virtual instruments should respond in ways that were complex and surprising at times, it was still expected that musicians would essentially want to maintain control. As it turned out, this was not necessarily the case.

It is important to note that it was not the case that each musician only approached the virtual instruments in one way, or that each instrument supported only one mode of interaction. It was common for the musicians to move between modes, even on one virtual instrument. While both J2 and Ben's interactions with *Partial Reflections II*, for example, seemed to be predominantly ornamental, musician D described using all three modes at different times with that virtual instrument. She spoke of how she first "responded" to *Partial Reflections II* (conversational interaction), then set it up into a particular state (instrumental interaction) and finally left it going while she played smooth phrases in counterpoint (ornamental interaction).

D [With] [*Partial Reflections II*] you sort of respond to it and set up something and then play along with it. So it is more of a duet.

...

D ... So even when you put it on a more subtle reception,¹⁰ I could still play a smooth melody over the top while it was still doing lots of activity underneath. Whereas with the other one it responded to everything I did- the speed and the... So it's almost the more simple response of this one allowed you to do something different in contrast with it, whereas the other one responded almost too much.

A Ok, more closely matched to what you were doing.

D Yeah, yeah. You couldn't get apart from it. Whereas this one you can get apart from it a bit and go, 'ok well I'm going to do something else that it can't follow me on'.

The "duet" that D speaks of here is different to a musical conversation. D's description of her approach indicated that she used the sounds produced by *Partial Reflections II* as a kind of background sonic wallpaper against which she would play in counterpoint. The fact that *Partial Reflections II* only responded to attacks, rather than notes in their entirety, enabled this approach, whereas the continuous control of *Partial Reflections I* and *Spheres of Influence* meant she "couldn't get apart from it" to play in this way.

The excerpts above illustrate the ways in which musicians described their use of the three modes of interaction: instrumental, ornamental and conversational. In summary, the musicians demonstrated and described interactions with the virtual instruments as:

- Instruments in the traditional sense
- Ornaments of their acoustic sound
- Conversation partners

As discussed, these modes of interaction are not exclusive in the sense that one musician always interacted with the virtual instruments in one mode, or that each virtual instrument was only used in one mode. Some instruments did tend to encourage particular interaction modes but not exclusively. These modes of interaction could best be seen as boundary points on a map of an individual's interactions with a particular virtual instrument (figure 5.5). As such, a musician may move from one mode to another, as musician D described above.

The three interaction styles will now be described in more detail.

Instrumental Interaction

When initially designing the virtual instruments, it was anticipated that the instrumental style of interaction would predominate. Musicians interacting with the virtual instruments in instrumental mode try to play them in a way that is analogous to the

¹⁰D is referring to microphone sensitivity, which had just been increased at her request.

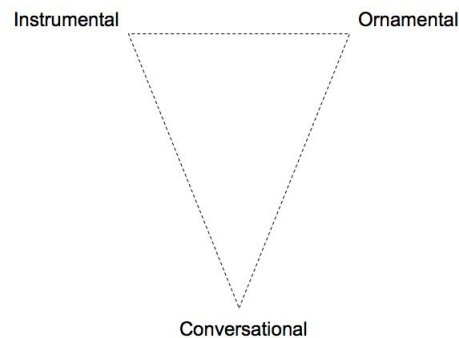


Figure 5.5: Three modes of interaction mark boundary points of a map of interactions with a virtual instrument.

traditional approach to acoustic instruments. They talk of controlling them and being able to guarantee that they can produce a particular musical effect that they like on demand. Musicians interacting in this way see the virtual instruments as *extensions* of their acoustic instruments. For these extensions to be effective, the link between the acoustic and virtual instruments needs to be clear and consistent.

Key issues that arise when musicians are in this mode are control, consistency, trust and proficiency. In this mode, musicians want to feel that the virtual instrument will do what they tell it - that they can trust it to respond consistently so that in performance they won't lose control. When the virtual instrument is consistent and controllable in this way, it allows the musician to build proficiency and facilitates the development of virtuosity.

J1 I don't think I'd be able to improvise to that. [A. Ok, that's interesting.] Because I think that I'd have to really understand like for instance if I played a G and it bent down to like an E or something like that. I'd have to know that was going to happen.

A OK, so you mean like it's more about really understanding exactly what it's going to do when you play something. That it's going to be very consistent as well. Because I guess the way it is now, pitch recognition on a computer is always a bit of an inexact science so it's pretty good but it's not perfect.

J1 I think maybe it could stem from the nature of my sort of performance as well. Because everything we do is so well prepared and so well rehearsed and practiced. Even just the way you practice the trumpet to be that sort of trumpet player is very methodical so for me to let go and not understand it and just... It's very difficult.

A Yeah ok. You say not understand it- this is the unpredictability thing

you're talking about?

J1 Yeah and I think the fear of not making it sound right. So, if I understood that if I played a G and it bent down to an E then I would be able to create an effect. So and then if it bent down to.. Say for instance I played G and it bent down to E, then I played E and it bent up to B then I could play another E or something like that and then i could start to create sort of like a... But without understanding that I find it too difficult maybe. Without a lot of...But then again the practice would just mean that I would understand it better.

It is interesting to note that when approaching the virtual instruments in this way that the musicians often tended to judge their own playing: they would try to control the virtual instrument by carefully playing a series of notes and then notice a technical issue in their playing that meant they did not quite get the effect they were intending. This indicates a desire to develop virtuosity on the virtual instrument and also points to a possible use of carefully designed virtual instruments of this type as tools to develop technical and aural skills in traditional music education.

P1 Yeah I quite like the way you can sort of build up a chord effect. I'd have to play more accurately. It wasn't quite what I wanted, but I didn't play accurately enough in tune at the beginning - that's the trouble.

Ornamental Interaction

When musicians use a virtual instrument as an 'ornament', they surrender detailed control of the generated sound and visuals to the computer, allowing it to create audio-visual layers that are added to their sound. A characteristic of ornamental mode is that the musicians do not pay active attention to the behaviour or sound of the virtual instrument. Rather, they expect that it will do something that complements or augments their sound without requiring much direct attention from them.

This mode of use was most prevalent with *Partial Reflections II*, which provided a fast rhythmical pulse and responded only to the beginnings of notes played by the musicians. Some of the musicians were happy to let go and use this effect as a colour or as a kind of background sonic wallpaper that they could play counterpoint too.

D I was sort of thinking of that one [*Partial Reflections II*], that provokes more of a duet sort of mode of thinking to me because you can set things up there and then play something quite different. Because it doesn't respond to long notes or sustained melodies without strong attacks, then you can actually set up two different things happening which is nice.

Other musicians found the pre-determined nature of the *Partial Reflections II* ornament overly restrictive. The fact that it tended to produce a constant, driving rhythmic pulse was seen as problematic by some because they felt was too dominant

and stifled their ability to take a more nuanced and fluid approach to rhythm. J2, for example, spoke of how he felt “alienated” by *Partial Reflections II*.

J2 If you want a feeling of domination and alienation, that’s certainly there with that one. I’m not being sarcastic. If you want the feeling that the machine actually is the dominant thing then that creates it quite strongly. All of that. It’s very strong, the feeling of alienation makes me uneasy.

It may be that the fact that the musicians approached *Partial Reflections II* as an ornament partly in response their perception of its behaviour as domineering. That is, they took an ornamental approach as a result of a kind of abdication of responsibility when the single-minded behaviour of the virtual instrument seemed to not really facilitate either instrumental or conversational interaction. That is, the virtual instrument seemed too far beyond their influence to facilitate instrumental interaction and its single-minded, domineering character made it unpleasant to deal with as a conversation partner. Faced with these problems some musicians simply left it to its own devices. J2, for example, hints that he may have taken this path:

J1 Like I find it amazing the whole thing, but your imagination is really dictated a lot by what you’re hearing back. You’re kind of affected by it and maybe you create that yourself because you’re playing it and it’s repeating it back to you but I don’t know I find it... I find that my imagination is limited because of that. You feel like you can’t expand more except for... But then again that’s a very early remark, so maybe I should just keep playing. Can I?

J1 starts by saying that he is “affected” by what *Partial Reflections II* does in response to his playing but that his imagination is “dictated” to by the sounds it makes. Later he speaks of his difficulties approaching it instrumentally:

J1 I don’t think I’d be able to improvise to that. [A. Ok, that’s interesting.] Because I think that I’d have to really understand like for instance if I played a G and it bent down to like an E or something like that. I’d have to know that was going to happen.

While some of the musicians liked the *Partial Reflections II* effect and saw it as a potential aid for establishing a connection with the audience (see the reference to the interview with Ben in section 5.3.1), the lack of control it allowed the musician was the most common observation. The other virtual instruments, *Partial Reflections I* and *Spheres of Influence*, were more flexible in that their design allowed them to be used as ornaments, but also as more controllable ‘instruments’ or more interactive ‘conversation partners’.

There are two aspects of *Partial Reflections II*’s design which seem most responsible for encouraging this ornamental approach. Firstly, the force which animates the *Partial Reflections II* simulated physical model is mapped to the first 100ms of each

note played by the musician. In other words, only the very beginning of every note, the attack, has any impact on the behaviour of the virtual instrument. Secondly, because these attacks must be distinguished from silence, there is a threshold volume. Acoustic notes below this threshold will therefore not impact on the virtual instrument. In contrast, the other virtual instruments both responded to every sound made by the musicians. The live sound acted like a continuous stream of force which acted on the physical model: even the softest sounds had noticeable effects. This continuous connection between live sound and force enabled fine-grained control. The discrete control afforded by the *Partial Reflections II* interaction scheme enabled the musical strategy of layering described by Musician D above, in which notes played above the threshold are used to put the virtual instrument into a particular state, against which notes below the threshold are played in counterpoint. However, this affordance came at the cost of reduced intimacy and controllability.

Conversational Interaction

A number of musicians talked of ‘conversing’ with the virtual instruments, that is, conducting a musical conversation with them as they might with another musician in a group. *Partial Reflections I*, with its more fluid style of movement and more complex sounds, was more likely to engage musicians in this way.

J2 [*Partial Reflections I*] gives you a feeling of conversation. Whereas the other one [*Partial Reflections II*] felt specifically like a direct response to what I just played, where this feels more like a conversation. (0:17:17)

A key characteristic of this mode of engagement is that musicians are not attempting to control the virtual instrument directly and are allowing it to more directly influence the direction of the music. This implies surrendering at least some of the control that characterises the instrumental mode. Musicians approaching the virtual instruments as conversation partners spoke of finding a balance between being able to control them and receiving rich sonic and visual responses in return. This ‘balance of power’ is a critical factor in facilitating conversational interactions. Simpler interaction styles, such as that used by *Spheres of Influence*, give a greater feeling of control but musicians can quickly lose interest. More complex interactions can be more satisfying but beyond a certain point the behaviour of the software starts to appear disconnected from the live music.

When musicians were engaged by a virtual instrument (generally when using *Partial Reflections I*), they spoke of the response as being ‘rich’ and ‘complex’, with the sounds and visuals having timbres, movements or colours that were multi-faceted. This allowed them to find new perspectives on the performance and led them to move in different directions. What differentiates the conversational mode of interaction is the sharing of control between the musician and the virtual instrument. The balance of power is in flux, allowing the virtual instrument to ‘talk back’ to the musician, re-

flecting and transforming the sonic input in ways that move the performance in new musical directions.

P2 discussed how as an improvising musician in particular, he was motivated to explore the boundaries of instrumental control. He spoke of a “spectrum” of the musician’s control over their instrument extending from absolute control over every detail to total lack of control.

P2 To me, it’s the whole spectrum of, it’s being able to have absolute control over everything that happens in your sound, from the moment it begins to the moment that it ends, and any brass player will tell you that that’s incredibly difficult to, so there’s that, um, having control over all those minute details of playing the instrument. That’s one end of the spectrum. At the other end of the spectrum I play things that I’m never quite sure... I alter my techniques so that I’m not quite sure how it going to come out. And that’s the spectrum. I think all people that are interested in improvisation, and interested in their instrument, have to have that spectrum, have to be able to have complete control over the instrument and be able to be interested in not having control over. Or at least being in situations where, could be through performing with other people, interacting with other performers. That can be an element of relinquishing control.

Thus it seems that for P2, being a good musician means having the skill to be able to completely control every aspect of his sound at the micro level. However, it also means seeking out situations where he is not in control. Continually putting himself in these situations seems to be a strategy to stimulate his creativity, a kind of tactic to ensure he is constantly being called upon to move beyond his comfort zone.

P2 I mean there’s all sorts of tactics that you can use. I mean, that’s essentially what jazz is, is never one person having complete control over the outcome of the music. It’s always, always, changing and always in flux, and that’s the thing that makes it such a difficult art form, but an interesting art form, and I think as a jazz musician I’m, you know, that’s my training is to be comfortable with those situations where I have no idea what’s coming up. I have no idea. And I am completely comfortable with that.

A It’s interesting. I wonder, it just occurred to me, is it something like it’s putting you in situations where you need to find control again. So pulling your trumpet apart, or doing something weird to it, so it’s now not the trumpet you knew, is putting you in the situation again a little bit like being a beginner, where you’ve got to go (P2 yeah) and find the way.

P2 Yeah, there are paradoxes in it. I mean, there are paradoxes in chance operations that, the more you want chance, the more rules you have to

impose upon things. That's part of the interesting thing about all Western arts, you know, the chance operations is basically another set of rules just to get you in a different spot. I think it's misinterpreted as being somehow, um, out of the realm of human control, that it is actually..human is the thing, I think.

P1 alluded to similar strategies, saying that he wished the virtual instrument would be more "adversarial" at times, "pushing back" in order to make the musical conversation more stimulating.

P1 It always has the note that you've done somewhere in the harmony and sometimes I think it would be nice if it did something contrary to what I just did. Rather than... It's quite complimentary isn't it? That's ok. I mean, some people are like that too, so... But yeah I like what it... It's certainly interesting.

A So you think there's too much of the note you play in the mix. You reckon there should be more...?

P1 Well I don't know whether it's almost a criticism, but it's an observation. That's the pattern that I notice. And the type of sound's it generates quite... Even when I did fairly harsh sort of sounds it was still quite smooth. It didn't suddenly go 'yearggghhh' at you you know.

A Yeah, you reckon that would be cool if it did do that?

P1 Yeah. I think it would be fun if it did something out of left field now and again.

A So when you pushed it it kind of... did something bizarre.

P1 Yeah, yeah. So if it gave the impression of throwing in an idea of its own. That would be nice.

A So as it was you didn't feel like there was much coming back from it?

P1 Yeah I did, but I felt that it was generally of a complimentary nature rather than, what's the word, adversarial. I didn't ever find it adversarial. And sometimes it'd be nice to have that [A Pushing back kind of..?] Yeah, like you did something and it deliberately did something opposite.

5.3.2 Focus of Attention

Musicians in the study directed their attention to different aspects of the virtual instruments at different times. The predominant focus for the musicians, perhaps unsurprisingly, was the sounds. Indeed, some musicians (J2, M, J) deliberately chose to ignore the visuals completely at times during their sessions.¹¹ However, musicians

¹¹In order to explore this issue further, during later sessions the musicians were asked to face away from the screen and comment on how this changed their experience.

also described taking a more visual approach, in which the audio was considered as something of a side-effect as they treated the virtual instrument as a *visual* instrument. In addition, musicians J1 and M demonstrated a slightly different focus at one stage during their sessions in which they approached the virtual instrument as a kind of mechanical entity, playing with it in a physical manner.

It was not the case that musicians always focused on one aspect of the virtual instrument. At different times they would 'tune in' to the visuals, the sounds or physical aspects of the virtual instruments. Thus the musicians did not always divide their attention rigidly between these different foci, and, in effect, often took a more holistic or integrated view.

These four different foci of attention, audio, visual, physical and holistic, are distinct from the three modes of interaction. The modes of interaction relate to the ways that the musicians manipulated and/or responded to the virtual instruments' behaviour. The foci of attention indicate which aspect of the virtual instruments - audio, visual, physical or all three (holistic) - the musicians attended to during their interaction. Thus it is possible that a musician can take an instrumental approach to the visuals of a virtual instrument or a conversational approach to the audio for example.

The foci of attention will be considered in more detail in the following sections.

Audio

The musicians in the study indicated that their predominant focus was on the overall sonic effect created by the acoustic and virtual instruments. Musician J2 went so far as to close his eyes at one point during his session with *Partial Reflections I*:

A It sounds like you're talking about it um... Most of the things you're saying are all aural things you know.

J2 Yeah...The visual?

A Yeah were you looking at it at all? [J Yeah] Was it just a distraction?

J2 No I don't find it a distraction. I quite like it.

A Did that help, to make it seem more person-like or creature-like?

J2 I'm having...I'm not certain. I was thinking about that. Cause I actually closed my eyes after a while to really concentrate on the sound. (0:17:49)

Musicians J1 and D both found the visuals for *Partial Reflections II* uncomfortable to look at for long periods. In particular, at times J found it difficult to attend to both the visual and aural aspects during performance with the virtual instruments:

J1 I think for me personally, I prefer not looking at the screen. Um, I think it will be awesome to have the screen there but not to actually be influenced by it.

A So you think - the audience can see the screen but you can't.

J1 Yeah. I think it- as a performer it takes your attention away too much. And I think it as a performance it would be awesome to have that there, but I just don't know if I could concentrate on trying to create a good performance whilst seeing... I don't know maybe it's because they say men can't multi-task. I don't know, but I'd find it very...

A Distracting? [J Yeah.] Do you think- it wouldn't distract the audience?

J1 No, no. I think that would be an extra part of it because they're not creating the music - they're having the experience. So, I think that would be a really good... (0:36:10)

J1's focus here is clearly on the audio. He seems to see the visuals as being separate from the 'performance' and sees them as potentially distracting him from the main task at hand. Thus, J1 here is describing an ornamental approach to the visuals. That is, he likes the fact that the audience can see the visuals as he feels this adds something to their experience, but he is not seeking to directly control the visuals or have them influence the musical direction of the performance.

M described a similar response to *Partial Reflections II*, describing how he felt that the visuals were 'incidental' and chose instead to focus on what sounds he could get out of it:

M I must say out of the 3 of them, this one I wasn't so much worried about the visuals. I was far more aurally involved than I was visually. The other ones were, 'ok how can I manipulate this with my sound' but this one wasn't so much that. It was like straight away, 'what could I do with the sound' I was concentrating on that, because for me the visuals were incidental to what I could do with the aural part of it.

P1 provides a clear example of a musician who has a general tendency to focus on the virtual instrument's sonic properties rather than the visuals. Here he describes how the visual's didn't engage him as completely as the audio. Once again, the visuals are described in ornamental terms, as a "fall out", "an outcome of the process rather than the other way around".

A Any comments about the visuals?

P1 Oh...yeah it's sort of attractive, but - it's hard to say... It didn't really do much for me like emotionally. So I tended to ignore it really. I was more interested in the sonic world. But I guess being a musician it's not unsurprising really... I mean I'm obviously less engaged, that's why I play music and I don't paint.

A Ok, if you were...

P1 ...Cause for me the sound's an emotional outlet. If I go to an art gallery it's much rarer that I would really be moved by a piece of painting. Usually I enjoy it and it's interesting to look at but...

...

A Did you find that at any stage were you trying to play in a particular way to influence the visuals?

P1 No I didn't no. No I was more trying to introduce...[A The sounds...] Yeah, the visuals was like an interesting fall out from what I was doing but I wasn't trying to actually do something. [A Kind of incidental really.] Yeah. It's a bit like you know how harmony is incidental in Xenakis' music. It's an outcome of the process rather than the other way around.

Visual

While the musicians had a general tendency to focus on sonic rather than visual or physical aspects of the virtual instruments, there were clearly times when they did choose to focus on the visuals. D gives a clear description of playing with a visual focus:

A Ok, tell me about the music you just played for [D For the round one?] the circle one.

D Um, I wasn't really attending to the music much, it was more like a technological approach to see what I could do to make the visuals change. So I can't remember what the music was like really. It was probably not musical at all. This [the clarinet] was a button, buttons [A A mouse or something like that.] sort of thing. Yeah, yeah. I was learning to use it.

A So is there a particular style of playing like legato playing or staccato playing or something like that, that is responds best to or...anything you particularly notice?

D Yeah I didn't notice that. I was sort of noticing what I could do with one note to make the different colours happen and how to play different notes in quick succession to make the thing bounce in certain kinds of ways. So I guess it made me aware of the, sort of the timbre thing, the timbre element, because of the visual response to that. Yeah I wasn't thinking in terms of duration much at all. So it was just more experimental. More visually experimental [A Than sonically?] Yeah. (2:08:50)

Several of the musicians (M, D and P1) spoke of their concern that having a visual focus had resulted in them producing music which lacked coherent musical structure and was thus potentially not as interesting to listen to. P1, for example, spoke of the 'disconnected' nature of the music he played while using *Spheres of Influence*:

A Tell me about the music you just played, any particular characteristics it had?

P1 I think in terms of music it's probably more disconnected than other things I've done. In terms of musical rationale, but it equated probably more closely with what was happening visually. (1:58:52)

D made the point that trying to attend to the visual and aural components at once was challenging and the visuals did tend to take her attention away from the music, but also that this was worthwhile because, "it makes you think differently".

A OK, do you have any comments about how easy or hard it was to play with it?

D Um, not really. But for me it was just a matter of getting my brain into that mode of playing. So it would get easier I suppose when I got to know it better. But it was easy because it sounded nice and it didn't matter what you did, it sounded pretty nice and it was fun. So it was easy in that respect. Hard to make something interesting but that's just because I'm not used to that mindset, being a performer, I'm used to having all the material that I do given to me. And I haven't done any improvising or stuff for a while so that just means my creative cogs are much slower. So that's maybe nothing to do with the program, it's more to do with me.

A Yeah ok. So getting into that mode of thinking - that's the mode you're talking about? The improvising mode? [D Yeah.] Or is it another kind of mode like a visual...mode or something?

D Yeah both. That's a good point, yeah. It's both. Because it's the improvising mode and a producing material kind of mode. But it's also the visual and the aural components at once. It's quite distracting - not in a bad way because it makes you think differently. But it does mean that you start to look at the gooby [the visuals] and respond to that, rather than listen to the sounds that are produced and respond to that. And there's not necessarily any necessary connection between the two. It's two different representations I suppose. So yeah, it's distracting but it's interesting so, that became a difficulty or a challenge maybe.

J1 described this challenge to traditional ways of thinking about composition/improvisation as a "totally different creativity". The visual response of the virtual instrument influenced his musical strategies, in this case encouraging him to do more "crescendoing and diminuendo" and "a lot of chromatic as well". Thus taking a visual focus has a significant impact on his music making. When he didn't look at the screen he returned to more traditional aesthetic.

J1 And I have to say that I get a totally different experience from not watching it.

A Ok, a better experience or just different?

J1 Just totally different. When I'm watching this I'm much more inclined to... in fact it does create a totally different creativity. So, like I'm doing a lot of crescendoing and diminuendo and I love the colours so I'm doing a lot of chromatic as well so that's very much affecting the way I play. But when I turn around I'm a lot more inclined to play something melodic and more traditional.

Physical

The physical behaviour of the virtual instrument is a defining characteristic of the software developed for this research. A particular kind of instrumental interaction was observed in some of the interactions with the virtual instruments in which the musician seemed to focus on these physical aspects. Musician J1 was most explicit in his description of this kind of interaction when playing with *Spheres of Influence*:

J1 And I think I might be just too much of symmetrical person. Because I find that I would love to... The effect that I originally went for... I automatically figured out where the D, the B the Ab and the F were. So, you can play with them. You know so you make the D bounce ((plays notes at 12 and 6 o'clock)) You know so they bounce off each other. And then ((plays notes at 3 and 6, 9 and 12 o'clock)).

A It's like you're playing pong!

J1 Yeah!

The physical focus that J1 is demonstrating here seems to be concentrated on the visual manifestation of the virtual instruments' physical characteristics - the way they bounce. This, of course, could be characterised as a visual focus. However, while there were not enough instances of this behaviour to saturate the 'physical focus' category, the behaviour appeared sufficiently different to the more typical examples of visual focus described in section 5.3.2 to warrant its own category. At one stage, musician M described a type of dance-like physical interaction when he tried to 'waltz' with *Partial Reflections I*:

M ((plays again- Blue Danube)) I was hoping I could make it swing back and forth.

A In what way swing back and forth? By playing the same phrase?

M I don't know. [A Like back this way? ((gestures))] Yeah... I was just thinking of doing a waltz with it.

While evidence for this physical focus was only seen in these two musicians, it nonetheless seemed distinct and interesting enough to warrant a separate category. One reason for this is that future virtual instruments are likely to use more

sophisticated physical modelling techniques which will provide greater opportunities for exploring this type of interaction. A tentative category here thus provides a starting point for further investigation.¹² At the least, the quotes above would seem to indicate that using physical models to mediate between live sounds and computer-generated sounds and visuals can facilitate musical approaches ranging from serious to whimsical.

5.3.3 Aspects of Musician - Virtual Instrument Interaction

In the previous section the musicians' interactions with the virtual instruments were characterised as having three modes: instrumental, conversational and ornamental. In addition, the musicians were shown to have demonstrated three foci of attention at different times during their sessions with the virtual instruments; namely audio, visual and physical. Here, aspects of the musician-virtual instrument interaction which were relevant to the musicians' experiences will be discussed. The focus will be on the relationship between the various characteristics of the virtual instruments, as identified by the musicians, and the modes of interaction.

Trust

Trust emerged as a critical issue in instrumental interaction. If musicians wished to interact with the virtual instruments as if they were extensions of their acoustic instruments, they needed to feel that the virtual instruments' response would be consistent so that they would be able to learn how to control them. They saw the development of skills as an important way to improve the quality of music they produced with the instrument. Skill development would not occur if the response of the virtual instruments varied to a significant degree given the same input. Musician G, for example, found that the level of unpredictability of *Spheres of Influence* was "challenging" at times.

G What I'm finding challenging - I'm not saying it's wrong - but what I'm finding challenging is if I have an effect that I like, to be able to guarantee that same effect on demand.

In this case, musician G was playing a rapid arpeggio into the microphone and the audio produced by the virtual instrument in response was subtly different each time. In fact, the pitch recognition would on occasion detect the root of the arpeggio an octave lower than what the musician was actually playing. He liked the response when this occurred but wanted to be able to produce it "on demand".

The fact that the virtual instrument provided a kind of visualisation of the sonic input from the musicians helped enhance the trust. Because the musician's control

¹²A more sophisticated virtual instrument which incorporated more realistic physical modelling has been built and used in live performances in late 2007 and 2008. It is expected that further research using this and other virtual instruments will be able to consider this type of physically-oriented interaction in more depth. See chapter 7, section 7.4 for a brief description of this virtual instrument.

over the virtual instrument was so dependent on pitch tracking and (for *Partial Reflections II*) onset detection, which are inherently inexact to some degree, the visuals were often used by the musicians to confirm that the software had correctly “heard” what they had just played.

J2 The visual reinforces the audio response. It lets you know you’re getting a response. So in a sense it could enhance the trust.

If trust was present to some degree, the visuals enabled the musicians to fine tune their technique so that the sounds they produced had the desired effect on the virtual instrument. Musician G found that while playing with *Partial Reflections II* he needed the visuals to help ensure that the articulations he was producing on the instrument were consistent enough to achieve the musical effects he was after.

G While I was looking for a musical and sound result, I felt that I needed the screen. The screen was very helpful to make sure that I articulated consistently.

This kind of interaction demonstrates that this type of software has potential in teaching, as a way of amplifying feedback from the acoustic instrument to help the musician become more aware of aspects of their own playing. Musician P1 talked about how the virtual instrument highlighted intonation problems when he was trying to get *Partial Reflections I* to produce particular chords:

P1 Yeah I quite like the way you can sort of build up a chord effect. I’d have to play more accurately. It wasn’t quite what I wanted, but I didn’t play accurately enough in tune at the beginning - that’s the trouble.

Often, trust was seen as a necessary first step towards a more conversational interaction. Musician P2, for example, was quite clear that he was not particularly interested in interacting with a virtual instrument in purely instrumental terms:

P2 I wouldn’t want to control it, I would want it to somehow control itself. That’s the thing that’s interesting about this to me. Otherwise it’s just a MIDI keyboard, I mean I don’t want to play a MIDI keyboard...

However, he was also adamant that, “the first stage in all these things is to be able to parallel a live performance and then once it can start paralleling it really accurately then, um, that obviously opens up scope for more interactivity.” That is, while he desired a degree of autonomy from the virtual instrument, so that it would lead him into interesting musical areas, trust was still an issue:

P2 Part of the thing of it is just like an intuitive thing of me knowing, like, being able to, I mean it’s a stupid word, but trust what the machine is, that it’s really accurate, and then I can kind of trust the randomness of it...

The trust in this sense is trusting that the machine is listening and ‘hearing’ properly what is going on, that it can ‘parallel’ the live performance. Beyond this though, the virtual instrument can (and should, in P2’s opinion) have a degree of autonomy or ‘randomness’ and produce musical material which is interesting and stimulating. Thus, the trust in this case operates on two levels. First, trust that the virtual instrument is listening and hearing what the musician is playing and second, trust that it will respond in ways that are musically interesting.

Richness of Response and Expressive Range

In addition to controllability, musicians wanted the sounds and visuals produced by the virtual instruments to be rich and complex. When musicians were critical of a virtual instrument, it was generally because they felt the sonic response was not complex enough to be musically interesting.

A So, can you tell me about the music you played- why it had the characteristics that it had?

M Ok. Well, first of all from first mucking around, with the sonic feedback from it, I thought, ‘ok well his is going to sort of suited to more the legato things’ - basically. I mean, while the doing staccato noises and stuff was all fun visually, but sonically it doesn’t really...wasn’t really cutting it. Just in terms of something interesting musically, that’s all.

M also described the visual display for *Spheres of Influence* as overly simple for use in performance:

M It’s interesting when you first learn it but after a while: ‘been there done that’. You can’t move it on from that point. For use in performance it needs to be more complex.

D described how she found that the sonic response from *Partial Reflections I* always had a similar quality even when quite different sounds were played into it. At one stage she asked for the microphone sensitivity to be increased so that the percussive sounds made by the keys of her clarinet would have more impact on the virtual instrument, but was disappointed that the sounds it made in response were not significantly different to those it made when she played more traditional clarinet sounds.

D When you turned the microphone up I could do more of these sort of things [*plays clarinet keys percussively*] but because I couldn’t hear it - all I could hear was what was coming back to me. And that wasn’t significantly different to what I was doing when I was playing louder.

Richness and complexity seemed to be a desirable feature of the virtual instruments, regardless of which mode of interaction the musicians took. For instrumental

interaction, musicians wanted a rich, complex sound which gave them the ability to explore a wide range of musical moods. The term 'complexity' in this case refers to the virtual instrument being able to produce a wide range of different sounds over which the performer exerts a high degree of detailed control. Improving the musical result when approaching the virtual instruments instrumentally involves becoming better attuned to the instrument, more aware of its capabilities and understanding the link between performer actions and audio-visual response of the virtual instrument. Musician G, for example, discussing the strategy he would use if writing a piece of music for *Spheres of Influence*, spoke of developing his understanding of the instrument so that he could more fully exploit its capabilities.

G My challenge would be to understand the capabilities of the software more. So I could perhaps get even more out of... that I would have more effects I would be able to choose from.

When taking a more ornamental approach, 'complexity' refers to compelling sounds and visuals which accentuate or exaggerate aspects of the musicians' sounds in order to support and enhance the audience's engagement with the work.

D I like how it sort of responded to my attack and exaggerated it in a kind of a way. I mean I might been making some duck noises but they sounded more duckish on that. So there was something a bit cartoonish or charicaturish going on there which I liked.

G found that the way that *Partial Reflections I* responded to lower notes was helpful for enhancing the dramatic elements in the live sound:

G I find also, from creative point of view, that this is quite dramatic. This is sort of software that I would like very much to use for very dramatic things. Just simply because it seems to pick those...the lower overtones and pump them out.

A complex response was also very important in more conversational interactions, where the term refers to the ability of the virtual instrument to inject musical ideas into the performance. A response which is musically complex can trigger deep listening and stimulate the musician to move into new musical directions. J2 specifically mentions the complexity of the response as being a key feature of *Partial Reflections I* that facilitated a more conversational style of interaction.

J2 When it responds in a more complex way... This seems to be a more complex response to the previous one. And, particularly from the point of view of... The other one seemed not to create other harmonies. It was just like a delay. So that was an effect. So it tended to keep me diatonic. Just me, other players would do different things. When I hear this, I don't know whether it's working off the overtone series or what and coming in

with...you know...not the specific root kind of idea, then it feeds me into other harmonic ideas. Cause I'm thinking of it as a player.

Drawing parallels with how people interact, and how a good conversationalist will introduce interesting and occasionally surprising ideas into a conversation, P1 expressed a desire for the response of the virtual instruments to do “something out of left field now and again”.

P1 It always has the note that you've done somewhere in the harmony and sometimes I think it would be nice if it did something contrary to what I just did. Rather than... It's quite complimentary isn't it. That's ok. I mean, some people are like that too, so... But yeah I like what it... It's certainly interesting.

A So you think there's too much of the note you play in the mix. You reckon there should be more...?

P Well I don't know whether it's almost a criticism, but it's an observation. That's the pattern that I notice. And the type of sound's it generates quite... Even when I did fairly harsh sort of sounds it was still quite smooth. It didn't suddenly go 'yeareggghhh' at you you know.

A Yeah, you reckon that would be cool if it did do that?

P Yeah. I think it would be fun if it did something out of left field now and again.

A So when you pushed it it kind of... did something bizarre.

P Yeah, yeah. So if it gave the impression of throwing in an idea of its own. That would be nice.

Interface Immutability

An aspect of interaction design which impacts on the richness of the audio-visual response is the degree to which the state of the interface changes during use. In order to facilitate instrumental interaction, the response of the software needs to retain a fairly high degree of consistency and this was one reason that the properties of the physical models that mediate between the live and computer generated sounds was fixed. That is, the weight of masses, tension of links and so on did not change during interactions.

This meant that the response of the virtual instruments would always tend to have a particular shape. Musician D spoke of wanting the virtual instruments' response to change over time so that the generated audio is more complex. Speaking of *Partial Reflections I*, she noticed that the nature of the physical model and the mapping between live audio and forces exerted upon it did, “...dictate that [the audio-visual response] will have a certain shape I suppose.” In other words, no matter what she played into the microphone, the “shape” of the response of the virtual instrument

would always be similar. This was due to both the physical model used for *Partial Reflections I* (a 'string' of masses suspended from the top of the screen) and the fact that the physical properties of the model itself did not change. It would move around the screen in response to the sounds played by the musician, but when they stopped playing it would always return to its resting place in a similar way as the simulated 'gravity' took effect.

Musician P1 spoke of all three of the virtual instruments having, "a particular set pattern in terms of length of sound and shape of sound" and suggested that it would be good:

P1 ...if you were able to change that within the piece and maybe have two or three other possibilities. I mean maybe with that one, when you stop playing the ball bounces back. What about if it just sort of stayed out there, or moved more slowly so you could actually have... You could change the speed at which the balls respond, you know. The speed at which they go back to normal again.

These comments suggest that changing the parameters of the physical model during musicians' interactions could lead to a more complex response which could facilitate more conversational interactions.¹³ If this is the case, the importance of instrumental interaction needs to be considered. The fundamental problem is that by increasing the complexity of the virtual instruments' response, we decrease the predictability of its behaviour and therefore reduce the scope for instrumental interaction.

The structure of the *Partial Reflections I* virtual instrument suggests that one strategy to combat this problem is to build simulated physical models in which the masses are 'cross-coupled' (Hunt et al. 2003). That is, models where individual masses cannot move in isolation. In *Partial Reflections I*, because the masses are linked into one big string, the movement of one mass will always impact, to a greater or lesser degree, on every other mass (see chapter 4, section 4.8.1).

Continuous vs. Discrete Control

As discussed in chapter 4 (section 4.8.2), *Partial Reflections I* and *Spheres of Influence* used mappings based on continuous excitation gestures, whereas *Partial Reflections II* used discrete excitation gestures. This discrete approach was something that musicians described as creating a greater feeling of distance or separateness from the virtual instrument. This was desirable at times, but also tended to reduce the scope for more instrumental or conversational interactions to develop. P2, for example, spoke of the need for the virtual instrument to be more responsive to the

¹³A fourth virtual instrument which enables the physical model to be modified during performance was recently developed by the author. During performances, link tension was adjusted and links between masses were cut at various pre-determined times. This appeared to give performers an opportunity to explore a wider range of sounds and visuals. See chapter 7, section 7.4 for a description of this software.

totality of what he was doing. He lost interest in what *Partial Reflections II* was doing because it was overly predictable and was not “interested” in anything beyond the attack of each note. This seemed to limit the extent to which intimacy could develop.

LM ... To some degree do you think it has some more of what you are after in the sense of having its own kinds of behaviour or has it actually got less?

P2 It definitely had its own behaviour, but the behaviour was predictable, and repetitious, so I lost interest in it. [L Mmmm] That's a crucial element in all of it. You want the behaviour to be interesting behaviour, not just [L Bad, like a naughty kid] [laughs] or the Sorcerer's Apprentice, where it's just uncontrollable and not that interesting. That's the difficult part I think, finding the right balance.

A So if it's possible... It's a hard question to ask, but if you could just imagine starting with something like this and you could do anything you like with it, could you imagine what would make this less repetitious and uninteresting, and more usable and interesting for you?

P2 Still mapping attack and pitch?

A No, this is a bit like brainstorming, so whatever occurs to you.

P2 (long pause) I felt that the way that it was reacting with this, was in a pretty limited way, in that it was only, like, so much of what I was doing, it wasn't interested in at all. (A yeah) So I felt that it was only going to be interested in ten percent of what it was that I doing, and to me it was the ten percent that wasn't that interesting anyway. So it was, um, the attack of the note is something that often programmers go for, like they want to get that attack of the note, but as a performer sometimes that's not, you know, the most interesting bit. Or, it's just one of many bits that, do you know what I mean?

J2's reaction to *Partial Reflections II* was similar, and expressed in no uncertain terms.

J2 ... Its response is... It's like a machine. It feels like a machine.

A Yeah, whereas the other ones felt like creatures.

J2 Absolutely.

L So it's driving you?

J2 It's ignoring me. It's responding to me but in a very limited fashion. It doesn't under... I'm talking to it and it doesn't understand. It's just coming back with the same old shit. Yeah. So when I said the word trust. It's interesting. Alright coming back to that other thing- my prejudice before [against computer music]. That displays it.

A It's like what you imagined.

J2 It's a machine with very limited understanding and very limited range of response. The other one - two, the other two - but particularly the middle one. I don't know what it is about it, seems to actually be more conversational. So instead of being this predictable. When I say predictable - I couldn't change the effect of the first one [*Partial Reflections I*]. I change intensity, I change range, I change... It didn't seem to recognise just sounds without pitch. Alright, it didn't appear to. No matter what I did, its response was basically the same. I know it wasn't but in so many ways it was. And the other ones may have been but they didn't appear to be. (0:54:57)

J2, frustrated with *Partial Reflections II*'s "very limited understanding" of what he plays to it, describes it as always "coming back with the same old shit". In contrast, both *Partial Reflections I* and *Spheres of Influence*, which use continuous control, give the impression that they are interested in the totality of his sounds and facilitate a more conversational, or complex instrumental, interaction.

While discrete control tended to reduce the feeling of connection with the virtual instrument for J2 and P2, D indicated that it did have advantages when interacting in a more ornamental mode. D, for example, found that the discrete approach used in *Partial Reflections II* allowed her to set up a kind of counterpoint. She spoke of how the continuous nature of the link between her playing and the movements and resulting sounds of *Partial Reflections I* did not enable her to "get apart from it." On the other hand, because only those notes that had a volume above the threshold had an effect on *Partial Reflections II*, she was able to set it up into a certain state and then play softly in counterpoint as it continued to make sounds.

D So it's almost the more simple response of this one [*Partial Reflections II*] allowed you to do something different in contrast with it, whereas the other one [*Partial Reflections I*] responded almost too much.

The issue of continuous vs. discrete control was only directly addressed by three musicians (J2, P2 and D), but would nonetheless appear to be a design feature which can impact significantly on the modes of interaction adopted by musicians. The evidence gathered in this study suggests that continuous control, used by *Partial Reflections I* and *Spheres of Influence*, tends to give musicians the impression that the software is 'listening' and 'interested' and therefore can lead to more successful conversational and instrumental interactions. However it seems that the feeling of separation between live sounds and computer generated sounds and visuals, which discrete control seems to lead to, can nevertheless be a useful feature which leads to a more ornamental mode of interaction. Given that only one virtual instrument which uses discrete control was developed for this research, this is an area that warrants further investigation.




Figure 5.6: Virtual instruments which support conversational interaction facilitate a shifting balance of power between musician and virtual instrument.

Shifting Balance of Power

The balance of power between musician and virtual instrument was a critical issue. Each of the three modes of interaction can be placed on a balance-of-power continuum (figure 5.6). At one extreme is instrumental mode, where the musician aims for absolute control of every aspect of the virtual instrument's behaviour. At the other is ornamental mode, in which the musician has no control over the virtual instrument, which adds audio visual 'layers' to the live sound in ways the musician does not seek to influence. Conversational interaction is an interesting hybrid, in which power is shared between musician and virtual instrument as they each take the lead in setting musical direction at different times during the performance.

Partial Reflections I, which was the virtual instrument that the majority of musicians preferred¹⁴, effectively supported a wider range of interaction styles than either *Partial Reflections II* or *Spheres of Influence*. Musician P2 explicitly discussed this:

P2 Whereas the other two, particularly the first one [*Partial Reflections I*], I didn't have any sense of [being locked into it]. I felt that I could actually ignore it if I wanted to, or interact with it. That level of engagement that I talked about earlier was more interesting to me in the first one [*Partial Reflections I*].

Later, he went on to say how this shifting balance of power was at the core of his approach to improvisation:

A So, you're not thinking of this, it seems to me, as an instrument (P No) at all?

P2 Well, an instrument that, well what would be interesting would be to be able to change the level of engagement that you have with the thing, in terms of controlling it, or, to have some, to have different levels of interactivity, and I wouldn't necessarily want to control what level of interactivity I was on at that point. I would rather be able to interact with it spontaneously. Do you know what I mean?

¹⁴Section 5.4.2 summarises questionnaire results that support this claim.

A Yeah. I think I know what you mean.

L But it also sounds like that's not just in terms of the relationship to this instrument. It's also in terms of, that you're kind of leaning towards chance and giving away a bit of control, starts even with your relationship with your trumpet.

P2 That's true. To me, it's the whole spectrum of, it's being able to have absolute control over everything that happens in your sound, from the moment it begins to the moment that it ends, and any brass player will tell you that that's incredibly difficult to [do]. So there's that, um, having control over all those minute details of playing the instrument. That's one end of the spectrum. At the other end of the spectrum I play things that I'm never quite sure... I alter my techniques so that I'm not quite sure how it's going to come out. And that's the spectrum. I think all people that are interested in improvisation, and interested in their instrument, have to have that spectrum, have to be able to have complete control over the instrument and be able to be interested in not having control over [it]. Or at least being in situations where - could be through performing with other people, interacting with other performers. That can be an element of relinquishing control.

This would suggest that in order to support conversational interaction, virtual instruments need also to support instrumental and ornamental interaction modes to a degree, and allow musicians to transition between them. If a musical 'conversation' is taking a direction the musician is not comfortable with, then they can move towards a more instrumental approach to shift the performance into a different area. Likewise, if certain musical ideas have been explored and the musician is looking for fresh input, they might give the virtual instrument more freedom and allow it to inject fresh musical material into the mix.

Musician J2 made a very interesting comment towards the end of his session which suggests that this kind of conversational interaction was fundamental to his approach to music. Discussing musical interactions within ensembles he said:

J2 You can have that alienating feeling with human beings you know. Vast, huge amounts of bands have that. They're individuals playing you know. They're all quite good but they're still individual. When you're move, when you're actually moved, when your heart is moved is when you know there's absolute... They're playing for each other and they're playing each other. You know, there's real communication.

The phrase, "they're playing for each other and they're playing each other" is quite poetic, suggesting that musicians at times approach *one another* instrumentally. This is something which is perhaps not often acknowledged. It is clear that musicians play *for* one another and in so doing contribute new musical material to the performance.

If the musical purpose of this playing is to add another layer or sonic element to the overall performance then this approach could be seen as 'ornamental' in the terms used in this thesis. However, J2's comment above suggest that at times the purpose of the musical contribution could be to change the playing of other musicians in specific ways, which suggests an 'instrumental' interaction. In J2's view, a successful, moving performance can only be generated when the members of an ensemble are able to engage with one another in both modes: ornamental ("playing for each other") and instrumental ("playing each other"). Skilled, aware musicians are able to shift easily, rapidly and frequently between these modes and in so doing enable a more 'conversational' interaction which takes the music to another level.

Degree of Intimacy

Both instrumental and conversational interaction require a level of intimacy to develop between musician and virtual instrument. Often, in instrumental interactions, the goal of the musician is to embody the virtual instrument so that they no longer perceive it as a separate entity. The detailed physical aspects of the interaction when this ideal state is achieved are managed by the subconscious. This approach can be seen as a kind of strategy to facilitate virtuosity and is a technique advocated by many instrumental pedagogues, perhaps most notably Arnold Jacobs (Stewart 1987, Frederiksen 1996).

Likewise, an ideal conversational interaction involves a high degree of intimacy. P1 spoke of the subtle and deep connection that can develop in musical interactions with other players:

L You used the word partner before. Do you imagine it was like a partner in some kind of musical expression?

P1 Yes I do. And that's what I want from it and if you're playing with a person there are many many subtle things that he can just do with a note which you know it's like having a conversation with him. And that's really satisfying when that actually happens. Because the thing about the musical relationships that you can develop is that they're much more...They're deeper and more personal than you would normally have with that person in real life. You know there's your partner, but you wouldn't normally have conversations where you're really relating to the person on that level in normal conversation.

When taking an ornamental approach however, musician and virtual instrument are more distinct. The virtual instrument acts as a kind of peripheral device which adds layers and visual effects to the live sound.

'Naturalness' of Interaction

The musicians in the study agreed that the response of the virtual instruments to their sounds in ways that seemed 'natural'.¹⁵ The response of *Partial Reflections I* and *Spheres of Influence* were felt to be more natural than *Partial Reflections II* which had a more mechanical character.

The musicians struggled to articulate exactly what it was that made the response of these instruments appear more natural but the slower movement of the model and the continuous (as opposed to discrete) link between live sound and force exerted on the model seemed to be important in producing a response that the musicians found more organically linked with the sound they made on their acoustic instruments. Musician J1, for example, spoke of how the response of the *Spheres of Influence* model seemed more appropriate than *Partial Reflections II*:

J1 And it feels like, like when you go like this, it actually feels like a note when you watch it. So you go: ((plays a note)). To me, in my head, that's ((indicating screen)) a note. Whereas the other one was more of a...[A More abstract somehow?] Ah, just like a line across the page and that, the round ball, just seems, rather than actually turning into a line, it actually seems to stay as a note.

J2 spoke in very clear terms about the importance of the interaction having a relationship to the 'natural' world, contrasting the response of *Partial Reflections I* (a "natural construct") with that of *Partial Reflections II*, which he found "so mechanical".

A You're talking about the effect of the weight of the sound?

J2 Yeah, my dynamics. It seemed to be very direct and related to the natural world. Yeah! Very much so. That one more than anything visually related to the natural world to me. It was a natural construct. You can see it. In fact I could make it physically and do it. You know, I'd have to find balls with the right density and size, but I could actually make it do that. [ie. physically build something similar.] You know trombone's interesting because it actually does put out sound and air and all of that kind of stuff and I could make the balls move in that way. I could play with them in that way. Whereas the other one is so mechanical.

P2 described *Partial Reflections I* similarly and seemed to enjoy the idea of his sounds being represented as "things slinking around".

P2 Also part of my attraction to this, is that it's an interface that I can... I guess I'm more familiar with some movement that's going on. There's something sort of archetypal about what's going on here, you know, it

¹⁵See sections 5.4.1 and 5.4.2 for evidence from the observer's notes and questionnaire findings which support this claim.

booms away when I blow loud, and it's, the responses are... I don't want to anthropomorphise it, but it seems more adhering to the rules of... I don't know...

A It seems a more natural response?

P2 Something sort of dancierly about it, the way it moves. There's something balletic. ((Plays)) I think aesthetically this one makes more beautiful shapes. It just moves elegantly for me. I guess that's nice to see that sort of representation of what I'm playing. Things slinking around...

A The visual style seems appropriate to..

P2 Not sure if it's appropriate, it just, that's attractive to me for..

L When you said you hate to anthropomorphise - why do you hate it? And where are you going with that sentence?

P2 I didn't want to, I thought that maybe it was moving like a humanistic thing, (L More organically?) more organically, yeah.

L But you don't like to say so because..?

P Um, did I say that?

L You said, "I don't want to anthropomorphise it", you said.

P2 Ok. To me that would limit... I wouldn't even want to be thinking that, because it would just, you know, load my interpretation of it. But it definitely has a more, kind of organic feel to it.

While these quotes illustrate the difficulty musicians had in articulating exactly why they felt the responses of these instruments was more natural, they seem to suggest that the design of the physical model, especially choices of mass values and spring tension, are very important in achieving a 'natural' effect. The comments of P2 above, that the movements of *Partial Reflections I* (which was most often characterised as having a natural response) were 'humanistic', 'dancierly', 'balletic', 'beautiful' and 'elegant' seem to quite strongly relate to human movement, despite his reluctance to anthropomorphise the virtual instrument.

In the late 1960s, Edward Ihnatowicz developed a remarkable robot creature named the *Senster* (Zivanovic 2005). Developed in partnership with the Philips corporation, the robot was displayed at the Evoluon, Philips permanent technological exhibition, where it captivated audiences with its lifelike movements and responses to audience behaviour. The *Senster* had an array of microphones on its head which allowed it to sense the direction of sounds, and it was programmed to shift and turn its head towards the source of any sounds it detected. This meant that by talking or clapping, members of the audience could attract the attention of the robot which would turn its head and move its body towards them. Ihnatowicz spent a great deal of time refining the control system and mechanics of the *Senster* in order that its movements would be as lifelike as possible. Of particular interest is the fact that

Ihnatowicz came to the conclusion that the style of its physical movements had far greater impact on how people interacted with it than its appearance. Despite the fact that it was 2.4 high and 5m long and made of welded steel tubes, its movements nonetheless convinced audiences that the *Senster* had intelligence and personality.

“He came to the conclusion that the shape and the general appearance of the structure were of very little significance compared to its behaviour, and especially to its ability to respond to the public. People seemed very willing to imbue it with some form of animal-like intelligence and the general atmosphere around it was very much like that in a zoo.” (Zivanovic 2005, p. 104)

Flexibility

Flexibility was a concept that arose often in the investigation. Musicians spoke of wanting the virtual instruments to be ‘flexible’ in that they should be able to:

1. Respond to a wide variety of input sounds;
2. Produce a wide range of complex sounds and visuals;
3. Be adjustable, in order to support a wide range of different performance situations;
4. Allow musicians to switch between modes of interaction.

Items 1 and 2 have been discussed previously (see section 5.3.3) and relate closely with the concept of complexity. Item 3, adjustability, was not widely discussed by musicians but this pragmatic concern did arise during the session with musician G and in the interview with Ben. G looked for confirmation that the software could be adjusted to suit a particular composition:

G I mean if you were saying as a general rule you were thinking of a piece and you would say, ‘Oh, I want a longer reverb’ or ‘I want a shorter reverb’, it’s possible to adjust that? (0:11:48)

Ben also mentioned this point, pointing out that a good virtual instrument should be “adaptable to the environment” (see appendix D).

The final type of flexibility is perhaps the most intriguing but also the least discussed by the musicians. It arose during the discussion with musician P2, during which he refers to a kind of flexibility in the way that musicians can interact with the virtual instrument. He talks of being able to choose to ignore the virtual instrument or actively engage with it.

P2 But I felt that I was locked into it [*Partial Reflections II*]. I felt that I was going to have to... All of my playing was going to get this thing to

do something that I would like. Whereas the other two, particularly the first one [*Partial Reflections I*], I didn't have any sense of that. I felt that I could actually ignore it if I wanted to, or interact with it. That level of engagement that I talked about earlier was more interesting to me in the first one.

As discussed in section 5.3.3, P2's comments suggest that virtual instruments which support conversational interaction should be flexible enough to effectively support all three modes of interaction. That is, the musicians should be able to ignore the virtual instruments (ornamental interaction), control them (instrumental interaction) or converse with them (conversational interaction).

5.4 Findings: Evaluating the Virtual Instruments

In the introduction (section 5.1), I said that this chapter was concerned with two questions:

1. How do expert musicians interact with the virtual instruments?
2. How well do the virtual instruments meet the design criteria identified in chapter 4?

The previous section (Findings: Interactions with the Virtual Instruments), has been primarily concerned with addressing question 1. This section will consider question 2.

5.4.1 Observer's notes

As discussed in chapter 3 and section 5.2.2, in addition to the musician and principal researcher, an observer attended all of the evaluation sessions. For six of the seven sessions the observer was the co-supervisor of this research, Dr. Linda Candy and for the final session, interactive art curator and PhD student Lizzie Muller took this role. During the session, the observer took notes and occasionally participated in discussions. These notes were used to support the analytic process by providing an additional perspective on events and discussion that occurred during the musician's sessions with the virtual instruments.

A short set of instructions were provided to the observers, detailing the kinds of issues this research was concerned with and what questions they should keep in mind during the sessions. (This document is included in appendix E.) In particular, observers were asked to consider how well or badly the virtual instruments met the design criteria. The observers took notes quite freely¹⁶ and at the conclusion of sessions with each virtual instrument provided a brief summary of whether, based

¹⁶A sample page of notes summarising the observer's findings in regard to one virtual instrument can be seen in figure 5.7.

Time	Event	Comment
6:15	A. intro	
6:16	Plays (12) 4 mins Comments on the fun of activating different roles using another.	could play a bit longer Visual-Attention is strong Conversation A admits he does not know why Motivation to play visually not necessarily (is this good?) confirmed by comments
6:28	A. suggests play without visual Plays (13) 4 mins D. liked this but still finds playing visuals more captivating.	therefore... "captivating"
6:31	So now STRING played without visual Plays (14) 3 mins	Refers back to (1) The attention to both visual & aural elements affects response & dual attention is difficult to achieve with party
Questions	Q1 Hugo: Birds not paying attention, leaning to use it. Style? - not noticed, instead concentrated on impacting the visuals.	
Q2	(Much easier to control, whereas string harder to control) predictable (this directly related to visual feedback as a result of the feedback.)	
Q3	'Complexity of string none arbitrarily related to sounds' Sound - not as good as other - BANNER? - pitch limitations, no decay...	
Q4	Visuals - Fun - excitement curiosity pt of view	
Q5	Playing? Response surprises to a single note	
Q6	Difference to normal playing. Attention to timbre & to visual aspects - never do that normally. DS is interested in phis music - ? : pushing ideas re performance & music - 'somatic' element.	

Figure 5.7: Sample page of notes taken by an observer present at the sessions.

on their observations of the musicians' interactions, they felt there was evidence that the virtual instrument had met, or failed to meet, seven of the design criteria.¹⁷ Figures 5.8, 5.9 and 5.10 show a summary of their observations.

¹⁷Criteria 1 ('The virtual instruments should be able to interact with a wide variety of acoustic instruments') was not considered here as each musician only played one instrument.

Partial Reflections I

Criteria	Musician							Total (sessions showing evidence of this criteria being met)
	G	M	J1	J2	P1	D	P2	
2. The virtual instruments should respond in a way that seems natural.	1	1	1	1	1	1	1	7
3. Virtual instrument response should be consistent.	1	1	1	1	1	1	1?	6
4. Virtual instruments should be simple but allow skilled musicians to create complex effects.	1	0	1	1?	0	1	1	4
5. The virtual instruments should be interesting, engaging and motivate the musician.	1	1	1	1	1	1	1	7
6. The musician should feel in control of the virtual instrument, but it should retain the ability to surprise.	1	1	1	0	1	1	1	6
7. The virtual instrument should encourage a playful, exploratory approach, especially in new users.	1	1	1	1	1	1	1	7
8. The relationship between live sound, the behaviour of the virtual instrument and the resulting sounds and visuals should be apparent to observers (eg. audience members).	1?	1	1	1	1	1	0	5
Total (number of criteria met during session)	6	6	7	5	6	6	5	
Average number of criteria met per session								6.00

Figure 5.8: Summary of findings by the observers of the musicians' interactions with *Partial Reflections I*. '1' indicates the observer found clear evidence the criteria was met; '0' indicates that the observer did not find evidence the criteria was met; '1?' indicates that evidence was found but was incomplete or unconvincing (these are not counted in the totals).

Criteria	Musician								Total (sessions showing evidence of this criteria being met)
	G	M	J1	J2	P1	D	P2	Total	
2. The virtual instruments should respond in a way that seems natural.	1	1	1	0	1	1	0	5	5
3. Virtual instrument response should be consistent.	1	1	1	1	1	1	0	6	6
4. Virtual instruments should be simple but allow skilled musicians to create complex effects.	1	0	1?	0	1	1	0	3	3
5. The virtual instruments should be interesting, engaging and motivate the musician.	1	1	1	0	1	1	0	5	5
6. The musician should feel in control of the virtual instrument, but it should retain the ability to surprise.	1	1	0	0	1	1	0	4	4
7. The virtual instrument should encourage a playful, exploratory approach, especially in new users.	1	0	0	0	1	1	0	3	3
8. The relationship between live sound, the behaviour of the virtual instrument and the resulting sounds and visuals should be apparent to observers (eg. audience members).	1	1	1	1	1	1	0	6	6
Total (number of criteria met during session)	7	5	4	2	7	7	0		
Average number of criteria met per session									4.57

Figure 5.9: Summary of findings by the observers of the musicians' interactions with *Partial Reflections II*. '1' indicates the observer found clear evidence the criteria was met; '0' indicates that the observer did not find evidence the criteria was met; '1?' indicates that evidence was found but was incomplete or unconvincing (these are not counted in the totals).

Criteria	Musician								Total (sessions showing evidence of this criteria being met)
	G	M	J1	J2	P1	D	P2	D	
2. The virtual instruments should respond in a way that seems natural.	1	1	1	1	1	1	0	6	6
3. Virtual instrument response should be consistent.	1	1	1	1	1	1	1	7	7
4. Virtual instruments should be simple but allow skilled musicians to create complex effects.	0	0	1	1	1	0	0	3	3
5. The virtual instruments should be interesting, engaging and motivate the musician.	1	1	1	1	1	1	0	6	6
6. The musician should feel in control of the virtual instrument, but it should retain the ability to surprise.	1	1	1	1	1	1	1	7	7
7. The virtual instrument should encourage a playful, exploratory approach, especially in new users.	1	0	1	1?	1	1	0	4	4
8. The relationship between live sound, the behaviour of the virtual instrument and the resulting sounds and visuals should be apparent to observers (eg. audience members).	1?	1	1	1?	1	1	0	4	4
Total (number of criteria met during session)	5	5	7	6	6	6	2		
Average number of criteria met per session									5.29

Figure 5.10: Summary of findings by the observers of the musicians' interactions with *Spheres of Influence*. '1' indicates the observer found clear evidence the criteria was met; '0' indicates that the observer did not find evidence the criteria was met; '1?' indicates that evidence was found but was incomplete or unconvincing (these are not counted in the totals).

The grids of numbers shown in figures 5.8, 5.9 and 5.10 can be viewed in rows or columns. Examining the numbers in columns makes it possible to examine the observers' findings in relation to each individual evaluation session. Thus, considering the column for the session with musician G in the grid for *Partial Reflections I* (figure 5.8), it can be seen that the observer found evidence that all criteria were met with the exception of criterion 8 ('The relationship between live sound, the behaviour of the virtual instrument and the resulting sounds and visuals should be apparent to observers'). Continuing across the grid, it is possible to see that the sessions with musicians J1 and J2 provided evidence that all seven criteria were met. (That is, the total number of criteria met during their session was 7.) In contrast, the sessions with P1 and P2 provided evidence that only 5 of the 7 criteria were met, according to the observers. However, given that on average 6 of the 7 criteria were observed as being met each session, it seems clear that the sessions provided substantial evidence that *Partial Reflections I* met the criteria overall, from the observers' perspective.

Considering the grid as a series of rows makes it possible to examine the observers' findings in terms of each individual criterion. Looking at the first row of the *Partial Reflections I* grid (figure 5.8), for example, it can be seen that criterion 2 ('The virtual instrument should respond in a way that seems natural') was observed as being met in all seven sessions. In contrast, evidence that criterion 4 ('Virtual instruments should be simple but allow skilled musicians to create complex effects') was met was found only in 4 of the 7 sessions.

Analysing the grids in this way helped to summarise the conclusions drawn by the observers. They are:

- Overall, the virtual instruments met the design criteria. Of the three virtual instruments, *Partial Reflections I* best met the design criteria, meeting an average of 6 criteria per session. *Spheres of Influence* averaged 5.29 criteria per session and *Partial Reflections II* was the least successful in these terms, meeting an average of 4.57 criteria per session.
- The musician's experiences with, and responses to, *Partial Reflections II* in terms of the design criteria were more variable than they were with the other virtual instruments, the number of criteria met during each session ranging from 0 to 7. In contrast, *Partial Reflections I* ranged from 5 to 7 and *Spheres of Influence* from 2 to 7.
- Criterion 4 ('Virtual instruments should be simple but allow skilled musicians to create complex effects') was least often met. *Partial Reflections I* was observed to meet this criterion in 4 out of 7 sessions (and with some incomplete evidence in a fifth session). *Partial Reflections II* was observed to meet this criterion in 3 out of 7 sessions (with some incomplete evidence in a fourth) and *Spheres of Influence* was observed to meet it in 3 out of 7 sessions.
- All three virtual instruments met criterion 2 ('The virtual instruments should

respond in a way that seems natural'), with *Partial Reflections I* most successful in this regard.

- Criterion 3 ('Virtual instrument response should be consistent') was met overall. P2 was the only musician whose session did not provide clear evidence of this (for *Partial Reflections I* and *Partial Reflections II*. This may be because P2 was not particularly concerned with this criteria (see section 5.3.3 for example).
- Criterion 5 ('The virtual instrument should be interesting, engaging and motivate the musician') was met by *Partial Reflections I* in all sessions; by *Spheres of Influence* in 6 of 7 sessions; and by *Partial Reflections II* in 5 of 7 sessions.
- Criterion 6 ('The musician should feel in control of the virtual instrument, but it should retain the ability to surprise') was generally considered to have been met by *Partial Reflections I* (in 6 of 7 sessions) and *Spheres of Influence* (all sessions). *Partial Reflections II*, in contrast, was less successful with the observer only finding evidence this criteria was met in 4 of 7 sessions.
- Criterion 7 ('The virtual instrument should encourage a playful, exploratory approach, especially in new users.') was met by *Partial Reflections I* (all sessions). *Partial Reflections II* and *Spheres of Influence* were less successful in this regard, with evidence they met this criterion being found in 3 of 7 and 4 of 7 sessions respectively.
- Evaluating the degree to which the virtual instruments met criterion 8 ('The relationship between live sound, the behaviour of the virtual instrument and the resulting sounds and visuals should be apparent to observers') was complicated by some ambiguity in the instructions given to the observers. When reviewing the notes at the conclusion of the study, it became clear that there was some confusion over whether the observers were looking for evidence that the *musician* thought this criterion was met or whether they were being asked to give their *own* opinion on whether it was met. So, with some qualification, it can be seen that there is at least some evidence this criterion was met. There was evidence that *Partial Reflections I* met this criterion in 5 of 7 sessions (with incomplete evidence in a 6th session); *Partial Reflections II* met it in 6 of 7 sessions; and *Spheres of Influence* met it 4 of 7 sessions with incomplete evidence in 2 more sessions.¹⁸

5.4.2 Questionnaire Findings

The questionnaire administered at the conclusion of each of the user studies is a good starting point to get a broad overview of how the musicians rated the various

¹⁸The degree to which this particular criterion was met is considered in more detail in section 5.4.3, in which a number of audience members are interviewed about how well they could identify the relationship between live sound and computer generated sounds and visuals.

virtual instruments in terms of the design criteria. The questionnaire asked the musicians to rank their level of concurrence with several statements on a five-point Likert scale from strongly disagree to strongly agree. The questionnaire was intended to elicit directly the musicians' opinion on whether each of the virtual instruments met the design criteria. The statements were:

1. The instrument responded to the music in a way that I would describe as 'natural'.
2. The instrument's behaviour was consistent. (ie. It responded to identical notes in the same way.)
3. The instrument was conceptually simple.
4. The instrument allowed me to create complex musical and visual effects.
5. I found the instrument interesting and engaging.
6. I felt in control of the instrument.

It can be seen that these questions correspond almost exactly with the list of design criteria with a few exceptions. Firstly, design criterion 1 (that the virtual instrument was suitable for use with a large variety of acoustic instruments) was not something the musicians individually were expected to answer. Secondly, criterion 4 (virtual instruments should be simple but allow skilled musicians to create complex effects), was addressed with two separate questions, 3 and 4 on the questionnaire. This was because it was felt that musicians may want to provide separate feedback on these two points. Thirdly, criterion 7 (the virtual instruments should encourage a playful, exploratory approach, especially in new users), was not addressed in the questionnaire as it was felt that this was an aspect more likely to be meaningfully explored by examination of qualitative data. Finally, criterion 8 (the relationship between live sound, the behaviour of the virtual instrument and the resulting sounds and visuals should be apparent to observers) was something that could be better explored in collaboration with audience members rather than performers.

In general, the musicians did not have difficulty completing the questionnaires, but as with any quantitative survey of this nature the musicians at times had difficulty rating their experiences with the virtual instruments in terms of the questions and categories provided. Some quotes from the evaluation sessions indicate this frustration:

J1 Can you agree and disagree with something? It says 'I felt in control of the instrument' and there was definitely aspects of it I felt totally in control of and there was others that I didn't. So I've got two answers to that question.

A It's always the problem with doing this kind of research. I guess that's probably neutral. With all of this questionnaire- the important thing is what

we talk about [it]. For the questionnaire I guess put neutral. Because we're recording all this I'll know what you mean.

Musician J2 was also quite frustrated by the questionnaire at times:

J2 Now this is... See this is why I find this question confusing. Cause there's no sense of satisfaction or dissatisfaction. So I find that even though it's a simple question it actually doesn't satisfy me at all. It doesn't give me a chance to say what I wanna say. It was pleasantly inconsistent or its balance of consistence and inconsistency. You know what I'm saying- it's too limiting.

A Yeah I know what you're saying. I mean, that's why we've got the video camera here... This [questionnaire] is very much a blunt instrument. [J2 It's a very blunt instrument.] So if you feel or disagree that it was consistent but you liked that...

J2 I don't... It was the perfect balance between consistency and inconsistency.

A Yeah I see what your saying, so to say 'neutral' [on the questionnaire] is probably the wrong thing isn't it.

J2 Totally wrong because - 'neutral' - how bloody positive is that? It's not positive at all. Anyway, moving on, I'll get this other one... 'The instrument was conceptually simple'. No, it was complex- it was beautiful.

The musicians' responses to the questionnaire provide us with some evidence of the degree to which they felt the software met the design criteria. Due to the small sample size and the simplistic design of the questionnaire this in itself is not a comprehensive study. If the sole aim of this research was to validate the fact that the design criteria were met, a far more rigorous process of questionnaire development and piloting would be required. However, the intention at this stage is to get a high-level indication of which interfaces the musicians felt did meet the design criteria and to what degree. For this purpose the questionnaire is adequate.

	Partial Reflections I	Partial Reflections II	Spheres of Influence
Musician G			
The instrument responded to the music in a way that I would describe as 'natural'.	3	4	4
The instrument's behaviour was consistent. (ie. It responded to identical notes in the same way.)	4	5	2
The instrument was conceptually simple.	3	4	5
The instrument allowed me to create complex musical and visual effects.	5	5	5
I found the instrument interesting and engaging.	4	5	4
I felt in control of the instrument.	3	4	3
Musician M			
The instrument responded to the music in a way that I would describe as 'natural'.	5	5	5
The instrument's behaviour was consistent. (ie. It responded to identical notes in the same way.)	5	5	5
The instrument was conceptually simple.	5	5	5
The instrument allowed me to create complex musical and visual effects.	4	4	3
I found the instrument interesting and engaging.	4	5	4
I felt in control of the instrument.	5	5	5
Musician J1			
The instrument responded to the music in a way that I would describe as 'natural'.	5	4	5
The instrument's behaviour was consistent. (ie. It responded to identical notes in the same way.)	4	4	3
The instrument was conceptually simple.	3	2	3
The instrument allowed me to create complex musical and visual effects.	5	3	5
I found the instrument interesting and engaging.	5	5	5
I felt in control of the instrument.	5	3	3
Musician J2			
The instrument responded to the music in a way that I would describe as 'natural'.	5	1	4
The instrument's behaviour was consistent. (ie. It responded to identical notes in the same way.)	4	4	4
The instrument was conceptually simple.	2	5	3
The instrument allowed me to create complex musical and visual effects.	5	2	5
I found the instrument interesting and engaging.	5	1	5
I felt in control of the instrument.	4	1	5

	Partial Reflections I	Partial Reflections II	Spheres of Influence
Musician P1			
The instrument responded to the music in a way that I would describe as 'natural'.	4	4	4
The instrument's behaviour was consistent. (ie. It responded to identical notes in the same way.)	4	4	4
The instrument was conceptually simple.	4	4	4
The instrument allowed me to create complex musical and visual effects.	5	5	5
I found the instrument interesting and engaging.	4	5	5
I felt in control of the instrument.	3	5	4
Musician D			
The instrument responded to the music in a way that I would describe as 'natural'.	4	4	5
The instrument's behaviour was consistent. (ie. It responded to identical notes in the same way.)	4	2	4
The instrument was conceptually simple.	4	4	4
The instrument allowed me to create complex musical and visual effects.	3	4	3
I found the instrument interesting and engaging.	5	5	5
I felt in control of the instrument.	4	4	4
Musician P2			
The instrument responded to the music in a way that I would describe as 'natural'.	5	2	3
The instrument's behaviour was consistent. (ie. It responded to identical notes in the same way.)	4	2	4
The instrument was conceptually simple.	4	2	2
The instrument allowed me to create complex musical and visual effects.	5	5	4
I found the instrument interesting and engaging.	5	3	3
I felt in control of the instrument.	3	3	4

Figure 5.11: List of questionnaire responses. Numbers indicate degree of concurrence with statements (1 = strongly disagree, 2 = disagree, 3 = neutral, 4 = agree, 5 = strongly agree).

Question	Mean response	Partial Reflections I	Partial Reflections II	Spheres of Influence
The instrument responded to the music in a way that I would describe as 'natural'.	4.4	3.4	4.3	4.3
The instrument's behaviour was consistent. (ie. It responded to identical notes in the same way.)	4.1	3.7	3.7	3.7
The instrument was conceptually simple.	3.6	3.7	4.0	4.3
The instrument allowed me to create complex musical and visual effects.	4.6	4.1	4.1	4.4
I found the instrument interesting and engaging.	4.6	4.1	3.6	4.0
I felt in control of the instrument.	3.9	3.6		
Standard deviation				
The instrument responded to the music in a way that I would describe as 'natural'.	0.8	1.4	0.8	0.8
The instrument's behaviour was consistent. (ie. It responded to identical notes in the same way.)	0.4	1.3	1.0	1.0
The instrument was conceptually simple.	1.0	1.3	1.1	1.1
The instrument allowed me to create complex musical and visual effects.	0.8	1.2	1.0	1.0
I found the instrument interesting and engaging.	0.5	1.6	0.8	0.8
I felt in control of the instrument.	0.9	1.4		

Figure 5.12: Mean values and standard deviations for each question.

Figure 5.11 shows the responses to the questionnaire statements by the seven musicians who participated in the study. Some patterns can be seen here which were backed up by the qualitative data. The simple statistical analysis shown in figure 5.12 would seem to indicate that in general:

- The musicians agreed that the virtual instruments met the design criteria overall, particularly *Partial Reflections I*.
- The musicians felt that *Partial Reflections I* and *Spheres of Influence* responded in ways they considered more 'natural' than *Partial Reflections II*.
- The musicians found *Partial Reflections I* better facilitated the creation of complex musical and visual effects.
- The musicians found *Partial Reflections I* and *Spheres of Influence* more interesting and engaging than *Partial Reflections II*.
- Musicians had more divergent views on *Partial Reflections II* than the other two virtual instruments. (It had consistently higher standard deviations for each question than the others.)

5.4.3 Audience Interviews

It was difficult to assess the degree to which the virtual instruments met design criterion 8 ('The relationship between live sound, the behaviour of the virtual instrument and the resulting sounds and visuals should be apparent to observers') during the evaluation sessions. This was because both researcher and observer were familiar with the virtual instruments and the musician was performing with them rather than watching a performance. For this reason a series of short interviews was conducted with audience members who attended a performance of *Partial Reflections* at the Queensland Conservatorium on May 17, 2007.¹⁹

The purpose of the interviews was to ascertain if the audience members could discern a relationship between the behaviour of the virtual instrument and the acoustic sounds being produced by the performer. Seven audience members, selected at random, were approached after the concert and a brief interview was conducted during which they were asked about their perceptions of the link between the performer's actions and the virtual instruments, if any. The interviews were recorded on minidisc and later transcribed.²⁰

Findings from the audience interviews

All seven interview participants agreed that they did see a link between the sound produced by the musician and the resulting sounds and visuals created by the virtual

¹⁹A copy of the consent form signed by audience members is included in appendix F.

²⁰The transcript can be found in appendix H.

instrument. However, some interviewees (AM1 and AM5²¹) were unsure whether the virtual instrument was actually responding in real-time or whether the audio-visuals produced by computer were in fact pre-recorded and carefully synchronised to the live performance. Interviewee AM1, for example got the impression that Ben was somehow watching the screen during the performance and following the visuals. When asked whether he thought the visuals were pre-recorded he responded that he was not sure, but that some of it might have been:

A ((Asks what he thought relationship was between what Ben did and what was on the screen.))

AM1 Not too many ideas. It seemed as if he had eyes in the back of his head and seemed to see what the screen was doing and seemed to follow what the screen was doing more than the other way around. I'm not sure how to say it, but I'm fairly sure Ben doesn't have eyes in the back of his head. It seemed as if it was really well written together.

A So do you think... Was your perception that what was happening on screen was kind of pre-recorded and that he was playing along with it?

AM1 um, yes and no. [A You're not sure?] Yeah I wasn't quite sure. I feel some of it might have been but other elements seemed to be a bit on the spot. (0:11)

From the point of view of evaluating whether the audience could see a relationship between the live sound and the behaviour of the virtual instrument, it doesn't really matter whether the audience perceives whether the instrument is responding in real-time or not. Clearly AM1 above perceives a relationship and thus criterion 8 seems to have been met in his case.

Two of the interviewees (AM4 and AM7) mentioned that they felt the connection between live sound and computer generated sounds and visuals was stronger in *Partial Reflections I* than *Partial Reflections II*. AM7, for example, found that she "lost the visual contact a little bit" during *Partial Reflections II*. With *Partial Reflections I*, on the other hand, she "found there was a huge relationship".

AM7 I actually found there was a huge relationship in the beginning between what I was seeing and what Ben was playing.

A So that was the string one [*Partial Reflections I*]?

AM7 The string one yeah. And I was lost, thinking it was like slow motion beads or something like that. Made of something you could see through or something like that. And I found I lost that visual contact a little bit - it wasn't maybe as strong in the second one [*Partial Reflections II*]. But still certainly some relationship. (11:59)

²¹Audience member interviewees are referred to as AM1-7

In summary, the brief interviews with audience members provides additional evidence that criterion 8 was met, with all interviewees seeing a relationship between live sounds and the behaviour of the virtual instrument. *Partial Reflections I* was specifically mentioned by two of the interviewees as making the link more obvious than *Partial Reflections II*. *Spheres of Influence* was not used in this concert, so no data was gathered relating to this virtual instrument.

5.5 Conclusion

This chapter has presented findings from the user studies and a short series of interviews with audience members. As discussed in the introduction to this chapter, the studies were primarily concerned with two questions:

1. How do expert musicians interact with the virtual instruments?
2. How well do the virtual instruments meet the design criteria identified in chapter 4?

The question of how the musicians interacted with the virtual instruments was addressed by using grounded theory techniques to analyse the video recordings of the musicians interacting with the virtual instruments and responding to a series of questions regarding their experiences during the user studies. The observers' notes were also drawn upon in order to provide additional perspectives during analysis.

The research indicates that the musicians interactions with the virtual instruments could be classified into three modes: instrumental, ornamental and conversational. In instrumental mode the musician seeks a high level of detailed control over all aspects of the virtual instrument's behaviour. Musicians taking an instrumental approach essentially see the virtual instrument as an *extension* of their acoustic instrument and want it to respond consistently so that they can trust it during performances.

In ornamental mode, musicians surrender detailed control of the generated sound and visuals and let the virtual instrument create audio-visual layers that are added to their acoustic sounds. Musicians taking an ornamental approach may not pay active attention to the behaviour of the virtual instrument, instead leaving it to its own devices and expecting (or hoping) that it will do something that complements or augments their sound without requiring directed manipulation.

Conversational interaction occurs when musician approaches the virtual instrument as a musical partner. In conversational interaction the musician allows the virtual instrument to 'talk back', at times directly influencing the overall direction of the music. The musical 'balance of power' is in flux as responsibility for shaping musical direction continually shifts between musician and virtual instrument.

It was noted that musicians, while predominantly focusing on the sonic aspects of the virtual instruments, also chose to focus on visual aspects at times. Some evidence was also found for what I have termed a 'physical' focus, in which musicians

were primarily concerned with the (simulated) physical characteristics of the virtual instrument.

In addition, aspects of musician - virtual instrument interaction which impact upon the overall effectiveness of the virtual instruments were identified and discussed. These included:

- The importance of trust in the relationship between musician and virtual instrument, particularly when musicians approach the interaction instrumentally.
- The musicians' desire for the sounds and visuals produced by virtual instruments to be rich and complex and provide scope for a wide expressive range. This was a key factor in all modes of interaction.
- The fact that keeping the properties of the physical model (mass, link tension, etc) relatively static helped enable instrumental interaction but also meant that the response of the virtual instrument was less complex.
- The tendency for instruments using continuous, as opposed to discrete, excitation gestures to more effectively facilitate instrumental and conversational interactions. The discrete excitation approach used by *Partial Reflections II* however was successful in some more ornamental interactions.
- The need for virtual instruments to allow musicians to easily and rapidly shift between instrumental and ornamental modes of interaction in order to enable conversational interaction.
- The importance of facilitating development of a degree of 'intimacy' between musician and virtual instrument in instrumental and conversational interactions.
- The fact that musicians perceived a 'natural' connection between their acoustic sounds and the simulated physical responses of the virtual instruments (particularly *Partial Reflections I* and *Spheres of Influence*).

The question of how well the virtual instruments met the design criteria was addressed by analysis of three sources of data:

- Notes made by an observer present during the user studies.
- Questionnaires completed by musicians who participated in the studies.
- Audio recordings of interviews with audience members attending a concert at which music for two of the virtual instruments was presented.

Key findings were that:

- Overall, the virtual instruments met the design criteria.
- Of the three virtual instruments, *Partial Reflections I* best met the design criteria.

- The musician's experiences with, and responses to, *Partial Reflections II* in terms of the design criteria were more variable than they were with the other virtual instruments.
- The musicians found *Partial Reflections I* and *Spheres of Influence* more interesting and engaging than *Partial Reflections II*.
- Audience members perceived a clear link between the sounds produced by the musician and the resulting sounds and visuals created by the virtual instrument.

In the next chapter a set of revised design criteria for virtual instruments will be presented which draw on these findings.

Chapter 6

A Framework for Criteria-Based Design of Virtual Musical Instruments

6.1 Introduction

In the previous chapter, evidence has been presented indicating the degree to which each of the virtual instruments met the eight design criteria described in chapter 4, section 4.3. In addition, the nature of the musicians' interactions with the virtual instruments have been investigated in depth and a theory which helps understand their experiences has been developed. In this chapter a set of design criteria which draw on findings from the user studies are presented. These criteria provide a framework for thinking about the relationships between virtual instrument characteristics and the modes of interaction.

6.2 Framework for Criteria-Based Design of Virtual Instruments

A major finding of the user studies was that the musicians demonstrated three modes of interaction with the virtual instruments: instrumental, ornamental and conversational. Various characteristics of the virtual instruments influenced which mode of interaction musicians tended to adopt (see chapter 5, section 5.3.3). *Spheres of Influence*, for example, which provided the musician with a high degree of control tended to encourage instrumental interaction. *Partial Reflections II*, on the other hand, which used a mapping based on discrete excitation gestures, gave musicians a greater feeling of separation from the virtual instrument and thus tended to facilitate ornamental interaction. Finally, *Partial Reflections I*, the virtual instrument which musicians tended to feel was most successful, allowed the balance of power between musician and virtual instrument to continually shift, facilitating a more conversational interaction.

The original design criteria which guided the development of the three virtual instruments were developed with a less sophisticated view of musicians' interactions in mind. While it was always the intention to develop virtual instruments that provided complex and interesting interactions, it is true to say that the initial concept was for virtual instruments with a more 'instrumental' character. A more comprehensive understanding of modes of interaction was not developed until the user studies were conducted. The initial design criteria therefore reflect a more simplistic understanding of the range of approaches to interaction that musicians might take with the virtual instruments.

The refined design criteria (shown in table 6.1) are explicitly linked to the three modes of interaction that have been identified. They reflect the author's developing understanding of designing musical applications of this kind, and are intended to provide others with a framework for thinking about and developing their own musical applications.

A number of the design criteria tend to relate to a particular mode of interaction.

Thus the following sections present the refined criteria grouped by the mode of interaction they are most closely associated with. Each of the modes of interaction are discussed briefly with reference to both the original criteria and the findings of the user studies. Finally, several criteria which apply across all modes are presented.

It should be pointed out that depending on the artistic goals of the virtual instrument designer, certain modes of interaction will be emphasised at different times. That is, not all criteria will be met for all virtual instruments. Rather, it is expected that some will be emphasised or de-emphasised depending on the desired interaction aesthetic for particular artistic projects. Thus, these criteria provide a framework which can guide development of virtual instruments and help clarify the relationships between design decisions and the modes of interaction.

Interaction Mode	Criteria
Instrumental	Ensure that the effect of performer actions on the virtual instrument can be accurately predicted.
	Use audio and visual feedback to enhance trust.
	Provide continuous (as opposed to discrete) links between performer actions and virtual instrument response.
Ornamental	Exaggerate, accentuate and/or enhance aspects of musicians' acoustic sounds.
	Operate without requiring direct attention from the musician.
Conversational	Allow the balance of power to shift between performer and virtual instrument during use.
	Respond in ways which are occasionally unexpected, but which are still clearly linked to performer actions.
	Allow the virtual instrument's behaviour to change over time.
All modes	Generate complex, multi-faceted sounds and visuals.
	Respond to the widest possible range of input sounds.
	Provide adjustability in order to support a wide range of different performance situations.

Table 6.1: A framework of design criteria informed by practice and research.

6.2.1 Instrumental interaction

Ensure that the effect of performer actions on the virtual instrument can be accurately predicted

Musicians taking an instrumental approach desire a high degree of control over all aspects of the virtual instruments' behaviour. In order for the musician to be able to develop this level of control, the effect that the actions of the performer will have on the virtual instrument should be able to be predicted. This might mean that if the musician plays two perceptually identical notes on their acoustic instrument, the effect on the virtual instrument will be the same.

This is not to say that the *response* of the virtual instrument should necessarily be the same however. One of the consequences of using physical models as a mediating mechanism between performer gestures and virtual instrument response is that the response of the virtual instrument to a given musical input will change over time. That is, two perceptually identical notes played at different times during the performance may cause the virtual instrument to move in different ways (and therefore produce different sounds). This is because the state of the physical model changes over time. The physical model starts in a resting state and when a note is played it moves as a result of force being exerted upon one of the masses. If the same force is exerted on the same mass before the model has returned to its resting point, the response of the virtual instrument will be different to when it was at rest, because the model is in a different state.

The response will be predictable to musicians however, because playing two identical notes will result in the same *forces* being applied to the same mass. It's just that because the mass will be in motion as a result of the force applied by the first note, subsequent forces will result in different movements and therefore sounds. Thus, the effect of the performer actions are predictable - they always result in the same forces being applied to the physical model - but the virtual instrument response is not always the same.

A virtual instrument which fulfils this criteria will also exhibit transparency of operation for audience members. The degree to which this is important is an aesthetic issue, but I agree with Fels et al. (2002) that in most cases transparency is desirable and is something that computer music in general continues to struggle with.

Use audio and visual feedback to enhance trust

In instrumental interactions, musicians should be able to trust that the virtual instrument will behave appropriately. This is especially important in performance situations where musicians taking an instrumental approach want to be able to guarantee that they can produce certain musical/visual effects on demand. In order to do this, it is desirable to foster a sense of trust between musician and virtual instrument. That is, the behaviour of the virtual instrument should make it clear to the musician that it is 'listening' and responding consistently. This behaviour can have audible and visible manifestations; using both helps give the musician confidence in the virtual instrument.

Provide continuous (as opposed to discrete) links between performer actions and virtual instrument response

The virtual instruments which best supported instrumental interactions in the user studies were those which used aspects of the live audio as *continuous* excitation gestures (*Partial Reflections I* and *Spheres of Influence*). As discussed in section 5.3.3, this helped to give the musicians the feeling that the virtual instrument

was responding to their live sound in its entirety and enabled more nuanced control over its behaviour.

The use of note onset volume as a *discrete* excitation gesture in *Partial Reflections II* tended to decrease the ability of the musicians to approach interaction instrumentally. Because the note onsets had to be distinguished from silence there is a threshold volume below which the musicians' acoustic sounds will not have any effect on the virtual instrument. If they gradually increased their volume then when they reached the threshold point the virtual instrument would suddenly respond. While the sound level at which the virtual instrument would suddenly begin to respond was quite consistent in a technical sense, a number of extraneous factors such as the musician's position relative to the microphone could lead the response to be unpredictable and therefore reduced its support for instrumental control. Musicians therefore tended to describe their interactions with this virtual instrument in ornamental terms.

6.2.2 Ornamental Interaction

Exaggerate, accentuate and/or enhance aspects of musicians' acoustic sounds

By visualising and transforming the musicians' acoustic sounds, the virtual instruments provide the musician and audience with additional perspectives on the live acoustic sound. When approaching the interaction ornamentally, the musicians do not seek to control the virtual instruments in specific ways. Instead, the virtual instruments complement or enhance the musicians' acoustic sounds in ways that the musician does not seek to directly influence. In order to support this mode of interaction the design of the virtual instrument should provide compelling sounds and visuals which help enhance the acoustic sounds.

The experiences of some of the musicians in the study with *Partial Reflections II* indicate that a risk in designing for ornamental interaction is that an overly predetermined and inflexible response can be perceived as musically domineering and can stifle creative expression. Evidence from the user studies suggest that ornamental interaction is a kind of fall-back position that musicians take when instrumental and conversational approaches fail (see section 5.3.1). Thus, designers need to be aware that if a particular virtual instrument seems to be encouraging purely ornamental interactions, it may be a sign that support for instrumental and conversational interaction is poor, not that the design for ornamental interaction is particularly good.

Operate without requiring direct attention from the musician

A defining characteristic of ornamental interaction is that the musician in a sense leaves the virtual instrument to its own devices. Thus, a virtual instrument intended to support this mode of interaction must be able to operate without direct attention from the musician.

6.2.3 Conversational Interaction

Allow the balance of power to shift between performer and instrument during use

As discussed in section 5.3.3, conversational interaction relies on the balance of power between musician and virtual instrument being constantly in flux. At times the locus of control will reside with the musician as they temporarily take a more instrumental approach in order to alter the musical trajectory of the performance. At other times the virtual instrument itself takes a dominant role, providing new musical material which the performer responds to. To enable conversational interaction then, virtual instruments must successfully facilitate and encourage this shifting balance of power.

The use of simulated physical models as an intermediate mapping layer between live sound and computer generated sounds and visuals is a technique that can facilitate conversational interaction. Advantages of using physical models in this way are:

- The behaviour of the models is predictable at a high level, because they behave in a way that is consistent with our experience of the physical world.
- If the physical models feature mobile masses which are coupled (see chapter 4, section 4.8.1), then the response of the virtual instrument can be complex or even surprising while retaining high-level predictability.

Respond in ways which are occasionally unexpected, but which are still clearly linked to performer actions

In conversational interaction, the virtual instrument will at times respond in ways that the performer was not expecting. As discussed above (section 6.2.3) the virtual instrument will at times be dominant in the musical relationship, leading the musician into new musical directions. In order for this to occur the virtual instrument needs to bring in new musical ideas which can surprise the musician in the same way that an ensemble player throws new musical ideas into the mix during performance.

The degree to which the behaviour can be unexpected, while still affording the musician some degree of control over the behaviour of the virtual instrument is a critical design issue. The virtual instruments in this study all produced sounds which were based to some degree on the acoustic sounds produced by the musicians and some musicians expressed the desire for the virtual instruments to be more autonomous - to bring in ideas from “left field” (see, for example, P1’s comments in section 5.3.3). On the other hand *Partial Reflections II*, which was the instrument that musicians felt had the greatest degree of autonomy, was also criticised for being less controllable. When the virtual instrument’s behaviour is continually surprising - when controllability is missing - the risk is that the musician falls back to a disengaged kind of ornamental interaction

Allow the virtual instrument's behaviour to change over time

The simulated physical models used as an intermediate mapping layer for the virtual instruments developed during this project did not change over time. This enabled the responses of the virtual instruments to be consistent and thus facilitate instrumental interaction. The use of physical models (which have a dynamic state) helps to provide somewhat dynamic responses while the virtual instruments retain high-level predictability (see section 6.2.1), but some musicians felt that the range of virtual instrument responses was not wide enough. In order to produce a wider range of sounds using this technique, there are two options:

1. Use more complex physical models; or
2. Have the structure of the physical models change over time.

This is an area for further artistic exploration and research. Creating more complex physical models is likely to involve trading controllability for musical/visual complexity to some extent. However, using simpler structures but allowing them to change over time may provide increased capacity for conversational interaction while retaining a higher degree of controllability and transparency (for audiences as well as musicians). A fourth virtual instrument, *Partial Reflections III*, which allows the structure of the physical model to be altered during performance has been developed. I will describe *Partial Reflections III* in chapter 7, section 7.4.

6.2.4 All Modes of Interaction**Generate complex, multi-faceted sounds and visuals**

Satisfying virtual instrument interactions enable the musician to produce and/or experience a wide variety of rich, complex sounds. The degree to which a sound can be 'rich and complex' while retaining support for some measure of instrumental interaction is difficult to quantify. There is an inherent trade-off between controllability and complexity. If sounds are too complex then the link between performer action and virtual instrument response becomes more difficult to discern and the mode of interaction shifts towards ornamental. In a nutshell, the response must be as complex as possible while ensuring the performer retains the ability to exert some degree of detailed control.

Respond to the widest possible range of input sounds

The range of acoustic sounds produced by the musicians during their interactions with the virtual instruments was extremely wide. Musicians made use of instrumental techniques such as multiphonics, split tones and air sounds. In addition, it was common to use their wind/brass instruments percussively at times (D, for example, used the keys of her clarinet to generate percussive sounds and G used his finger

to flick the bell of his trombone). Virtual instruments which respond only to sounds produced using 'traditional' techniques limit the capacity of musicians to explore new musical areas.

Provide adjustability in order to support a wide range of different performance situations

A pragmatic concern which arose during performances with the virtual instruments as well as during setup for the user studies was the need for the virtual instruments to be adjustable to support a wide range of different physical/acoustic environments. Once the physical model characteristics are determined, minimal adjustability may be provided by a simple slider for adjusting the input level from the microphone. Naturally, during the early stages of virtual instrument design, a larger number of adjustments are likely to be necessary in order to find an optimal settings for the various characteristics of the physical model, such as spring tension, mass values, etc.

6.3 Conclusion

In this chapter, a set of revised design criteria have been presented in the form of a framework which links them to the modes of interaction identified in chapter 5. This framework draws on the enhanced understanding of musicians experiences with the virtual instruments provided by the user studies. It can be seen that there is tension between some of the design criteria, reflecting inherent trade-offs between controllability and complexity of response. By linking the criteria to the modes of interaction, this framework helps to clarify the impact of design decisions on the experiences of musicians

The decision to design to encourage a particular mode of interaction is, of course, an *aesthetic* one. While some performers in the user studies have appeared to prefer one mode or another, it is not the intention here to argue that one mode of interaction is in some way superior to others.

The revised design criteria are informed by artistic practice and user studies. I have attempted to make the criteria as generally applicable as possible but I do not claim that they are 'universal' in the sense that they will apply to artists creating music of any kind with virtual-instrument-like software. This is, in a sense, a limitation of this research, a consequence of its interpretive nature. The virtual instruments were designed in a particular artistic context, and the musicians who participated in the user study had their own pre-existing highly developed aesthetic sensibilities. The findings from the user study and thus these design criteria could be seen as a kind of synthesis of these different aesthetic perspectives and as such represent a snapshot of music practice in this small group of people at this time.

In the following chapter I will consider some implications of the work described in this thesis and describe a new virtual instrument, *Partial Reflections III*, in order to illustrate the likely direction of future work.

Chapter 7

Implications, Ongoing and Future Work

7.1 Introduction

The contributions of this thesis have been presented in previous chapters. They include:

1. Development of a series of virtual instruments which use a unique interaction paradigm;
2. A theory of musician-virtual instrument interaction; and
3. A set of virtual instrument design criteria informed by 1 and 2 above.

The virtual instruments described in chapter 4 use a unique interaction paradigm which uses simulated physical models to mediate between live sounds produced on acoustic instruments and computer generated sounds and visuals. This technique is a key contribution of this thesis. In chapter 5, I described how a set of virtual instruments using this technique facilitated creative interaction in three modes: instrumental, ornamental and conversational. Finally, in chapter 6 I presented a framework of design criteria for virtual instruments informed by the creative work and user studies.

In this chapter I will look ahead and consider some implications of this work. The ideas I present here are more speculative and are intended to illustrate some likely directions for future work in this area.

In chapter 2 (section 2.5), I pointed out that there were several taxonomies which group virtual instruments into various categories based on technical characteristics or on how musical interactions are structured. However, I deferred discussion of these until they could be placed in context with the findings of the user studies. Having presented these findings in chapter 5, in this chapter I will give an overview of some important taxonomies of virtual instruments (or, more broadly, interactive music systems) and place them in context with the modes of interaction I have identified.

In addition, I consider how the use of simulated physical models in musical interfaces may develop in future work and outline some opportunities I see for making improved virtual instruments. Sound synthesis techniques which may help provide a richer palette of sounds for virtual instruments of this type are identified. Next, I describe how using physical models as a mapping layer between acoustic sounds and computer generated visuals might be a fruitful technique for providing feedback to musicians on aspects of their instrumental technique. Finally, I provide a brief description of *Partial Reflections III*, a new virtual instrument developed in late 2007 and intended to further develop some of the ideas presented in this thesis.

7.2 Taxonomies of Interaction with Virtual Instruments

Analysis of the musicians interactions with the virtual instruments in the user studies indicates that they interacted with the virtual instruments in three different modes of interaction: instrumental, ornamental and/or conversational. To the best of my

knowledge there is no previous study which uses rigorous qualitative techniques to generate explicit theories of performers' interaction with new musical instruments. However, a number of taxonomies of musical instruments have been proposed and I will present these in the following sections.

A key difference between the modes of interaction I presented in chapter 5 and the taxonomies I present here is that the modes of interaction were derived from a structured study of musicians. The taxonomies I will discuss in this section, in contrast, usually arose from their author's considerable experience designing and using new musical instruments. I do not suggest that there is more value in either approach but I do point out that studies of the kind I have presented in this thesis can compliment personal experience reports and can be valuable in generating new perspectives.

Other taxonomies of musical interactions with software include:

- Robert Rowe's classification of different instruments: instrument vs. player, etc.
- Todd Winkler's extension of Rowe's taxonomy to include conductor, chamber group, etc.
- Joel Chadabe's continuum from deterministic to non-deterministic.
- Andrew Brown's modes of compositional engagement.

In this section I will briefly summarise these taxonomies and place them in context with the modes of interaction.

7.2.1 Rowe's Taxonomy

Robert Rowe proposes what he calls a "rough classification system for interactive music systems" (Rowe 1993, p.6). His taxonomy is not intended to be categorical but instead to provide a set of dimensions which can help to, "distinguish and draw relations between interactive programs." (Rowe 1993, p. 7)

The first dimension Rowe proposes separates 'score-driven' and 'performance driven' systems:

"Score-driven programs use predetermined event collections, or stored music fragments, to match against music arriving at the input. They are likely to organize events using the traditional categories of beat, meter, and tempo. Such categories allow the composer to preserve and employ familiar ways of thinking about temporal flow, such as specifying some events to occur on the downbeat of the next measure or at the end of every fourth bar." (Rowe 1993, p.7)

An extreme example of a 'score-driven' program would be software which provides automatic musical accompaniment to pre-composed music. The software

SmartMusic (*SmartMusic* 2005), for example, if provided with a score can track an acoustic musician as they perform that music and supply additional accompanying parts in real time. Any variations in tempo will be recognised and adjusted for by the software. Philippe Manoury's composition *Jupiter* is an early landmark composition which makes use of this type of technology.

'Performance driven' systems on the other hand do not track aspects of a musician's performance with respect to a stored score. That is, the system has no expectation that the performer will play a pre-determined composition. My virtual instruments fall into this category.

A second dimension in Rowe's scheme classifies systems as transformative, generative or sequenced. Transformative systems take "musical material and apply transformations to it to produce variants" (Rowe 1993, p.7). Rowe notes that often such transformations are applied to live audio. By this definition then, the virtual instruments I have developed are transformative.

Generative systems, by contrast, "use sets of rules to produce complete musical output from the stored fundamental material" (Rowe 1993, p.7). The source material for these systems then is "elementary or fragmentary" and is embedded within the system rather than being input to it from external sources. An example might be a system which uses an algorithmic process to select notes from a set of scales.

Sequenced systems allow the performer to play back pre-recorded musical fragments, providing controls to vary various parameters such as, "tempo of playback, dynamic shape, slight rhythmic variations, etc" (Rowe 1993, p.7).

The third and final dimension of Rowe's taxonomy separates systems which behave like human musicians or 'players' from those which act like 'instruments'. Systems which follow the instrument paradigm behave like enhanced acoustic instruments:

"performance gestures from a human player are analyzed by the computer and guide an elaborated output exceeding normal instrumental response" (Rowe 1993, p.8)

Systems which follow the player paradigm, "try to construct an artificial player, a musical presence with a personality and behavior of its own" (Rowe 1993, p.8). The human musician effectively plays duets with player paradigm systems.

My virtual instruments would appear to be between these two poles, and just where to place them depends on who is using them. This is partly because of the nature of Rowe's taxonomy: *it categorises musical systems based on the characteristics of the systems themselves, rather than how musicians interact with them.* The modes of interaction I present in this thesis complement Rowe's taxonomy by taking a perspective which is player focussed rather than system focussed. Studies of the kind I have conducted consider technical aspects of the virtual instruments in the context of the impact they have on the experiences of the musicians who use them. In this way they help to bridge the gap between system features and player

experience.

7.2.2 Winkler's Performance Models

In his book *Composing Interactive Music*, Todd Winkler (1998) takes a closer look at the interaction between interactive music systems and the people who use them. He argues that *control* is a crucial issue in shaping the nature of this interaction and proposes four models from three different traditions (symphony orchestra, string quartet and jazz combo) which can suggest various approaches to the design of new systems.

The first of these models is the conductor model, in which the human performer is in control of high-level performance parameters but is not directly responsible for sound creation or detailed shaping of timbre. Winkler notes that the real-world relationship between an orchestra and its conductor is highly complex and shaped by social and historical forces as well as purely musical ones. The computer systems which Winkler cites as examples of conductor systems tend to ignore much of this complexity and interpret 'conducting' to mean giving the performer control over tempo, phrasing and dynamics.

Musical interaction in a chamber group is quite different. In Winkler's 'chamber music model' the interplay between musicians in a group is more sophisticated:

"The interaction in the chamber music model is more complex since several musicians reciprocally influence each others' performance. In a string quartet, for example, even though the first violinist is often considered the effective "leader" (i.e., conductor) of the group, in reality the interplay between musicians demonstrates shared control... This taking and yielding of control, which makes the string quartet so dynamic, is a strong feature built into the composition itself." (Winkler 1998, p.25)

Implicit in both the composer and chamber music models is the assumption that performances are realisations of a pre-composed score. In the 'improvisation model' this is not the case. Winkler argues that, "traditional jazz pieces provide structure and a shared conceptual framework in which musicians interact with each other, influencing both the interpretation of written music (the head), and the improvisation of primary compositional material (solos)" (Winkler 1998, p.25). Winkler notes that the musical and cultural knowledge that the musicians collectively draw on to ensure the music is meaningful in that context is considerable:

"What makes this relationship function to produce music that does not sound like random babbling is that there are a huge number of shared assumptions and implied rules based on years of collective experience. This kind of musical intelligence can be simulated with interactive software on a very simple level." (Winkler 1998, p.26)

Winkler's final performance model is 'free improvisation', exemplified by the free jazz movement of the 1960's. Musical interactions in this context are highly idiosyncratic and unpredictable. Notions of who is "in control" are highly variable, but both performer and computer will influence one another to at least some degree.

Examining the modes of interaction in relation to Winkler's performance models is interesting as they do not map clearly to one another. The difficulty is that the first three performance models (conductor, chamber group and improvisation) imply a macro-structural or phrase-level approach (see chapter 2, section 2.2.6). That is, an approach which is centred on notes and phrases rather than timbre. The computer system references some pre-composed score (in the case of the conductor and chamber group models) or has embedded within it 'rules' which enable it to 'understand' musical phrases in some way in order to respond in a musically sensible fashion (in the improvisation model). These 'improvising' systems contain musical theories concerned with the organisation of musical events at the note level and above which guide them in the generation of their own notes.

Because of this, Winkler's model is less relevant to composers and improvisers who work with timbre. I contend that the modes of interaction framework is broader and can therefore be more usefully applied to a wider range of musical contexts. Instrumental mode can be seen as an extreme version of the conductor model, with the musician directing a performance at a high level of detail. Conversational mode subsumes the chamber music, improvisation and free improvisation models, as these models are all concerned with a shifting balance of power between musician and computer system. The key distinguishing feature between these three Winkler models is not control, in fact, but the degree to which a pre-composed score is involved. Finally, ornamental mode seems to fall outside the scope of Winkler's models. Interactions in this mode could perhaps be seen as occurring at the fringes of the free improvisation model, where the actions of the performer are almost de-coupled from the sounds produced by the computer.

7.2.3 Chadabe's Continuum

Joel Chadabe has described a continuum of virtual instrument behaviour from deterministic to non-deterministic. At the deterministic end, the instrument "is defined by the complete predictability of its output relative to a performer's controls" (Chadabe 2002, p.2). That is, the instrument itself, "provides no information beyond that supplied by the performer" (Chadabe 2002, p.2).¹ Chadabe argues that such an instrument is likely to be suitable for professional (or budding professional) players as opposed to novices, as the performer is required to control every aspect of the sound it produces.

¹Of course no actual instrument, electronic or otherwise, can reach this ideal as any mechanism which links physical action to sound will inscribe something of its character on the sonic results, no matter how small this effect might be.

As we move along the continuum towards non-deterministic instruments, the instruments begin to produce unpredictable sounds. Chadabe argues that instruments with a small amount of indeterminate behaviour can behave like, “a performer’s talented assistants, automatically supplying creative details while the macro-music remains completely under the performer’s control” (Chadabe 2002, p.2). At the non-deterministic end of the continuum the instrument, “outputs a substantial amount of unpredictable information relative to a performer’s controls”, as the performer “shares control of the music with algorithms as virtual co-performers such that the instrument generates unpredictable information to which the performer reacts, the performer generates control information to which the instrument reacts, and the performer and instrument seem to engage in a conversation” (Chadabe 2002, p.2).

Clearly, Chadabe’s degree of determinism continuum has a lot in common with the modes of interaction described in chapter 5. In particular, his description of the conversational approach that non-deterministic instruments stimulate matches the conversational mode of interaction described in this thesis. A finding from the user studies was that the control over the music was indeed “shared” as Chadabe observes, and that a shifting balance of power was the mechanism by which power was shared (see section 5.3.3, p.136). That is, the sharing was not a simple division of responsibilities but a dynamic process of claiming and surrendering control during performance.

It is interesting that the findings of this study seem to indicate that professional musicians do not necessarily seek instrumental control and in many cases wished to avoid it (see P2’s comments in section 5.3.3, p.128, for example). This is not to discount Chadabe’s conclusion, because the virtual instruments I have developed are designed for musicians who use acoustic instruments and are not used on their own. It does indicate however that professional musicians’ desire for full control should not be assumed in all contexts.

The mid-point on Chadabe’s continuum - those instruments with a small amount of indeterminate behaviour which can act as ‘assistants’ - seems similar to the ornamental interaction observed in the user study. However, none of the musicians in the study described any of the virtual instruments as assistants and none of them suggested this as a direction to pursue. The closest was D’s use of *Partial Reflections II* to provide background sounds which she could play counterpoint to (see section 5.3.1, p.118).

In summary, Chadabe’s description of a continuum of control between deterministic and non-deterministic instruments is clearly concerned with many of the same issues as those which arose in the user studies. The findings presented in this thesis confirm many of his arguments but also provide more detail on the nature of the relationship between control and conversational interaction.

7.2.4 Modes of Compositional Engagement

An examination of the work practices of composers who used computerised tools to support their work led Andrew Brown (2003) to identify what he called ‘modes of compositional engagement’, which reflect the ways that composers interact with compositional material on the computer. The five modes are:

“Audient The composer notices, but does not influence, the music.

Director The composer manages technologies to achieve their vision.

Player The composer improvises with their resources to create music.

Explorer The composer undertakes a search for appropriate musical materials and structures.

Selector The composer chooses from computer-generated musical material.” (Brown 2003, ch.7)

This framework was based on observations of composers, not performers working in real-time, but there are clearly modes of engagement which share characteristics with modes of interaction. Brown describes the ‘audient’ mode of engagement as a state where, “the composer stands apart from the compositional process- detached, in a state almost of dis-engagement” (Brown 2003, ch.7). The mode shares with the ornamental interaction observed in the user studies an emphasis on detachment and separation from the behaviour of the software. Likewise, the ‘director’ mode of engagement, characterised by the composer, “consciously manipulating musical materials to shape them into a desirable form” (Brown 2003, ch.7) has similar characteristics to instrumental interaction in which musicians seek detailed control over all aspects of the virtual instruments’ behaviour.

Brown’s ‘player’ mode of engagement shares many characteristics with the conversational mode of interaction:

The mode of composer as *player* acknowledges the medium as contributor to the compositional process and works with it as a partner. When experienced composers are engaged as a *player*, composition appears like a dance where a developed understanding between the partners makes collaboration appear effortless.” (Brown 2003, ch.7, italics in original)

Clearly, the interaction described above has a conversational character. Brown describes the computer as a partner in the compositional process, just as players in the user study (chapter 5) did.

The fact that non-real-time interactions such as those examined by Brown exhibited similar characteristics to the real-time interactions with the virtual instruments studied here indicates that the modes of interaction may be applicable to a wider range of creative activities. Consideration of the relationship between composition and performance, and how characteristics of the tools and/or instruments being used affect the nature of the interactions, may provide further insights into the creative process and how best to support it.

7.3 Physical Models and Musical Interfaces

The use of physical models as an intermediate mapping layer between acoustic sounds and computer generated sounds and visuals is a defining characteristic of the virtual instruments I have developed. In this section I identify some opportunities for future exploration in this area.

7.3.1 Mapping Between Acoustic Sounds and Physical Forces

Because they react to acoustic sounds, as well as generating their own sounds, the virtual instruments I have described in this thesis act as a kind of visualisation of live music. In fact, because the sounds are often mapped to forces which act on the physical models, perhaps it might be more accurate to say that the effect is a 'physicalisation' of their sound. Musicians often felt a natural connection between the sounds they played on their acoustic instruments and the physical manifestation of that sound in the simulated physical environment (see section 5.3.3). J1, for example, was very clear about this in relation to *Spheres of Influence*:

"And it feels like, like when you go like this, it actually feels like a note when you watch it. So you go: *((plays a note))*. To me, in my head, that's *((indicating screen))* a note. Whereas the other one [*Partial Reflections II*] was more of a...[A More abstract somehow?] Ah, just like a line across the page and that, the round ball, just seems, rather than actually turning into a line, it actually seems to stay as a note." (Musician J1)

As he goes on to say, J1 did not, however, feel the same way about *Partial Reflections II*. An area for future work is to investigate the nature of the perceived links between acoustic sounds and forces acting on simulated models. Some musicians spoke of the 'naturalness' or 'organic' nature of the link between their sounds and the physical models, especially in relation to *Partial Reflections I* and *Spheres of Influence* (see section 5.3.3). The fact that both these virtual instruments used sound as a continuous (rather than discrete) excitation gesture seems to be significant here, but this needs further investigation.

Perhaps approaches to music which emphasise the link between music and movement such as Dalcroze *Eurhythmics* (Juntunen and Hyvönen 2004) will help by suggesting strategies for finding links between music and movements which are likely to resonate with musicians. This is an area for further work.

7.3.2 Using Simulated Physical Models to Control Audio Synthesis

During the time I was working on this thesis, a number of projects have emerged which provide more advanced physical modelling simulation environments than the

ones used to create the virtual instruments. The ArtiSynth software developed by Fels et al. (2006), for example, is a “a general purpose biomechanical simulation platform focused toward creating integrated 3D models of the vocal tract and upper airway, including the head, tongue, face, and jaw” (Fels et al. 2006, p.419). The complexity of this undertaking is considerable and, in addition to the use of this software in medical research, it offers a number of interesting possibilities for use in musical contexts. The instruments described in this thesis used extremely simple physical models, partly to contain the complexity of development, but also because it did not seem necessary to use complex models to generate interesting musical responses. In the future however, I anticipate that more sophisticated modelling techniques such as those provided by ArtiSynth provide greater scope for creating virtual instruments which have a wider expressive range.

Another promising software project is *DIMPLE* (Sinclair and Wanderley 2009), which simplifies the creation of 3D physically modelled environments which are coupled to haptic controllers. *DIMPLE* has been designed with visual programming environments such as Max and Pure Data in mind and is especially well-suited to multimedia applications. The extension of the software developed in this thesis into the haptic domain would allow the behaviour of the physical models to become tangible. Currently the models behave in physically plausible ways but users only experience this visually - through the screen. It would be interesting to see what effect allowing them to ‘feel’ the response of the physical model would have on the way musicians interact with the system.

7.3.3 Other Audio Synthesis Methods

Both *Partial Reflections I* and *Spheres of Influence* used additive synthesis to produce sounds. As described in section 4.5.2, in additive synthesis sounds are created by combining simple waveforms - sine waves in this case. Additive synthesis was chosen primarily because changes in the volume of individual sine wave oscillators could be clearly heard and associated with the movement of on-screen masses. That is, the changing timbre could easily be linked with the movements of the simulated physical model. This aided transparency (for both performer and audience) and helped facilitate instrumental interaction.

In *Partial Reflections II*, the movements of the physical model triggered playback of short samples of the acoustic sounds produced by the musician (see section 4.6.2). This meant that once playback of a sample had been initiated, musicians had no further control over the sound produced and this tended to push them in the direction of ornamental interaction (see section 5.3.3).

It is likely that other synthesis methods could be suited to control via simulated physical models. One interesting technique which makes use of low-frequency physical models as a kind of mapping layer between the user and computer generated sounds is *scanned synthesis*. Developed by Verplank, Shaw and Mathews (Verplank




Figure 7.1: Overview of scanned synthesis. The performer manipulates a dynamic system (usually a mass-spring model) to change the shape of the audio waveform, thus directly affecting the timbre of generated sounds (Verplank et al. 2000, p. 368).

et al. 2001, 2000) at Interval Research, this technique allows the performer to directly manipulate the timbre of a computer generated sound. The idea is that a physical model (or other dynamic system) which moves at what they term “haptic rates” (0-15Hz) has virtual forces exerted upon it by the actions of the performer. In response, masses in the model move around and the model changes shape. The shape of the model is then scanned and forms the basis for the wave shape of synthesised audio. The pitch of the resulting sound can be changed by altering the rate that the model is scanned (ie. faster for higher pitches, slower for low ones). The timbre thus depends on the shape of the physical model (see fig. 7.1).

Scanned synthesis is another example how physical models can be used to bridge the gap between human actions and synthesis parameters in virtual instruments. A number of instruments have been developed which make use of this technique (eg. Couturier 2002, Verplank et al. 2002), and this would be a logical area to explore further in my work.

Granular synthesis techniques are another area for further exploration. In granular synthesis timbres are composed by combining large numbers of elementary sound ‘grains’ with durations between 1 and 50 ms (Roads 2001). A mapping layer using physical models could be used to visualise and control various grain parameters during performance, for example.

Likewise, subtractive synthesis, a kind of complement to additive synthesis in which filters are used to alter characteristics of a complex sound, would be relatively simple to implement and may help provide a richer sound palette for virtual instruments of this type.

7.3.4 Instrumental Control and Technical Feedback

Any virtual instrument which affords instrumental control of visuals in real time also has potential as a practice or teaching aid. Once the musician becomes aware of how the virtual instrument responds and develops the requisite trust in its ability to respond consistently, they are able to use the virtual instrument to check up on aspects of their (acoustic) instrumental technique. For example, on the trombone it can be difficult to slur between two notes while maintaining a constant volume. If a trombonist knows that the amount of force exerted on masses in the simulated physical model is proportional to how loudly they play, then the physical responses of the virtual instrument to a scale played on their acoustic instrument will help make bulges or blurts in the sound more apparent.

Of course the effectiveness of this approach is dependent on the accuracy of the audio analysis software, as the physical model really acts as a visualisation of the data it generates. Given this is the case one might ask why volume should not simply be graphed against time. Is there anything to be gained by filtering the data through a physical model? There are a large number of music-learning applications designed to give real-time feedback to musicians (eg. Thorpe 2002, Miller and Schutte 2002, Howard et al. 2007, Smith and Johnston 2008) but the feedback is generally in the form of graphs, spectrograms and other well-known mathematical visualisations. My feeling is that the use of physical models to help visualise (or 'physicalise') aspects of a musician's technique has potential to lead to a more playful and perhaps less judgemental approach to music learning. This is an opportunity which warrants further investigation.

7.4 Partial Reflections III: Exploring Conversational Interaction

Of the three modes of interaction, conversational interaction is perhaps the most interesting and difficult to design for. In this thesis I have demonstrated that the use of simulated physical models can support this kind of interaction. The challenge for future work is to build on this foundation in order to encourage more complex conversations.

In late 2007 a new virtual instrument, *Partial Reflections III*, was developed and used in concert.² This virtual instrument had a number of new features which were intended to improve support for conversational interaction. These included:

- The ability to respond to two performers simultaneously.
- The ability to alter the structure of the simulated physical model during performance.

²This instrument was featured in an improvised performance presented at the Seymour Centre *Sound Lounge* on November 27, 2007.

- A more complex physical model.

I will describe these features in the following sections. The intention is to illustrate how the findings from this thesis are informing ongoing work.

7.4.1 Responding to Two Performers Simultaneously

Like the instruments described in chapter 4, *Partial Reflections III* was designed for use in live performance in collaboration with expert musicians, in this case the clarinetists Diana Springford and Jason Noble. The intention was to create a virtual instrument which would respond to the sounds of both players simultaneously but also independently. That is, the musicians would have separate channels (ie. microphones) through which they could act upon the virtual instrument, but they both interacted with the one instrument. The idea was that part of the musicians' musical conversation would be mediated by the virtual instrument, and that the virtual instrument itself would facilitate conversational interaction with the musicians. We were not interested in supporting purely instrumental or ornamental interactions.

This approach provides increased opportunities for musical conversations to occur. Firstly, as with any duet, musical conversation can occur between the two acoustic musicians. Secondly, each of the musicians can converse musically with the virtual instrument just as the musicians in the user study did. Thirdly, because both musicians are able to interact with the same virtual instrument through their own microphone, they could converse with one another *via* the virtual instrument. In effect the virtual instrument becomes another channel over which conversation could occur. Finally, because the person operating the software (on laptop) is able to alter the structure of the physical model during performance, the interaction between that person and the acoustic musicians may become conversational also.

7.4.2 A More Complex Physical Model

The physical models used for the virtual instruments described in chapter 4 were each made up of only twelve masses. In addition, the only one of these virtual instruments, *Partial Reflections I*, coupled these masses to one another (see section 4.8.1).

In order to provide a more complex and dynamic response and thus stimulate conversational interaction the simulated physical model at the core of *Partial Reflections III* was comprised of 48 masses arranged in a large circle (figure 7.2). Also, as proposed in chapter 6 (section 6.2.3), each of the masses was linked to its neighbour masses. This meant that forces acting on individual masses would propagate through the network of connections and cause linked masses to also move to some degree.

In order that the masses remained in a circle each mass was linked to an invisible




Figure 7.2: The physical model for PR3 was made up of 48 masses arranged in a circle.

mass which was fixed in position.³ Finally, links were put in place which acted only when masses were effectively in contact with one another. The effect of this was to allow masses to bounce apart when they collided with one another.

The simulation itself was developed using Pure Data (Puckette 1997), GEM (Danks 1996) and the Mass-Spring Damper (msd3D) object by Nicolas Montgermont.⁴ Some helper objects written in Python were also used when the visual programming style of Pure Data was found unnecessarily clumsy.

As with *Partial Reflections I* and *Spheres of Influence*, the physical model acted as both visualisation of the musicians' acoustic sounds and as a controller for additive re-synthesis of those sounds.

The fiddle~ object (Puckette et al. 1998) was used to analyse the audio streams coming from the two microphones. This was used to provide continuous data streams containing:

- Current volume.
- Estimated current pitch (and derived from this, pitch class).
- The three most prominent peaks in the harmonic spectrum.

As with *Partial Reflections I* and *Spheres of Influence*, the live acoustic sounds were used as excitation and selection gestures. The current volume was mapped

³If this was not the case then the floating (ie. non-fixed) masses would drift away from their starting positions in the circle as soon as forces were applied.

⁴The Mass-Spring Damper objects are improved versions of the 'Physical Modelling for Pure Data' objects used for the earlier virtual instruments. They allow simpler creation of more complex physical structures and facilitate dynamic changes to model structures while the simulation is running.




Figure 7.3: Screenshot showing effect on the physical model when a middle C is sounded on an acoustic instrument. As with the other virtual instruments, the mass which was currently having force exerted upon it was also made to glow in order to aid transparency of operation.

to the amount of force exerted on the physical model and the current pitch class determined which of the 48 masses would be the target of that force. In order to map the octave onto 48 masses each semitone was simply divided into 4. That is, the mass at the top of the model was associated with the pitch class C, the mass immediately to its right with a C an eighth tone sharper than C, the next mass to the right with C a quarter tone sharper and so on around the circle. Thus, every fourth mass would be associated with a pitch-class from the standard 12 tone equal temperament scale (see figure 7.2). Forces always acted in an outward direction, pushing masses away from the centre of the circle.

An example should help to illustrate how this worked in practice. If a musician played a concert C on their acoustic instrument, the mass at the top of the physical model (ie. at the 12 o'clock position) would have force exerted on it. The amount of force would be proportional to the volume of the sounded note. In response the C mass would be pushed outwards from its resting position while the note was sounding (figure 7.3)⁵. Because each mass in the model is linked to its neighbour masses, the masses closest to the C mass are also dragged out of their resting positions.

7.4.3 Altering Physical Model Structure During Performance

In chapter 6, section 6.2.3, I suggested that making the *structure* of the physical models at the core of the virtual instruments more dynamic - changing the physical properties of the model during performance - was one way to provide more complex

⁵As with the other virtual instruments, the mass which was currently having force exerted upon it was also made to glow in order to aid transparency of operation.




Figure 7.4: During performance the structure of the physical model was altered. This screenshot shows the model after a number of links have been cut and the tension in some springs relaxed.

audio-visual responses while retaining controllability. This approach was used in *Partial Reflections III* and so at several points during performance the structure of the physical model was changed. The approximate points at which this would occur were pre-arranged with the musicians. The changes involved altering tension in some of the links between the masses and cutting others. The effect was that the circle would be seen to gradually lose shape as some of the masses broke loose (figure 7.4). This also resulted in a greater number of collisions between masses and thus a corresponding increase in more percussive sounds generated by the synthesis engine.

Altering the physical model during performance in this way was something not attempted with any of the earlier virtual instruments. Our experience with using *Partial Reflections 3* in performance suggests that this is a technique which can help sustain conversational interaction over longer periods by allowing the virtual instrument to exhibit a wider range of behaviours. The challenge in future work will be in developing techniques (musical and computational) for altering structures in this way while retaining transparency and providing sufficient support for instrumental interactions.

Tache and Cadoz (2006) describe ongoing work towards developing what they term an ‘instrumentarium’ - a library of physical models which have interesting musical properties. In particular they are interested in building simple virtual instruments whose physical structure changes over time. I see this approach as being potentially very fruitful. The four instruments described in this thesis could be seen as the beginning of an instrumentarium. I envisage that future works could involve physical structures which seamlessly transition from one state to another. The structure might begin as a circle which gradually dissolves and reforms as a writhing string, which in

turn transforms into another structure, all the while responding to live sounds played by musician(s). *Partial Reflections III* gives an indication of how such a work might begin.

7.5 Conclusion

In this chapter I have placed the modes of interaction identified in chapter 5 in context with other taxonomies of musical interaction. In addition I have outlined some key areas for future exploration and described *Partial Reflections III*, a new virtual instrument which further explores conversational interaction.

In the next chapter I will summarise the contributions of this thesis.

Chapter 8

Conclusion

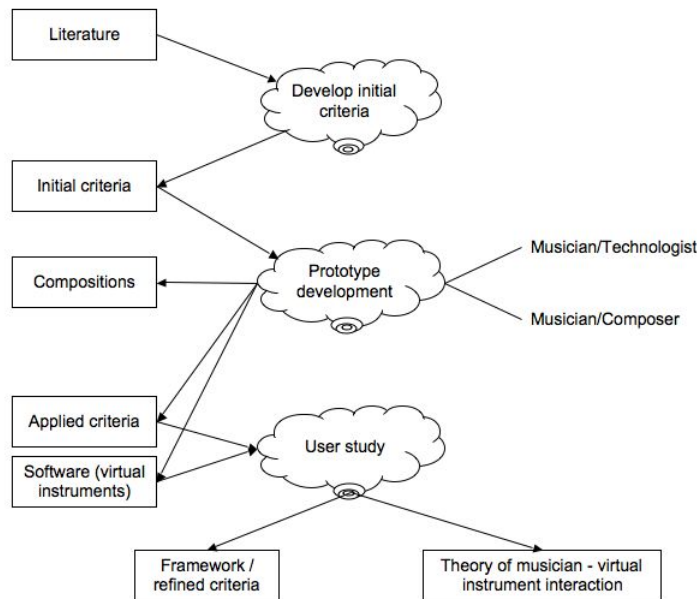


Figure 8.1: Overview of the research project. Processes are shown as clouds, concrete outcomes as boxes.

In the introduction I stated that this thesis is concerned with the design and use of interactive virtual musical instruments which respond to, augment and ‘talk back’ to musicians playing acoustic instruments. In order to explore this area I adopted a research approach which draws on participatory design, action research and design science. The basic structure of the project (shown diagrammatically in figure 8.1) was:

1. Examine the literature to identify initial design criteria for virtual instruments.
2. Develop a series of virtual instruments using iterative, participatory design methods.
3. Conduct a series of user studies to examine how musicians interact with the virtual instruments.
4. Derive a set of design criteria for virtual instruments of this type informed by practice (step 2) and research (steps 1 and 3).

Examination of the literature (chapter 2) revealed the complex nature of designing for musical expression. I concluded that while there were no guaranteed recipes for success, there were a number of commonly accepted guiding principles for virtual instrument designers. Key issues include the need for instrument designers to balance controllability and complexity in order to maintain transparency of operation but also

motivate and inspire performers. I also examined the use of simulated physical models in musical applications and proposed their use as an intermediate mapping layer in order to provide a musical interface which balanced controllability and complexity.

Working with composer/trombonist Ben Marks I developed three new virtual instruments: *Partial Reflections I*, *Partial Reflections II* and *Spheres of Influence* (chapter 4). These instruments use physical models to mediate between live sounds played by musicians on acoustic instruments and computer generated sounds and visuals, as proposed. Compositions by Ben for trombone and the *Partial Reflections I* and *Partial Reflections II* instruments were premiered at the Sydney Opera House Studio in August 2006. *Spheres of Influence* was exhibited at Beta_Space in the Powerhouse Museum, Sydney in September-October 2006.

The literature review and design process resulted in the development of an initial set of eight design criteria. In order to determine how well the virtual instruments met these criteria, the relevance of the criteria, and more broadly examine the experiences of musicians who used them, a series of user studies were conducted. Seven professional musicians used the virtual instruments, made comments, responded to interview questions and completed a questionnaire. These sessions were video recorded, transcribed and grounded theory techniques used to generate a theory of musician - virtual instrument interaction. In addition, a short series of interviews with members of the audience for a concert featuring *Partial Reflections I* and *Partial Reflections II* were conducted to help ascertain whether they found the operation of the virtual instruments transparent.

The key finding from the study was that musicians interact with the virtual instruments in three different 'modes of interaction' : instrumental, ornamental and conversational. In purely instrumental mode, musicians attempt to control the virtual instrument in every respect. In purely ornamental mode the musicians surrender control to the virtual instrument, which adds audio-visual layers to the live performance. Conversational mode is characterised by the virtual instrument 'talking back' to the musician. The balance of power is in flux as responsibility for shaping musical direction continually shifts between musician and virtual instrument.

Additional findings from the user study were:

- Overall, the virtual instruments met the stated design criteria.
- Trust is an important factor in the relationship between musician and virtual instrument, particularly when musicians approach the interaction instrumentally.
- Musicians wanted the sounds and visuals produced by virtual instruments to be rich and complex and provide scope for a wide expressive range. This was a key factor in all modes of interaction.
- Keeping the properties of the physical model (mass, link tension, etc) relatively static helped enable instrumental interaction but also meant that the response of the virtual instrument was less complex.

- Virtual instruments which used continuous, as opposed to discrete, excitation gestures more effectively facilitated instrumental and conversational interactions. The discrete excitation approach used by *Partial Reflections II* however was successful in some more ornamental interactions.
- Virtual instruments should allow musicians to easily and rapidly shift between instrumental and ornamental modes of interaction in order to enable conversational interaction.
- Facilitating development of a degree of ‘intimacy’ between musician and virtual instrument is very important in instrumental and conversational interactions.
- Musicians perceived a ‘natural’ connection between their acoustic sounds and the simulated physical responses of the virtual instruments (particularly *Partial Reflections I* and *Spheres of Influence*).
- Of the three virtual instruments, *Partial Reflections I* best met the design criteria.
- The musician’s experiences with, and responses to, *Partial Reflections II* in terms of the design criteria were more variable than they were with the other virtual instruments.
- The musicians found *Partial Reflections I* and *Spheres of Influence* more interesting and engaging than *Partial Reflections II*.
- Audience members perceived a clear link between the sounds produced by the musician and the resulting sounds and visuals created by the virtual instrument.

Drawing on these findings, the initial criteria which guided the development of the virtual instruments were refined and a framework linking design criteria and the modes of interaction was developed (chapter 6). The framework (shown in table 8.1) is intended to support the design of new virtual instruments and provides a starting point for examining other virtual instruments in relation to the kinds of interactions they are likely to facilitate and encourage.

The framework groups design criteria by mode of interaction. Thus, a virtual instrument which primarily meets the three criteria in the ‘instrumental’ grouping is likely to encourage interactions which have a more instrumental, as opposed to ornamental or conversational, flavour. In addition, a number of criteria which are applicable across all modes are provided.

Finally, in chapter 7, I outlined some implications of this work and suggested some areas for further creative exploration and research. In addition, I presented a new virtual instrument which demonstrates a more sophisticated approach to designing for conversational interaction based on the findings of this research. The ideas and techniques which informed development of *Partial Reflections III* illustrate how the work described in this thesis relates to ongoing creative work.

Interaction Mode	Criteria
Instrumental	Ensure that the effect of performer actions on the virtual instrument can be accurately predicted.
	Use audio and visual feedback to enhance trust.
	Provide continuous (as opposed to discrete) links between performer actions and virtual instrument response.
Ornamental	Exaggerate, accentuate and/or enhance aspects of musicians' acoustic sounds.
	Operate without requiring direct attention from the musician.
Conversational	Allow the balance of power to shift between performer and virtual instrument during use.
	Respond in ways which are occasionally unexpected, but which are still clearly linked to performer actions.
	Allow the virtual instrument's behaviour to change over time.
All modes	Generate complex, multi-faceted sounds and visuals.
	Respond to the widest possible range of input sounds.
	Provide adjustability in order to support a wide range of different performance situations.

Table 8.1: A framework of design criteria informed by practice and research.

Overall, I have demonstrated that the use of physical models in musical interfaces is a powerful technique for enabling expressive and inspiring musical interfaces which effectively support all three of the modes of interaction I have identified. I hope the work I have described here demonstrates their potential and inspires others to explore.

References

- Abrahamsson, P., Salo, O., Jussi, R. and Warsta, J. (2002), *Agile Software Development Methodologies: Review and Analysis*, VTT Electronic, Espoo.
URL: <http://www.inf.vtt.fi/pdf/>
- Arfib, D., Couturier, J.-M. and Kessous, L. (2005), 'Expressiveness and digital musical instrument design', *Journal of New Music Research* **34**(1), 125–136.
- Bally, L., Brittan, J. and Wagner, K. H. (1977), 'A prototype approach to information system design and development', *Information and Management* **1**(1), 21–26.
- Beck, K. (2000), *Extreme Programming Explained: Embrace Change*, Addison-Wesley, Reading, MA.
- Beck, K. and Fowler, M. (2001), *Planning Extreme Programming*, The XP Series, Addison-Wesley.
- Bødker, S., Greenbaum, J. and Kyng, M. (1992), Setting the stage for design as actions, in J. Greenbaum and M. Kyng, eds, 'Design at Work: Cooperative Design of Computer Systems', Lawrence Erlbaum Associates, Inc., Mahwah, NJ, USA, pp. 139–154.
- Boehm, B. (1986), 'A spiral model of software development and enhancement', *SIG-SOFT Software Engineering Notes* **11**(4), 14–24.
- Bongers, B. (2006), *Interactivation: Towards an e-cology of people, our technological environment, and the arts*, PhD thesis, Vrije Universiteit Amsterdam.
- Brooks, F. P. (1987), 'No silver bullet: Essence and accidents of software engineering', *Computer* **20**(4), 10–19.
- Brown, A. R. (2003), *Music Composition and the Computer*, Phd thesis, The University of Queensland.
- Burtner, M. (2002), 'The metasaxophone: concept, implementation, and mapping strategies for a new computer music instrument', *Organised Sound* **7**(2), 201–213.
- Burtner, M. and Serafin, S. (2002), 'The exbow metasax: Compositional applications of bowed string physical models using instrument controller substitution', *Journal of New Music Research* **31**(2), 131–140.

- Buur, J. and Bagger, K. (1999), 'Replacing usability testing with user dialogue', *Communications of the ACM* **42**(5), 63–66.
- Cadoz, C. (1988), Instrumental gesture and musical composition, in 'Proceedings of the 14th International Computer Music Conference', International Computer Music Association, San Francisco, pp. 1–12.
- Cadoz, C. (2002), The physical model as metaphor for musical creation: "pico.TERA", a piece entirely generated by physical model, in 'Proceedings of the 2002 International Computer Music Conference', pp. 305–312.
- Cadoz, C., Luciani, A. and Florens, J. L. (1984), 'Responsive input devices and sound synthesis by simulation of instrumental mechanisms: The CORDIS system', *Computer Music Journal* **8**(3), 60–73.
- Cadoz, C. and Wanderley, Marcelo, M. (2000), Gesture-music, in M. Wanderley and M. Battier, eds, 'Trends in Gestural Control of Music', Ircam, Paris, pp. 71–94.
- Canazza, S., De Poli, G. and Rodà, A. (1997), Analysis by synthesis of the expressive intentions in musical performance, in 'Proceedings of the International Computer Music Conference', International Computer Music Association, pp. 113–120.
- Castagne, N. and Cadoz, C. (2003), 10 criteria for evaluating physical modelling schemes for music creation, in 'Proc. of the 6th International Conference on Digital Audio Effects (DAFX-03)', London, UK.
- Chadabe, J. (1984), 'Interactive composing: An overview', *Computer Music Journal* **8**(1), 22–27.
- Chadabe, J. (2002), The limitations of mapping as a structural descriptive in electronic instruments, in 'NIME '02: Proceedings of the 2002 conference on New interfaces for musical expression', National University of Singapore, Singapore, Singapore, pp. 1–5.
- Charmaz, K. (2006), *Constructing Grounded Theory: A Practical Guide through Qualitative Analysis*, Sage Publications Ltd.
- Choi, I. (2000), Gestural primitives and the context for computational processing in an interactive performance systems, in M. Wanderley and M. Battier, eds, 'Trends in Gestural Control of Music', Ircam, Paris, pp. 139–171.
- Clarke, E. (1985), Structure and expression in rhythmic performance, in P. Howell, I. Cross and R. West, eds, 'Musical structure and cognition', Academic Press.
- Cockburn, A. (2002), *Agile Software Development*, Addison-Wesley.
- Collins, N. (2003), 'Generative music and laptop performance.', *Contemporary Music Review* **22**(4), 67 – 79.

- Cook, P. (2001), Principles for designing computer music controllers, in 'NIME '01: Proceedings of the 2001 conference on New interfaces for musical expression', National University of Singapore, Seattle, Washington.
- Costello, B., Muller, L., Amitani, S. and Edmonds, E. (2005), Understanding the experience of interactive art: lamascope in beta_space, in Y. Pisan, ed., 'Interactive Entertainment (IE2005)', Creativity and Cognition Studios Press, Sydney, Australia, pp. 49–55.
- Couturier, J.-M. (2002), A scanned synthesis virtual instrument, in 'New Interfaces for Musical Expression', Dublin, Ireland, pp. 176–178.
- Crabtree, A. (2003), *Designing Collaborative Systems: A Practical Guide to Ethnography*, Springer-Verlag New York, Inc., Secaucus, NJ, USA.
- Crotty, M. (1998), *The Foundations of Social Research: Meaning and perspective in the research process*, Allen & Unwin.
- Csikzentmihalyi, M. (1996), *Creativity: Flow and the psychology of discovery and invention*, Harper-Collins, New York.
- Danks, M. (1996), The graphics environment for max, in 'Proceedings of the International Computer Music Conference', pp. 67–70.
- Dobrian, C. and Koppelman, D. (2006), The 'e' in nime: musical expression with new computer interfaces, in 'NIME '06: Proceedings of the 2006 conference on New interfaces for musical expression', IRCAM - Centre Pompidou, Paris, France, pp. 277–282.
- Ehn, P. and Kyng, M. (1992), Cardboard computers: mocking-it-up or hands-on the future, in J. Greenbaum and M. Kyng, eds, 'Design at work: cooperative design of computer systems', Lawrence Erlbaum Associates, Inc., Mahwah, NJ, USA, pp. 169–196.
- Ericsson, K. A. and Simon, H. A. (1993), *Protocol Analysis: Verbal Reports as Data*, revised edn, MIT Press, Cambridge, MA.
- Fels, S. (2004), 'Designing for intimacy: creating new interfaces for musical expression', *Proceedings of the IEEE* **92**(4), 672–685.
- Fels, S., Gadd, A. and Mulder, A. (2002), 'Mapping transparency through metaphor: towards more expressive musical instruments', *Organised Sound* **7**(2), 109–126.
- Fels, S. S. (2000), Intimacy and embodiment: Implications for art and technology, in 'Proceedings of the ACM Conference on Multimedia', pp. 13–16.
- Fels, S. S., Lloyd, J. E., van den Doel, K., Vogt, F., Stavness, I. and Vatikiotis-Bateson, E. (2006), Developing physically-based, dynamic vocal tract models using artisynth, in 'Proceedings of ISSP 06', pp. 419–426.
- URL:** <http://hct.ece.ubc.ca/publications/pdf/fels-et-al-ISSP2006.pdf>

- Floyd, C. (1984), A systematic look at prototyping, in 'Approaches to Prototyping', Springer-Verlag, pp. 1–18.
- Frederiksen, B. (1996), *Arnold Jacobs: Song and Wind*, Windsong Press Limited, Gurnee, Illinois.
- Gergen, K. J. (1999), *An Invitation to Social Construction*, SAGE Publications Ltd., London.
- Glaser, B. G. (1978), *Theoretical Sensitivity*, The Sociology Press.
- Glaser, B. G. (1992), *Basics of Grounded Theory Analysis*, Sociology Press.
- Glaser, B. G. (1998), *Doing Grounded Theory: Issues and Discussions*, Sociology Press.
- Glaser, B. G. and Strauss, A. L. (1967), *The discovery of grounded theory: strategies for qualitative research*, Aldine de Gruyter, New York.
- Goto, S. (2000), Virtual musical instruments: Technological aspects and interactive performance issues, in M. Wanderley and M. Battier, eds, 'Trends in Gestural Control of Music', Ircam, Paris, pp. 217–230.
- Grønbaek, K. (1989), 'Rapid prototyping with fourth generation systems - an empirical study', *OFFICE: Technology and People* 5(2), 105–125.
- Gurevich, M. and Treviño, J. (2007), Expression and its discontents: Toward an ecology of musical creation, in 'Proceedings of the 2007 Conference on New Interfaces for Musical Expression (NIME)', pp. 106–111.
- Henry, C. (2004a), Physical modeling for pure data (pmpd) and real time interaction with an audio synthesis, in 'Sound and Music Computing '04', Paris.
- Henry, C. (2004b), 'Using physical modelling for pure data (pmpd) with an audio and video synthesis'.
- URL:** <http://puredata.info/community/projects/convention04/lectures/tk-cherry/pmpd.pdf>
- Henry, C. (2006), Basic physical modeling concepts, in F. Zimmer, ed., 'Bang: Pure Data', Wolke Verlag, Hofheim, pp. 49–59.
- URL:** <http://pd-graz.mur.at/label/book01/bangbook>
- Hevner, A. R., March, S. T., Park, J. and Ram, S. (2004), 'Design science in information systems research', *MIS Quarterly* 28(1), 75–105.
- Hewett, T., Czerwinski, M., Terry, M., Nunamaker, J., Candy, L., Kules, B. and Sylvan, E. (2005), Creativity support tool evaluation methods and metrics, in B. Shneiderman, G. Fischer, M. Czerwinski, B. Myers and M. Resnick, eds, 'Creativity support tools: A workshop sponsored by the National Science Foundation', National Science Foundation.

- Holmes, T. (2002), *Electronic and Experimental Music*, 2nd edn, Routledge.
- Howard, D. M., Brereton, J., Welch, G. F., Himonides, E., DeCosta, M., Williams, J. and Howard, A. W. (2007), 'Are real-time displays of benefit in the singing studio? an exploratory study', *Journal of Voice* **21**(1), 20–34.
- Howard, D. M., Welch, G. F., Brereton, J., Himonides, E., Decosta, M., Williams, J. and Howard, A. W. (2004), 'Winsingad: a real-time display for the singing studio', *Logopedics Phoniatrics Vocology* **29**(3), 135–144.
- Hunt, A. and Wanderley, M. M. (2002), 'Mapping performer parameters to synthesis engines', *Organised Sound* **7**(2), 97–108.
- Hunt, A., Wanderley, M. M. and Kirk, R. (2000), Towards a model for instrumental mapping in expert musical interaction, in 'Proc. International Computer Music Conference'.
- Hunt, A., Wanderley, M. M. and Paradis, M. (2003), 'The importance of parameter mapping in electronic instrument design', *Journal of New Music Research* **32**(4), p429 – 440.
- Järvinen, P. (2007), 'Action research is similar to design science', *Quality and Quantity* **41**(1), 37–54.
- Jordà, S. (2002), 'Fmol: Toward user-friendly, sophisticated new musical instruments', *Computer Music Journal* **26**(3), 23–39.
- Jordà, S. (2004), 'Instruments and players: Some thoughts on digital lutherie', *Journal of New Music Research* **33**(3), 321–341.
- Jordà, S. (2005), Digital Lutherie: Crafting musical computers for new musics' performance and improvisation, PhD thesis, Departament de Tecnologia, Universitat Pompeu Fabra.
- Juntunen, M.-L. and Hyvönen, L. (2004), 'Embodiment in musical knowing: how body movement facilitates learning within dalcroze eurhythmics', *British Journal of Music Education* **21**(2), 199–214.
- Kemmis, S. and McTaggart, R. (1988), *The Action Research Planner*, 3rd edn, Deakin University Press, Geelong.
- Krefeld, V. and Waisvisz, M. (1990), 'The hand in the web: An interview with Michel Waisvisz', *Computer Music Journal* **14**(2), 28–33.
- Kyng, M. (1988), Designing for a dollar a day, in 'CSCW '88: Proceedings of the 1988 ACM conference on Computer-supported cooperative work', ACM, New York, NY, USA, pp. 178–188.

- Leman, M. (2007), *Embodied Music Cognition and Mediation Technology*, The MIT Press.
- Levin, G. (2000), Painterly interfaces for audiovisual performance, Master's thesis, Massachusetts Institute of Technology.
- Levin, G. (2008), personal communication with Andrew Johnston.
- Levin, G. and Lieberman, Z. (2004), In-situ speech visualization in real-time interactive installation and performance, in 'NPAR '04: Proceedings of the 3rd international symposium on Non-photorealistic animation and rendering', ACM, New York, NY, USA, pp. 7–14.
- Levin, G., Lieberman, Z., Blonk, J. and La Barbera, J. (2004), 'Messa di voce', World Wide Web electronic resource.
URL: http://www.tmema.org/messa/report/messa_report_600.pdf
- Lewin, K. (1946), 'Action research and minority problems', *Journal of Social Issues* **2**, 34–46.
- Lewis, C. (1982), Using the 'thinking-aloud' method in cognitive interface design, IBM Research Report RC 9265 (#40713), IBM Thomas J. Watson Research Center, Yorktown Heights, NY.
- Lewis, G. E. (2000), 'Too many notes: Computers, complexity and culture in voyager.', *Leonardo Music Journal* **10**(1), 33–39.
- Luciani, A., Cadoz, C. and Florens, J. L. (1994), 'The crm device: a force feedback gestural transducer to real-time computer animation', *Displays* **15**(3), 149–155.
- Machover, T. (1992), Hyperinstruments: A progress report, Technical report, MIT Media Laboratory.
- March, S. T. and Smith, G. F. (1995), 'Design and natural science research on information technology', *Decision Support Systems* **15**(4), 251–266.
- Menzies, D. (2002), 'Composing instrument control dynamics', *Organised Sound* **7**(3), 255–266.
- Miles, M. B. and Huberman, A. M. (1994), *Qualitative Data Analysis: An Expanded Sourcebook*(2nd Edition), Sage Publications, Inc.
- Miller, D. G. and Schutte, H. K. (2002), 'Documentation of the elite singing voice'.
URL: <http://www.vocevista.com/contents.html>
- Miranda, E. R. and Wanderley, M. M. (2006), *New Digital Musical Instruments: Control and Interaction Beyond the Keyboard*, A-R Editions, Inc.
- Mogensen, P. (1992), 'Towards a prototyping approach in systems development', *Scandinavian Journal of Information Systems* **4**, 31–53.

- Momeni, A. and Henry, C. (2006), 'Dynamic independent mapping layers for concurrent control of audio and video synthesis', *Computer Music Journal* **30**(1), 49–66.
- Moore, F. R. (1988), 'The dysfunctions of midi', *Computer Music Journal* **12**(1), 19–28.
- Mulder, A., Fels, S. and Mase, K. (1997), Mapping virtual object manipulation to sound variation, in 'USA/Japan intercollege computer music festival', Tokyo, Japan, pp. 63–68.
- Myatt, T. (2002), 'Strategies for interaction in construction 3', *Organised Sound* **7**(2), 157–169.
- OED Online (1989), *The Oxford English Dictionary*, 2nd edn, Oxford University Press.
URL: <http://dictionary.oed.com/>
- Oliver, W., Yu, J. and Metois, E. (1997), The singing tree: Design of an interactive musical interface, in 'Proc. Designing Interactive Systems', ACM Press New York, NY, USA, Amsterdam, The Netherlands, pp. 261–264.
- Omodei, M. and McLennan, J. (1994), 'Studying complex decision making in natural settings : using a head mounted video camera to study competitive orienteering.', *Perpetual and motor skills* **79**, 1411–1425.
- Paine, G. (2002), 'Interactivity, where to from here?', *Organised Sound* **7**(3), 295–304.
- Paradiso, J. A. and O'Modhrain, S. (2003), 'Current trends in electronic music interfaces.', *Journal of New Music Research* **32**(4), 345–349.
- Portillo, M. and Dohr, J. H. (1994), 'Bridging process and structure through criteria', *Design Studies* **15**(4), 403–416.
- Puckette, M. (1991), 'Something digital', *Computer Music Journal* **15**(4), 65–69.
- Puckette, M. S. (1997), Pure data, in 'Proceedings of the International Computer Music Conference', pp. 224–227.
- Puckette, M. S., Apel, T. and Zicarelli, D. D. (1998), Real-time audio analysis tools for pd and msp, in 'International Computer Music Conference', International Computer Music Association, San Francisco, pp. 109–112.
- Rapoport, R. N. (1970), 'Three dilemmas of action research', *Human Relations* **23**, 499–513.
- Roads, C. (1996), *The Computer Music Tutorial*, MIT Press, Cambridge, MA, USA.
- Roads, C. (2001), *Microsound*, MIT Press.

- Rovan, J., Wanderley, M., Dubnov, S. and Depalle, P. (1997), Instrumental gestural mapping strategies as expressivity determinants in computer music performance, *in* 'Proceedings of the AIMI International Workshop KANSEI - The Technology of Emotion, Genova', pp. 68–73.
- Rowe, R. (1993), *Interactive Music Systems*, The MIT Press, Cambridge, Mass.
- Royce, W. W. (1970), Managing the development of large software systems: concepts and techniques, *in* 'Proceedings, IEEE WESCON', IEEE Computer Society Press, pp. 329–339.
- Rubine, D. and McAvinney, P. (1990), 'Programmable finger-tracking instrument controllers', *Computer Music Journal* **14**(1), 26–41.
- Ryan, J. (1992), Effort and expression, *in* 'Proceedings of the 1992 International Computer Music Conference', International Computer Music Association, San Jose, California, pp. 414–416.
- Schnell, N. and Battier, M. (2002), Introducing composed instruments, technical and musical implications, *in* 'New Interfaces for Musical Expression (NIME-02)', University of Limerick, Department of Computer Science and Information Systems, pp. 138–142.
- Schön, D. A. (1983), *The Reflective Practitioner: How Professionals Think in Action*, Basic Books.
- Settel, Z. and Lippe, C. (2003), Convolution brothers' instrument design, *in* 'Proceedings of the 2003 Conference on New Interfaces for Musical Expression (NIME-03)', Montreal, Canada.
- Simon, H. A. (1981), *The Sciences of the Artificial*, 2 edn, The MIT Press.
- Sinclair, S. and Wanderley, M. M. (2009), 'A run-time programmable simulator to enable multi-modal interaction with rigid-body systems', *Interacting with Computers* **21**(1-2), 54–63. Special issue: Enactive Interfaces.
- Smalley, D. (1997), 'Spectromorphology: explaining sound-shapes', *Organised Sound* **2**(2), 107–126.
- SmartMusic* (2005).
URL: <http://www.smartmusic.com/>
- Smith, G. and Johnston, A. (2008), Interactive software for guitar learning, *in* S. Wilkie and A. Hood, eds, 'ACMC08 SOUND : SPACE Proceedings of the Australasian Computer Music Conference 2008'.
- Smith, J. O. (2004), 'Virtual acoustic musical instruments: review and update', *Journal of New Music Research* **33**(3), 283–304.

- Sol, H., G. (1984), Prototyping: A methodological assessment, in 'Approaches to Prototyping', Springer-Verlag, Berlin.
- Stewart, M. D. (1987), *Arnold Jacobs: The Legacy of a Master*, The Instrumental Publishing Company, Northfield, Illinois.
- Suchman, L. A. (1987), *Plans and situated actions: the problem of human-machine communication*, Cambridge University Press, New York, NY, USA.
- Suchman, L. A. and Trigg, R. H. (1991), Understanding practice: Video as a medium for reflection and design, in J. Greenbaum and M. Kyng, eds, 'Design at Work: Co-operative Design of Computer Systems', Lawrence Erlbaum Associates, Hillsdale, NJ.
- Tache, O. and Cadoz, C. (2006), Using evolving physical models for musical creation in the genesis environment, in 'Sound and Music Computing', Marseille, France.
- Tanaka, A. (2000), Virtual musical instruments: Technological aspects and interactive performance issues, in M. Wanderley and M. Battier, eds, 'Trends in Gestural Control of Music', IRCAM - Centre Pompidou, Paris, pp. 389–406.
- Thorpe, W. (2002), 'Visual feedback of acoustic voice features in singing training'.
URL: http://www.med.rug.nl/pas/Conf_contrib/Thorpe/Thorpe_bio_touch.htm
- Verplank, B., Gurevich, M. and Mathews, M. (2002), THE PLANK: Designing a simple haptic controller, in 'New Interfaces for Musical Expression', Dublin, Ireland, pp. 33–36.
- Verplank, B., Mathews, M. and Shaw, R. (2000), Scanned synthesis, in 'Proceedings of the 2000 International Computer Music Conference', International Computer Music Association, Berlin, pp. 368–371.
URL: <http://www.billverplank.com/ScannedSynthesis.PDF>
- Verplank, B., Mathews, M. and Shaw, R. (2001), 'Scanned synthesis', *The Journal of the Acoustical Society of America* **109**(5), 2400.
- Vlimki, V., Pakarinen, J., Erkut, C. and Karjalainen, M. (2006), 'Discrete-time modelling of musical instruments', *Reports on Progress in Physics* **69**(1), 1–78.
- Wanderley, M. M. and Orio, N. (2002), 'Evaluation of input devices for musical expression: Borrowing tools from HCI', *Computer Music Journal* **26**(3), 62–76.
- Wessel, D. and Wright, M. (2002), 'Problems and prospects for intimate musical control of computers', *Computer Music Journal* **26**(3), 11–22.
- Winkler, T. (1998), *Composing Interactive Music: Techniques and Ideas Using Max*, MIT Press, Cambridge, MA, USA.
- Wishart, T. (1996), *On Sonic Art*, Routledge.

Woods, D. and Fassnacht, C. (2007), 'Transana v2.22', **Madison, WI: The Board of Regents of the University of Wisconsin.**

URL: *http://www.transana.org*

Zivanovic, A. (2005), The development of a cybernetic sculptor: Edward ihnatowicz and the Senster, *in* 'C&C '05: Proceedings of the 5th conference on Creativity & cognition', ACM Press, New York, NY, USA, pp. 102–108.

Appendix A

Memos

The process of memoing is described in Chapter 3 and in section 5.2.3 of chapter 5. All the memos that were produced during analysis of the video recordings of the user study are shown here.

These memos are reproduced here in order to demonstrate how grounded theory analysis was conducted during this study. It should be noted that memos are 'quick and dirty' made during analysis and as such they are not intended to be representative of the final conclusions of the research. Rather they help document my emerging understanding of the situation under study as the analysis progressed.

Memos : Musician D

D talks of the use of string spheres "fixing" a particular kind of awareness. One that is more visually driven. ("I realise I'm not listening like I would if I wasn't watching it if you know what I mean. [A That's interesting.] Because it fixes a certain kind of awareness I suppose.")

D wants the s/w to change over time, so that response is more complex. ("So maybe the timbre, because I think of some timbres as dirty or scungy, maybe a textural shift or a shape shift or something.")

For D simple visuals are 'legible'. ("I like the fact that it's simple so you can sort of read it. It's legible.")

D speaks of initial experiments to figure out what was happening. Experimental mode? Switching between instrumental/conversational?

D talks about the s/w responding "in the moment" and contrasts musical approaches which could work with this. One approach focussed on creating "short elements" which the s/w responds to "within a period of seconds" and the big picture approach which uses the elements the s/w responds to ("overtones, speed and dynamic") as building blocks in "bigger structure".

The static nature of the s/w is an issue for D. ("But also I suppose there's a rhythm. It's more about duration or something than a beat or something. It doesn't respond to that, so if I maybe set up a beat, it feeds it back a little but it doesn't transform that. It transforms my sound quite a bit but it doesn't transform the beat or duration much.") This also means that the s/w responses always have a particular shape. (D. So that means what I do isn't going to be amplified infinitely, which is good, because it means new material can come in. But that does dictate that it will have a certain shape I suppose. The event will have a certain shape and that will be a... A A very...The shape will always be the same.)

Using string spheres made D more aware of partials in her sound. ("One I sort of became more aware of how the quality of a note was related to the pitches in it. You sort of become a bit more aware of that.")

D found that playing string spheres when they were set to output square wave suited the percussive approach to the clarinet (using keys only, etc). ("The sound quality just seemed to match the percussive elements better.")

D often acknowledges trade-off between complexity and simplicity. She likes the simplicity and 'elegance', at the same time as desiring more complexity. D pointed out that the simplicity of string spheres put the onus on performer: "the simplicity of it means that you have to think more about what you're doing. So in that sense maybe I wouldn't change that."

2/7/2008 3:57:30 pm
simplicity challenges musician.

Software being dictatorial can stimulate musician. I think J2spoke about the software 'poking' / provoking the musician. D says, "That sings da da da da da has quite a dictatorial effect but that's good for me. It makes me think about things to do." Seems to be about the power balance. Giving power to the software is not all bad! eg. D makes interesting use of the word 'provokes'. The software has "dictatorial effect" which "makes me think about things to do". It "provokes more of a duet sort of mode of thinking".

D describes doing "duet" with spinner. I wonder is this the same as the conversational approach J2spoke of? D says: "It (string spheres) was more doing something to create a sound whereas that (spinner) you sort of respond to it and set up something and then play along with it. So it is more of a duet." For D, the fact that the set up of spinner allowed her to set it up into a certain state – by playing loudly directly into the mic – and then play along with it away from the mic, playing smooth phrases, etc. was important.

With spinner D seems to feel that visuals are more separate to the audio than with string. It's hard to express this idea: "Whereas that (the visuals) is something from outside the... It's like a different sort of... It's just like a feedback rather than a... I can see how that responds to me, but not how it has a relationship with me that's sort of a function... Ah– it's been a long day!" Keywords here to me are: "outside", "different", "just...feedback". She can see "how [the visuals] responds to me" but can't see "how it has a relationship with me..." The visuals respond only to D's playing – not to the wholeness of the sounds created by D and the computer together. Could contrast this with string spheres in which movement of string is 'two-way'. ie. it responds to live music, but also directs the generated sounds.

D comments that the link between performer musical gestures and computer generated sounds/visuals does not have to be direct. In fact, with spinner, the fact the the link could be severed at times was quite useful to her. ("Yeah cause I could do something quite different to what it was doing but it still responded to me. So even when you put it on a more subtle reception, I could still play a smooth melody over the top while it was still doing lots of activity underneath. Whereas with the other one it responded to everything I did– the speed and the... So it's almost the more simple response of this one allowed you to do something different in contrast with it, whereas the other one responded almost too much.") With string, "you couldn't get apart from it. Whereas this one [spinner] you can get apart from it a bit and go, 'ok well I'm going to do something else that it can't follow me on'."

2/7/2008 3:50:45 pm

Being able to get separation at times– go off and do counterpoint.

For D a drawback of spinner was that "you couldn't differentiate between pitches and I suppose, like the sound, the speed seemed to be set up from the very beginning". She also found that with spinner she didn't have to "produce as much interesting stuff, because the relational thing made it more interesting just to start with". That is, the independent nature of the spinner allowed her to

play in a more 'duety' kind of way. The 'close coupling' (aj's words) put more onus on her to play musically interesting things.

This issue of being able to 'read' the response of the software comes up often with D. ("I suppose the simplicity of it – you could always imagine it doing more but then you wouldn't be able to read it as well and it wouldn't...your perception of it would be too confused or something.")

Consistency was less of an issue for D, although she noticed when it was lacking. A response which was consistent ON AVERAGE seemed fine for her. ("I couldn't figure out what made it different. I was trying to do different things to see if that changed it, but I didn't work it out. So... That wasn't a bad thing though, because it generally if I tried it a few times it came back to what I expected it to do, so...You if you try hard enough or long enough it will do what you want.")

Interesting thought: D didn't seem to speak of 'trust' as an issue. I wonder why.. Consistency was not so important for her I don't think, maybe there's a link there..?

D talks of taking a "technological approach to see what I could do to make the visuals change" with circle spheres:("This was a button, buttons [A A mouse or something like that.] sort of thing. Yeah, yeah. I was learning to use it.")

D found circle spheres drove her to take a more visually driven approach, partly because the arrangement of spheres in a circle was more congruent with her view of musical pitches and their relationships, and partly because of the colour change in response to timbre change. ("So I guess it made me aware of the, sort of the timbre thing, the timbre element, because of the visual response to that.")

Visual awareness vs. aural awareness (see around line 39,40). Perhaps this is another dimension of the three interaction modes...so could have 'instrumental visual', 'conversational visual', etc, or 'instrumental aural', etc. Or a mix?

Memos : Musician G

Consistency and control were big issues for G. (See 'to sum up' clip.) "what I'm finding challenging is if I have an effect that I like, to be able to guarantee that same effect on demand."

The virtual instruments seem good at "highlighting a lot of stuff" – making particular sonic effects more visible and audible (reinforcing). But, the sound colours (of circle spheres) are limited. He asks whether the other two virtual instruments have different sounds colours.

For circle spheres, G did not feel it was necessary for him to look at the visuals while playing. c/w J2, who felt that the visuals helped enhance the trust. Later, when using spinner, he found visuals WERE needed to be sure he was getting what he wanted. (Because spinner was less consistent when we tested it with G?)

It seems that with G the focus was on exploring and controlling the software rather than being influenced by it. eg. "I was concentrating more on the effects my tone would have on the response. I was really concentrating on what I could do to change the colours and if I could play in a more mellow or brassy way, how that would affect it."

G also talks about 'effects'. In contrast to J2' thinking of some of the interactions as 'conversations'.

"The screen was very helpful to make sure that I articulated consistently." – ties in with 'trust' idea of J2. The visuals confirm that computer is listening, responding.

NEW KEYWORD: TRUST (as both G and J2 mentioned this idea)

Memos : Musician J1

Interesting that he didn't want any introductory explanation– preferred to learn on his own.

J1 didn't seem to try to find the limits/explore extremes initially. Straight away he played single notes in middle range, arpeggios.

SUGGESTION: The faster you play, the faster the repetitions (ie. speed of spinners) become. (around 10:30)

J1 prefers not looking at screen for spinner. (36:00)

From J1's session I got the idea of physical play. eg. When he describes pushing string spheres back and then filling them in with colours as they swing back. (1:22:00 approx.) This seems to me to tie in with the idea of counterpoint. Setting up the model – getting it into an interesting state – and then playing over the top of it. This leads to suggest an important characteristic of the virtual instruments:

CHARACTERISTIC: Instruments should have 'momentum', which allows the musician to get the instrument into an interesting/useful state and then play in counterpoint with it. Simple example would be setting up a drone which can be improvised over. Thinking about it, maybe this isn't a characteristic as such but rather a way of playing or compositional/improvisational technique which the virtual instrument can help or hinder.

NEW KEYWORD: CONTROL (1:25:00 approx.) Seems separate to 'consistency' keyword here.

NEW KEYWORD: FLEXIBILITY (9:18 approx.) "I find the fact that you can't actually control the tempo of it..." He doesn't like the spinner because it dictates the tempo of everything you play. "If you could build up something a little more ethereal..." Less mechanical maybe?

The pitch change in spinner really seemed to throw J1 (see around 26:00 +). Because he was unfamiliar with the software's response, he didn't feel confident in his ability to control it. He seems to associate this lack of control with the "fear of not making it sound right." (c. 25:00) This is something that as a classical musician he has trouble with. ie. It's difficult for him to "let go and not understand it" (c. 24:45)

NEW KEYWORD: ENJOYMENT (c.35:00) While J1 did not show huge enjoyment at this point, I realised that there was no keyword for someone enjoying playing! The 'interest' keyword is more for someone finding the software intellectually interesting – a different feeling to enjoying it in my view.

NEW KEYWORD: LIKE/DISLIKE (c.35:00) Like 'enjoyment' I realised I didn't have any keyword to indicate the musician liked something (or not)!

NEW KEYWORD: DISTRACTION (c.36:00) J1 found the visuals for spinner 'took his attention away too much'. ie. They were distracting for the performer. However, he was interested in working with the visuals in another way. ("...if I played with it a lot I would love to create pictures" c. 48:00) In contrast perhaps to J2, he did not want to be influenced by the visuals in performance, but rather would figure out how to make the software do certain visual things in advance.

2/7/2008 4:42:04 pm

Diana also found the spinner visuals distracting – or at least they didn't create a 'duet' feeling in the same way that the sounds did.

Like J2, J1 found the spinner kept him in a single tonal area. "Like I would have been changing keys a lot more often and like, you can't, unless you want that particular effect and I didn't. So that's why I chose the fanfare because you can just play arpeggios and still make it sound good." (c. 48:00)

J1 pointed out that he "...lost interest in the dynamics thing pretty quickly" (c. 57:25) because the software "...wasn't doing anything". Continuous control – so the software listens to your sound in its entirety – seems important. DESIGN CRITERIA: Maybe 'continuous link' between live sound and model is the design criteria.

"in fact it does create a totally different creativity" – the impact of visuals on music-making. (c.1:09:00)

DESIGN CRITERIA: maybe 'forgiving' – J1 mentions how the spinner is not very forgiving of split notes, etc, whereas with string spheres, "it's out and it goes and that's it" (c. 1:11:00)

PATTERN OF USE: (c.1:15:00) My interpretation: When I was watching it, the visuals dominated. Given time I would work out strategies to make it more unified. Initially though I was exploring the visual effects.

NEW KEYWORD: NATURAL. (c.1:15:00) Ties in with the design criteria – 'natural response'. J1 talks of the 'liquidy' nature of string spheres. "the movement in it wasn't rigid at all."

Musical strategies: J1 found string spheres was more 'scale-based' than the spinner.

"it actually feels like a note when you watch it." (c.1:38:00) Feels that circle spheres response (and also string spheres I think) is more appropriate, fitting than response of spinner.

2/7/2008 4:55:37 pm : Audio-visual congruence. Ties in with Diana, J2.

NEW KEYWORDS: SIMPLICITY and COMPLEXITY (c.1:42:15) Tie in with criteria: Instrument is simple but allows musician to create complex effects.

"so it's in some ways changing an F as a note on the page to F as an object that actually has a life." (c.1:45:58) Interesting way of thinking about process of making music on brass instrument and using the s/w in teaching. As musicians we take note on page and translate it into physical movement. Software takes the resulting sound and maps it to physical forces – gives it another kind of life.

Memos : Musician J2

Considers the spinner to be 'an effect' (see around 16:00) whereas string spheres has more complex response- 'creates other harmonies'. This is a really key idea I think.

I added some more keywords based on the conversation at around 15:50. The constructs seem to be something like: 'person -- effect' or perhaps 'direct response -- conversation' (to use J2's terms).

J2 really dislikes the spinner. The 'pulsing thing' - the pulsing, continuous audio is what most annoys him. It is too big a 'thematic idea', too 'dominating'. It was not interactive enough. He felt that no matter what he did ('I change intensity, I change range...'), '...its response was basically the same.' (see around 54:30 onwards). 'The other ones seemed very interactive in a way I could understand.'

It's interesting that J2 talks about 'weight' as being important. (I wonder if this is the fact that the spinners are so light and respond so quickly, have no damping, etc which makes them 'unnatural'. The string and circle spheres have much more physical substance in a way. ie. the masses are heavier! Babies would be startled by spinner but not by string/circles perhaps?)

J2 thought the good instruments were 'natural'. eg. 'It seemed to be very direct and related to the natural world. Yeah! Very much so. That one more than anything visually related to the natural world to me.' (see around 1:00:00)

'Clicking points' in spinner created some interest (1:01:00). Design ideas/opportunities arising from bugs.

Relationship between live sound and generated sound/visuals is still critical with J2. eg. One of his criticisms of the spinner was that he 'never sensed [the relationship between live sound & generated sound] because of the pulsing.' The constant pulse (and short sample length?) meant he didn't pick up on this relationship in the same way he did with string spheres.

Suggestion for spinner - much more direct link between dynamic and speed of rotation - not constant. (1:14:00)

For J2 'equilibrium' was important. Mentioned the need for the virtual instruments to 'come to rest'. All the instruments do this to some extent but spinner continues to make noise even when there is no input sound- it never stops! The effect of 'gravity' on string spheres was appealing to J. It made apparent the inevitability of coming to rest? SUGGESTED DESIGN CRITERIA: Instruments should have a resting point that is readily apparent. They can take time to reach this point. Gravity can be used to indicate inevitability of coming to rest.

Visuals enhance the trust: "The visual reinforces the audio response. It lets you know you're getting a response. So in a sense it could enhance the trust." (c.33:30)

Also, an aspect of spinner that did not enhance trust for J2 was the fact that although the visuals came to rest, the audio kept going. (c.48:31)

The 'pulsing' and 'delay' of spinner seemed to be the thing that most annoyed him. (c. 1:05:00)

(c.1:26:00) "They're playing for each other and they're playing each other." I really like this phrase. The idea that the musicians are playing each other– not 'with' each other.. Playing each other like an instrument... I like this because it suggests that musicians can be 'instruments'. Implies you can have a 'conversation' with your instrument and you can 'play' other musicians. A good musician can be like an instrument – a good instrument can be like a musician.

"We're constantly as an improvising musician trying to find the balance between structure and creativity." (c.1:36:21) Relates to the balance between control and surprise in software for musical expression I think.

"So this kind of program or instrument is a very different relationship if you've got a really developed improviser." Perhaps improvisers think of the interaction as a conversation whereas non-improvisers think of it as an instrument to be controlled?

J2 mentioned the difference between Bachs and Yamahas. Bach is more difficult initially but has bigger expressive range. (c.1:47:53)

J2uses the term 'positive discomfort' to describe the feeling of using circle spheres. It was 'stimulating' him to try harder because it was less harmonically complex than string spheres. It was 'provocative'. (c.1:54:13)

Memos : Musician M

M liked spheres much more when the sensitivity was higher. ie. When there was a lot of onscreen movement in response to sounds. (see around 44–45 minutes in video).

M preferred spinner when sliding pitch was removed. This gave him more control over the harmonies generated by the software. ie. He could load up pitches into the spheres to play chords etc. With sliding pitch this was not practical.

Also preferred it when links from spinner to central point were made invisible. The visuals were more 'random' than which he liked. Perhaps by 'random' here he means rich, complex...?

Found spinner much more aurally involving. Was far less interested in visuals for this one.

Linda observed (around 1'38") that M found circle spheres best for teaching, string spheres more interesting visually and spinner more interesting aurally. Do M's comments during use back this up? I agree circle was best for teaching but I thought he found that spheres and spinner were both interesting aurally.

"I was just doing a waltz with it." – M talking about spheres. To me this shows greater engagement and a move towards the 'conversational' experience. On the other hand: "You'd have little moments where you'd just manipulate the audio" (Talking about composing for string spheres) Phrases like this seem to indicate a more 'instrumental' approach.

M suggested that the string spheres 'suspension point' could be moved around in order to get more interesting shapes.

Seems to like spinner best..

M found that he could make the circle spheres do interesting things visually by playing with a wide range of different attacks, etc but the sonic result "wasn't really cutting it" (0:26:41) in terms of producing music that was interesting to listen to.

Variation, richness was important to M. The fact that string spheres and circle spheres moved/responded in the same way to different sounds was seen as limiting. He was often seeking "more variation". eg:

M So that would be another way of getting more variation into the visual aspect. But I do like the playback. I'd have to sit down and think about it a lot, but I reckon I could get something a little bit more presentable from that than the first one.

A Are there any types of music that this responds best to?

M I felt like this was less limiting than the first one. Do you know what I mean? So, playing legato, doing the ((sings be bop bup)) jumpy stuff, you know, all worked fine sort of thing. I liked the feedback better on that too. The sonic, sort of, part of it, yeah. So I think that's less limiting, yeah.

A So if you were composing a piece for it can you imagine what characteristics that piece would have? If you can say...

M Yeah I mean I've sort of like got a vague idea what it would be... But you know, I'd be thinking more in terms of like, this is something new that we really haven't... You know I wouldn't try to compare it to other things. Yeah, a new style of composing. I mean, of course it's always going to have some sort of 1960's electronic sort of feel about it or whatever but yeah...

A I'm thinking about the trombone part. What would that look like do you imagine? Lot's of pauses, lot's of...?

M Yeah I'd imagine, um... Well, with this one you could do lots of legato but you'd have like pauses in it and you'd have little moments where you'd just manipulate the audio thing... Um, what's coming back – the audio. I'd go for some harmonies of some sort by doing that. And then I'd be doing things above it. But, you know I could imagine that working with non-legato things and, you know, actually quite fast moving or slowing with that one. Yep, so I can see lots of possibilities I guess, yeah. (0:56:53)

and

M Perhaps something like that. That could be interesting. I think there is something really cool that we could do here I reckon. I don't know what it is yet but I can see something.

M I can see something really cool, but just to have some variance here, because always going up in pitch will get tiring.

A So the basic idea is good but it needs more variation before you could really...

M Yeah, because you're just locked into the same sort of thing and you'd end up... (1:14:30)

Like D, M used the counterpoint technique with spinner:

A The staccato obviously works. Why the legato?

M To contrast against it. Because you can set up like your harmonies and stuff with the staccato and then play stuff legato over the top.

A Ok, a bit like Ableton Live. [M I guess so, yeah, it was a little bit like that.] Did you play differently than you normally would?

Memos : Musician P1

It's interesting that P1 wanted the string spheres to be 'adversarial' at times, rather than 'complimentary'. Ties in with the idea that the s/w should be able to surprise the user.

Another interesting point that P1 raises is that when you know the software well, you can make it appear to be more intelligent than it really is. "One of the things you can do when you know it really well is you can make it appear more like it is being an improvising partner..."

Also:

P1 Well you can, and when you know it well you as a performer can make it look like that better too. Because I've done those pieces that Martin Wesley-Smith wrote for clarinet and tape which dates...It's a pretty classic piece from the 70s. And the trick with that, because it's just on a CD, is if you're really absolutely on top of it rhythmically you can make it sound like it's bouncing off you. Technically it's very difficult to do because you've got to be incredibly precise yourself.

A So you don't feel that it's bouncing off you it's just you give the impression to the audience.

P1 Yeah if it's not an effective performance the performer just sounds out the effect of the recording.

Seeing the software as a 'musical' instrument as opposed to a visual one (or hybrid). Visuals were 'an interesting fall out...' (see below)

"A Did you find that at any stage were you trying to play in a particular way to influence the visuals?

P1 No I didn't no. No I was more trying to introduce...[A The sounds...] Yeah, the visuals was like an interesting fall out from what I was doing but I wasn't trying to actually do something. [A Kind of incidental really.] Yeah. It's a bit like you know how harmony is incidental in Xenakis' music. It's an outcome the process rather than the other way around."

L suggests 'adaptive' as a word to describe the kind of interaction that P1 was interested in:

L Maybe that's the word you were looking for. 'Adaptive'. It adapts to... You were talking about it responding and you described it as 'complimentary'. But maybe if it was more adaptive rather than complimentary, that would be more interesting.

P1 Yeah!

L I'm thinking that it could adapt in ways that would be perhaps unexpected.

P1 It would be interesting if it had a sort of process where it picked up on certain idiosyncratic things that you tended to do as an improviser. Because you do - you tend to have little ideas that you keep doing. And if the same thing is repeated often enough then it would start to remember that.

Perhaps this 'adaptive' style could lead to more 'conversational' interactions?

Very hard to get P1 to comment on spinner! Questions did not really seem to engage him.

P1 mentions that the drive to make things look good may detract from the music: "Because sometimes things that look good might not sound good." G also mentions this. Leads to:
NEW KEYWORD: musical-visual conflict

P1 wants more complexity in spinner ("maybe some way of making it do some long sounds as well.") so that response can be richer.

P1 mentions that getting students to practice advanced technique is good for fundamental technique- software can encourage them to try these advanced techniques. Also goes on to talk about how important it is that students exaggerate in order to get musical ideas across to audiences. Software can help to remove focus on technique. ("But the fact that you're focused on an image as well tends to make you. You wouldn't be as focussed on problems with the instrument necessarily.")

Interesting that during initial play with circle spheres P1 didn't approach the interaction 'instrumentally'. ie. He wasn't trying to control the colours ("didn't get that far") rather he "just noticed that [it] changed". Could indicate 'conversational' interaction perhaps? Certainly not a 'controlling' mode of interaction. eg.

A Was the colours...did you find you had control over them? Like could you make a purple sound wildly (?) or was that not an issue?

P1 I didn't quite get that far... I just noticed that I changed. I wasn't connecting a particular colour with a particular sound. I hadn't got that sophisticated yet.

P1 talks about the music he played for circle spheres was "disconnected" (and I seem to remember Josh talking about music being "disjointed"). Impact of visuals on music- tendency to play things that look good (sometimes at expense of sounding good).

Several musos mentioned that the music was "disconnected", "more disjointed", etc., because of the impact of the visuals.

Pete had faint praise for the independence of the masses in circle spheres (compared w/string spheres): "I think I quite like the fact that it moved independently."

P1 liked colour change in circle spheres- thought it encouraged full use of clarinet's alternative fingerings.

P1 mentions that s/w encourages musician to respond to visual images. His mention of 'response' here perhaps indicates response to the software 'talking back' - conversational mode? ("it encourages you to sort of respond, you know, with sound to a visual image which is something you actually don't often do really. You normally respond to an acoustic image.")

P1 would like more complexity in all of them– and the ability to change over time, but also recognises the trade off between complexity and control ("The more complex they are, of course in some ways the harder it would be to perform with because the relationship between...You'd need more skill in the way that you actually manipulated it.")

Memos : Musician P2

P2 mentions the software "paralleling" the live sound as a "first step". Once it can do that then it "opens up scope for more interactivity". Ideally it needs to be as "flexible as a live instrument". I think he's saying that the paralleling is equivalent to control. It is a necessary first step, but is not enough to sustain interest. Going on, P2 talks about string spheres "accurately reflecting what I was doing" and being "sensitive" but also says that trying to control it is "never really interesting to me anyway." (line 22). He goes on to say that "if it can actually complement" his playing "in a kind of reciprocal way, that's the thing that is sort of interesting to me" (line 22). Also, he mentions it should be 'not autonomous but...'

P2 liked the fact that he could play more or less 'intuitively' and he could be confident that it would react in a way that was 'interesting'. He finds 'control' and 'moving the dot around the screen' not particularly interesting and is looking for a more conversational style. For him, for conversational interaction to occur, he needs to feel trust that the software will do interesting things without requiring instrumental style instructions. In the performance "mental state" (line 26) he didn't want to have to think about how he was going "to adjust what I am doing to change it to make it [the software/model] better" (line 26).

Sine wave sound was "one dimensional" (line 29)

The novelty of the virtual instrument was important. The fact that the virtual instrument "changes the context with which I'm using the tools that I use all the time" (ie. trumpet techniques) is a good thing which leads him into "interesting areas of performance" (line 35).

Interesting that while James liked the fact that string spheres came to rest, P2 found that this was too predictable. P2 wanted it to be more "generative itself to varying degrees", to have "an element of randomness to it", to "manipulate what it was that I was doing in a more generative way" (23:49, line 49). He talks of the "stupid word" "trust". He needs to be able to trust that the machine is "really accurate, and then I can kind of trust the randomness of it".

P2 mentions several times that beyond accuracy, latency issues is "going to the next step. Some sort of interactivity or randomness or..." (24:56, line 51). Accuracy, "paralleling" performance is the first step, followed by more generative processes. Implication I think is that generative processes on their own are not enough – first step is necessary. Interesting though that he also says of "trust", that "that's not a high priority" and it's a "stupid word" (24:56)

It's interesting that P2 is unsure whether it is necessary for the audience to be able to discern a relationship between live audio and computer generated sound/visuals. "It's hard to know [whether it's important] isn't it?" (26:00, line 55)

P2 wanted the sound to be "flexible" and to have "nuance". He wanted the software to be more "autonomous" (new keyword). It seems that this autonomy is like a higher degree of unpredictability. Of all the musicians, P2 seemed more prepared to surrender control to the software processes. "I wouldn't want to control it, I would want it to somehow control itself. (A yeah.) That's the thing that's interesting about this to me. Otherwise it's just a MIDI keyboard, I mean I don't want to play a MIDI keyboard, I want this thing, I would want it to be as, you know, far more interesting than what I can think of, you know!" (30:59)

"It could have something to do with a combination of pitch and dynamics and, um, somehow it comes up with a 70% chance it's going to be a nice sound and 30% chance it's going to be granulated or whatever." (31:50) He seems quite comfortable with less predictable, loose relationship between his playing and the computer's playing.

He is not quite talking about a completely ornamental interaction though. He wants the "continuous feedback loop of interactivity", but seems more comfortable with interaction closer to the random, ornamental side.

Interesting suggestion: chain virtual instruments together (ie. output of one linked to input of another). 36:15

The rhythm of the bouncing balls in circle spheres was too "predictable" for P2. "The tempo's the same all the time, isn't it?" (45:40)

There are some contradictions implied in P2's comments at times. He's not so interested in control but he comments that "it was difficult to make just little subtle movements with the balls in this one". He doesn't actually say he prefers to be able to make subtle movements, but all the same...

He says that "some people" might like to make patterns in their playing visible with circle spheres, but this is not so interesting to him. (47:11) The one on one mapping for him seemed dull. He says that "I lose interest in it pretty rapidly, as a performer" (47:44)

Strangeness, randomness, autonomy – put performer in position of having to find control. Music-making is about the process of finding, developing control. Skill in improvisation is being able to respond quickly to challenging musical situations. Instruments themselves throw out challenges – physical ones. They are unpredictable to some degree. Professionals have less unpredictability – they are trained to deal with everything that comes up. But they still search for challenges, to put them back in the situation where they need to find the way again, get challenged to 'escape' from a situation. Demonstrate mastery? Very interesting discussion around these issues around 55:00 and on.

"I'm rather more interested in what it can do badly, or the error of it, or you know, like what sort of new aesthetic thing can come from it. And often it's the edges of what the thing can do, or maybe (L Mmmm) what it's not supposed to do. Whether you can find interesting areas. And it's often they're the bits that you can't plan for, or you don't necessarily, um, you can't develop those. They're

just found." This idea is interesting– you can't design or plan for the interesting areas at the edges – performers have to find them somehow. This could be a variation on the instrumental approach – the musician is still approaching the software as an instrument but rather than controlling it, they are trying to find the boundaries, glitches, the things it doesn't do well.

P2 doesn't see the trumpet as having autonomy – even though he consciously tries to find the edges, extremes of the instrument (spit, air sounds, etc), because "I have to control it. It's inert." He does acknowledge they are "similar" though because what he is doing is mapped to control the virtual instrument.

He jokes that he would love to have "a valve on this instrument that I could just, that it would just go down by itself while I was playing. That would be fantastic." He talks about how, "you do things as a player to, kind of, encourage error. You know, you can do things just to get you in a different spot, musically."

This use of "error" or glitches to deliberately get himself out of the comfort zone is interesting. For me, it is possible to see this as a perpetual process of struggle against circumstances. Musical expression in this aesthetic area is about getting yourself into difficulty and responding.

"what would be interesting would be to be able to change the level of engagement that you have with the thing, in terms of controlling it, or, to have some, to have different levels of interactivity, and I wouldn't necessarily want to control what level of interactivity I was on at that point. (L Mmmm) I would rather be able to interact with it spontaneously." This is close to talking about modes of interaction I think. Different levels of engagement that the musician can move between. He doesn't want to control the change of mode but would prefer to be able to be free to change modes spontaneously.

Lizzie makes an interesting comment that P2's "leaning towards chance and giving away a bit of control, starts even with your relationship with your trumpet".

P2 talks about a "spectrum" ranging from absolute instrumental control to almost total non-control. Great quote: "To me, it's the whole spectrum of, it's being able to have absolute control over everything that happens in your sound, from the moment it begins to the moment that it ends, and any brass player will tell you that that's incredibly difficult to, so there's that, um, having control over all those minute details of playing the instrument. That's one end of the spectrum. At the other end of the spectrum I play things that I'm never quite sure, I alter my techniques so that I'm not quite sure how it going to come out. And that's the spectrum. I think all people that are interested in improvisation, and interested in their instrument, have to have that spectrum, have to be able to have complete control over the instrument and be able to be interested in not having control over. Or at least being in situations where, could be through performing with other people interacting with other performers. That can be an element of relinquishing control."

The colour change- making the harmonic structure of the sound visible was interesting and mildly engaging for P2 initially, but he tired of it, mainly because the "sound thing's a little bit annoying" and the movement of the balls (and resulting 'tempo' of the sounds) was too simple and consistent. "you look away and you know exactly how it's going to be moving"

For spinner, P2, "felt that I was locked into it. I felt that I was going to have to... All of my playing was going to get this thing to do something that I would like. Whereas the other two, particularly the first one, I didn't have any sense of that. I felt that I could actually ignore it if I wanted to, or interact with it."

P2 found spinner to be "predictable, and repetitious". Also, it only responded to "10 percent" of what he was doing and ignored the rest:

"I felt that the way that it was reacting with this, was in a pretty limited way, in that it was only, like, so much of what I was doing, it wasn't interested in at all. (A yeah) So I felt that it was only going to be interested in ten percent of what it was that I doing, and to me it was the ten percent that wasn't that interesting anyway. So it was, um, the attack of the note is something that often programmers go for, like they want to get that attack of the note, but as a performer sometimes that's not, you know, the most interesting bit. Or, it's just one of many bits that, do you know what I mean?"

While P2 wanted the earlier virtual instruments to be more autonomous, he found that while spinner did have a kind of autonomy it was too limited and repetitious. "It definitely had its own behaviour, but the behaviour was predictable, and repetitious, so I lost interest in it. (L Mmmm) That's a crucial element in all of it. You want the behaviour to be interesting behaviour, not just (L Bad, like a naughty kid)(laughs) or the Sorcerer's Apprentice, where it's just uncontrollable and not that interesting. That's the difficult part I think, finding the right balance." Once again, it comes back to the balance between controllability, interest and autonomy.

P2 came back to string spheres at the end of the session. When he played without looking at it, he noticed that he originally "was focussing more on the visuals and the audio thing, and I think the visuals are more interesting than the audio thing. So I was definitely latching on to those." Partly this was because "so much of the aural perception for me is probably subliminal anyway, so it's not something that I need to necessarily... You know I can quite easily be totally engaged with it, but unconsciously so, whereas the visual thing for me is something that, with regard to responding to music, you know, it's not something that, doesn't come naturally for me to see something moving to the sound of the instrument, so I was focussing on that." The aural perception is almost habitual, but visual responses to music were more novel for him.

Appendix B

Coding Scheme

The ongoing development of a coding scheme is a key part of the grounded theory method. An example illustrating this process is provided in chapter 5, section 5.2.4. The following 'Keyword Summary Report' generated by Transana shows the coding scheme developed for this research. It should be noted that the development of a coding scheme is part of a process to *facilitate* analysis; it is not the end result of analysis (see section 5.2.4).

Keyword Summary Report

behaviour

can I make it...?

User tries to make the software do something specific.

discovery

Musician discovers something about their playing/music-making as a result of using the software.

exaggerates/sensitivity

Software highlights or is particularly responsive to some musical characteristic or style. Can also refer to 'sensitivity' of software.

finding limits

Musician deliberately tries to find what the limits are. (eg. plays very high, low, soft, loud...)

physical play

Musician describes pushing, prodding or otherwise playing in a physical way with the virtual model.

playing

Musician plays with the virtual instrument.

request

Musician requests an adjustment to the software or eval. procedure. Differs to 'suggestion' in that a request is for an immediate change, whereas a suggestion is for a later change. Suggestion is likely to be more substantial and request more trivial.

software affects music

The music the musician makes is affected in some way by the behaviour of the software.

sonic exploration

Musician explores the audio possibilities of the system. eg. tries to make new sounds, attempts to modify the sound in particular ways, etc.

visual exploration

Musician tries to make different visual effects.

what does it do if I...?

Musician explores the behaviour of the software. eg. Makes various sounds to see what effect they have on the sound and/or look of the software.

what's going on

Musician explores what the software does.

feature

audio

Reference to the audio produced by the software.

bug

A bug in software is identified - possibly leading to ideas for future directions. Bug can be good as well as bad!

mapping

Mapping between musical features and behaviour of the software.

model

Reference to an aspect of the physical model underlying the software.

visuals

Reference to some aspect of the visual output.

Q&A

about piece

Tell me about the piece of music you wrote. Why did it have these characteristics?

characteristics of playing

While you were interacting with the work, did you become aware of any particular characteristics of your playing?

done differently?

Do you have some comments or questions about what we've done today? Is there anything we should have done differently?

easy/hard

Do you have some comments about how easy or hard it was to write for or perform with the software?

play differently

Did you play differently today than you normally would?

should have asked

Is there anything I should have asked you?

sound

Do you have some comments about the sound produced by the software?

suggestions

Do you have any suggestions or proposals for how we might improve or extend this software?

uses

Can you think of any uses for this software?

visuals

Do you have some comments about the visual display?

response

autonomy

Musician discusses the level of autonomy that the software/virtual instrument has.

complexity

Evidence the software enabled/did not enable the musician to create complex effects.

consistency

Comments on software consistency. ie. Given the same input the software behaves in the same way.

control

Comments on trying to control the software and/or overall sound.

distraction

Evidence of being distracted by an aspect of the software.

ease of use

Comments on how easy/hard the software was to use.

effect vs. conversation

Musician describes playing with software as being on the continuum between playing with an 'effect' (such as delay, etc) and playing with another person/having a musical conversation.

enjoyment

Musician shows evidence of enjoying playing with the software.

expression

Comments on how effective/ineffective the software is at facilitating expression.

flexibility

Comments on the flexibility of the software.

how it works

User perceptions of how the software works (in a behavioural sense - not a technical one).

instrumental

Musician describes or uses the software in an 'instrumental' fashion. That is, plays it (or tries to play it, or talks of it) as if it were a musical instrument.

intensity

Musician observes that some aspect of the software/experience is intense.

interest

Musician comments that something is interesting/not interesting.

like/dislike/preference

Musician shows evidence of liking/disliking some aspect of the software, or of having a preference for software/characteristic/setting.

limitation

Musician identifies a limitation of the software.

link

Comments on links between input/behaviour/visuals/audio output, including mapping.

metaphor

Musician makes connection between the behaviour of the software and other items of personal experience.

motivation

Musician is motivated/not motivated.

musical strategies

Musician describes/discusses musical or compositional strategies.

musical-visual conflict

Evidence of/comments on the conflict between visually-driven and musically-driven playing. (eg. P1: "things that look good might not sound good")

natural

Evidence of the software responding a ways that musician considers 'natural'.

novelty

Musician comments on novelty of the system. That is, it's newness or strangeness.

power balance

Comments on the balance of power between musician and software.

proficiency

Musician comments on becoming proficient using software.

simplicity

Evidence the musician found the instruments conceptually simple.

suggestion

Musician makes a suggestion.

surprise

Musician is surprised by some aspect of the software.

trust

Evidence that an aspect of the software engendered a feeling of trust - or not...

use/purpose

Uses or potential uses for the software.

Appendix C

Episode Report

When doing grounded theory analysis in Transana, I broke the video recording of each musician's session with the virtual instruments into sections or 'clips'. Each clip contains material which was found interesting or relevant. In order to facilitate the 'constant comparison' of incidents in the data, each clip is labelled with codes ('keywords' in Transana parlance). This appendix contains a sample 'Episode Report' (for musician J2), generated by Transana. It shows:

- the transcripts of each of the clips created;¹ and
- the codes (keywords) assigned to each clip.

This report is included here to help illustrate how analysis was carried out.

¹The numbers next to each line in the transcript are time codes, which, unfortunately, cannot be removed from the report.

Transana Episode Report

Episode: j2

Clip: spinner: playing

Collection: j2

Time: 0:03:39.2 - 0:06:23.7 (**Length:** 0:02:44.5)

Episode Transcript: aj-notes

Clip Transcript:

J plays softish long Bb, cresc. (no response) Increases volume until spinner responds. Improvises. Melodic, smooth, contained.

Clip Keywords:

behaviour : playing

Clip: So, yeah...enjoyable

Collection: j2

Time: 0:06:35.5 - 0:07:40.3 (**Length:** 0:01:04.8)

Episode Transcript: aj-notes

Clip Transcript:

J Without much enthusiasm. It's enjoyable. The regular pulse creates that trance-like idea. I don't know if it's... It's obviously my taste/aesthetic but when I know that the sound is going to continue like that for a long time I tend to stay in the same tonal area, you know. So it's like a long... I don't know... I don't want to over simplify it but it's like a delay - a specific delay with an unending...no decay it feels like. So, yeah...enjoyable. But I tend to play, at least at 10 in the morning, in the same tonal area.

Clip Keywords:

behaviour : exaggerates/sensitivity

feature : audio

response : complexity

response : enjoyment

response : how it works

response : musical strategies

Clip Comment:

In the light of future comments, his statement that 'it's enjoyable' should probably not be taken at face value.

Clip: string playing

Collection: j2

Time: 0:08:58.2 - 0:13:12.0 (**Length:** 0:04:13.9)

Episode Transcript: aj-notes

Clip Transcript:

J improvises w/software. Much sparser than previous improv. Uses wider pitch range.

Clip Keywords:

behaviour : playing

Clip: that's lovely, lovely

Collection: j2

Time: 0:13:12.0 - 0:15:27.0 (**Length:** 0:02:15.0)

Episode Transcript: aj-notes

Clip Transcript:

J mmm...that's lovely, lovely. The fact that there's still no decay at the end.

«<808073>A You like that or you don't like that?

«<810942>J um, I prefer that it has a length.

«<818047>A Yeah. Are you talking about when it comes to complete rest like now [J Yeah], you would like it to be totally silent?

«<824514>J At time..within the time.

«<826962>A Yes, gradually [J yeah] come to rest.

«<832664>J So there's a tangible... (...point where it ends is what I think he means here.) Where you've got a sense of time.

⌘<840541>A Yeah I know what you mean.

⌘<841027>J ...That after you've created something, or started something - cause that's what it feels to me... It's like I'm starting something. It's then reacted upon by the - I don't want to use 'machine'; I don't like the word - by the 'idea' you've created. I like the way it responds...It seems like I stimulate it and it responds. So I can actually play off it.

<effect vs. conversation> I start to get a sense of... I don't know the logic specifically but I can get... I tend to notice that in the lower registers it doesn't seem to stimulate the 'idea' very much. **<exaggerates>** It generates it's own lower register I notice that. You know I could (give sound (sometimes)... It wasn't an octave either. It seemed like sometimes it was like, I don't know, like a tritone or a...

⌘<902989>A Below what...you were playing?

⌘<904646>J Yeah, yeah. It wasn't diatonic and I quite liked that because if you've got good ears... **<surprise>** My pitch is pretty good so then it opens up ideas and I can play off your program. **<software affects music>** I like that a lot. So it created more... It stimulated me to try other areas.

Clip Keywords:

behaviour : exaggerates/sensitivity
behaviour : software affects music
feature : audio
response : effect vs. conversation
response : like/dislike/preference
response : suggestion
response : surprise

Clip Comment:

Having a resting point <- could be a design criteria?

Clip: effect vs. conversation

Collection: j2

Time: 0:15:48.8 - 0:17:49.7 (**Length:** 0:02:01.0)

Episode Transcript: aj-notes

Clip Transcript:

J When it responds in a more complex way... This seems to be a more complex response to the previous one. **<complexity>** And, particularly from the point of view of... The other one seemed not to create other harmonies. It was just like a delay. **<limitation>** So that was an effect. So it tended to keep me diatonic. **<musical strategies>** Just me, other players would do different things. When I hear this, I don't know whether it's working off the overtone series or what **<how it works>** and coming in with...you know...not the specific root kind of idea, then it feeds me into other harmonic ideas. **<software affects music>** Cause I'm thinking of it as a player.

⌘<997693>A Yep, that's interesting.

⌘<1002915>J Yeah, well I have to otherwise I don't sound very good.

⌘<1005746>A I'm just curious, did you think of the first one as a player as well or not so much.

⌘<1011220>J I thought of it more as an effect. As if I had...*stamps foot*.

⌘<1014960>A A delay pedal.

⌘<1017478>J Yeah yeah. It felt more like that. This feels like... I can imagine a person if you want.

⌘<1027562>A Ok, so you mentioned about the more complicated harmonic response. Was there anything else about it that made you think of it more as a person?

⌘<1037062>J It's not obviously a pulse. Um...The nature of a swell kind of thing. It's not linear. Maybe I'm not using the right terms. [A I know what you mean though.] But it swelled at times. [A Yes] So it gives you a feeling of conversation. Whereas the other one felt specifically like a direct response to what I just played, where this feels more like a conversation.

Clip Keywords:

behaviour : software affects music

feature : audio
response : complexity
response : effect vs. conversation
response : expression
response : how it works
response : limitation
response : musical strategies

Clip Comment:

A key idea here I think- spinner is an 'effect' to James whereas playing with the string spheres is more like a conversation - like playing with another person.

Clip: like lanterns, hanging lanterns

Collection: j2

Time: 0:17:49.7 - 0:24:32.2 (**Length:** 0:06:42.5)

Episode Transcript: aj-notes

Clip Transcript:

A It sounds like you're talking about it um... Most of the things you're saying are all aural things you know.

⌘<1077826>J Yeah...The visual?

⌘<1080281>A Yeah were you looking at it at all? [J Yeah] Was it just a distraction?

⌘<1083600>J No I don't find it a distraction. I quite like it.

⌘<1087333>A Did that help, to make it seem more person-like or creature-like?

⌘<1092195>J I'm having...I'm not certain. I was thinking about that. Cause I actually closed my eyes after a while to really concentrate on the sound. However, yeah...Alright cool. I'll just play for a little bit watching.

⌘<1111861>J *Plays. More free ranging- bigger range of dynamics(?)*

⌘<1326155>J Yeah I see really clearly. Yeah, it's great. I don't know how you did it but it's great. You can see - um, I'm trying to imagine what it is - like lanterns, hanging lanterns, and they respond to the flow. And they come back to an equilibrium. That's what I meant by the decay in a sense.

⌘<1355814>A Yes, I know what you mean.

⌘<1359843>J So then you can be... And it's quite beautiful to see. Yeah it actually does have a really clear effect. **<ease of use><link>**

⌘<1370498>A Did looking at the visuals change the way you played?

⌘<1375429>J I think so. You'll look back at the video and you'll be able to tell. Cause I'm in it so it's my perception's coloured. But I very much enjoyed it. Very much. You'll notice I'm changing harmonic areas at times. Sometimes... If we were working together on it, the previous harmonic area could still be ringing on sometimes so I can then layer some different...yeah? **<musical strategies>** [A yeah..] Cause it's suggesting to me... My ears are quite good. It's suggesting ideas. It comes out with an overtone of some sort you know, oh and it leads me into a new area.

<software affects music> So yeah I really enjoy this program. It could be maybe, sometimes maybe not consistently I'm not sure, if maybe that harmonic area would continue through even though I've gone to a new one.

<suggestion>

⌘<1446219>A Yeah, so it lingers for while so you can play off it and then fades again.

⌘<1451547>J So I can get a couple of colours at the same time. The visual thing is really effective and I think it would be really effective for the audience also for a performance. Yeah, I really liked it. Good!

Clip Keywords:

behaviour : playing
behaviour : software affects music
behaviour : sonic exploration
feature : visuals
response : distraction
response : enjoyment

response : like/dislike/preference
response : link
response : musical strategies
response : suggestion

Clip: that's the equilibrium

Collection: j2

Time: 0:24:35.8 - 0:24:52.9 (**Length:** 0:00:17.1)

Episode Transcript: aj-notes

Clip Transcript:

J An interesting idea that comes to my mind with this. You know how I talked about the sound having a decay. [A yeah] If people come into the room and that sound is already existing, then that's cool because that's the equilibrium. That's the resting point.

Clip Keywords:

response : suggestion
response : use/purpose

Clip: visual has a huge impact

Collection: j2

Time: 0:26:03.4 - 0:35:10.8 (**Length:** 0:09:07.5)

Episode Transcript: aj-notes

Clip Transcript:

J plays. *Starts with breath sounds only. Moves between two pitches.*
«<1886578>J laughing. That's funny. Yeah, this, the visual has a huge impact compared to the other one. I don't know why...because it's linear and I visualise that quite light response. This immediately feels... It looks busier. So it encouraged me to play busier, both harmonically and in time. It's interesting. I thought I noticed the scale around the circle so I tried some quarter tone stuff to see how it would respond. **<software affects music><musical strategies>** I wasn't sure I couldn't hear... Does it recognise microtones?

«<1952545>A Yeah.

«<1953657>J Because it was hard... I was watching its transition.

«<1959983>A Yeah, it does recognise it. Maybe not as accurately as your ear though, but as close as a computer can get.

«<1967667>J It's interesting I imagined... By the end I imagined a circle of people, bald heads, looking down at them and they're all nodding.

<metaphor> *Laughs.* It's quite wonderful. But it distinctly has a different effect.

«<1981480>A The sounds seemed to be - especially at the beginning the sounds were a lot subtler - that you were making I mean.

«<1990161>J Yeah, it seemed to be more... I don't know... I seemed to notice...Maybe you had it cranked over the sensitivity up. But I noticed that I didn't do very much and it responded.

«<2002970>A Yeah, I was adjusting a little bit as we went, so as you were playing softer I might have cranked it up a bit...

«<2010097>J Yeah but right at the beginning it seemed like I did very little and there was a visual response. So that's interesting. I don't know what other people's experiences are, but it kind of... The visual reinforces the audio response. It lets you know you're getting a response.

So in a sense it could enhance the trust. **<control><consistency>**

«<2038351>A Between you and it?

«<2040898>J Yeah. The response, of the second one... This seemed more diatonic, I don't know, in tone generation. I may be wrong.

«<2055045>A Yeah it is.. *(actually it's not!)*

«<2056905>J Alright cool. This encouraged more, in the beginning, sounds and things of a low level but then once the balls started bouncing around it kind of incites more full on interaction. The other one had that beautiful equilib... That idea you were waiting for it to come back to equilibrium so it encourages that kind of playing. This comes to an equilibrium again, but there's something about that linear thing. For me

anyway.

⌘<2094287>A Is it maybe a gravity thing or something?

⌘<2095014>J Yeah absolutely, absolutely and it's coming to rest. Yeah, cool. Do you think we've spent enough time with this without me knowing anything?

Clip Keywords:

behaviour : playing
behaviour : software affects music
behaviour : visual exploration
feature : audio
feature : visuals
response : consistency
response : control
response : enjoyment
response : metaphor
response : musical strategies

Clip: I've had a prejudice I imagine

Collection: j2

Time: 0:35:15.9 - 0:38:33.1 (**Length:** 0:03:17.2)

Episode Transcript: aj-notes

Clip Transcript:

J It's funny this. I haven't spent much time on this kind of activity. It's always been with people. And I've had a prejudice I imagine.

⌘<2137095>A Against computer music?

⌘<2139361>J Yeah absolutely. Very much so.

⌘<2143629>A What is it about your understanding of computer music that prejudiced you against it?

⌘<2149470>J I'm an improvising musician; that's my nature. This [*what we are doing now*] is improvised. I know it's pre set; it's a program that's pre-set. But inbuilt in it allows for the human element - at your point to create it - but then at my point to actually work it as an instrument, so that's what I really enjoy about it.

⌘<2180383>A The fact that you don't feel as constrained by it?

⌘<2182663>J No I don't feel constrained at all. You know it's like more a concert kind of idea for me where it would be a work and this [*virtual instrument*] is the environment of this piece. And then if we were to play another piece it would be a different environment that you would create.

⌘<2211007>A So you're basically creating a specific instrument for each work in a way.

⌘<2217100>J Yeah. That would be my idea...at the moment. Cause I really enjoy this. I have no idea how long the pieces were.

⌘<2228621>L Well you started 2 minutes, with the spinner one - the first one. And you went on with spheres - the string one - was 3 minutes and then you played 4 minutes when you started to focus on the visuals. And then for this one you played 5 minutes. It was the longest piece. Which is interesting, each time you were getting longer.

⌘<2255047>J I'm getting more familiar with it. Yeah. Great fun! I told you that Herbie [*Hancock*] is interested in a lot of this. So that was interesting. And Phil Slater has just done a performance in the Auckland Festival when we were down there with the Art Orchestra. That group's opened my mind up to what is music. You know. Amazing individuals with all this different experience. So my idea of what beauty and music is has opened up heaps. But this feels very interactive. The second one and the third one in particular. The first one felt more like an effect that I was working with. That's like that other (way). This is like I'm playing with someone.

Clip Keywords:

response : effect vs. conversation
response : enjoyment
response : musical strategies
response : use/purpose

Clip: very beautiful basically

Collection: j2

Time: 0:39:44.0 - 0:40:33.8 (**Length:** 0:00:49.8)

Episode Transcript: aj-notes

Clip Transcript:

J I can't really remember the visuals at all. I don't know why. I remember there was a spot when I walked in so I went, 'that's interesting so we're working with visuals as well.' But I can't remember what it did. I remember very strongly the linear, gravity one (*string spheres*). Very strongly. For some reason I really relate to that. I find it very beautiful basically. This one I like but no where near as much. It brings out a different part of my nature.

⌘<2425460>A The circular one?

⌘<2427005>J The circular one yeah.

⌘<2427818>A So of the three..

J (*I prefer the string?*)

Clip Keywords:

feature : visuals

response : enjoyment

response : like/dislike/preference

Clip: responds far more to the upper end

Collection: j2

Time: 0:40:33.8 - 0:41:24.0 (**Length:** 0:00:50.1)

Episode Transcript: aj-notes

Clip Transcript:

A Could I maybe put the first one, the spinning one.

⌘<2437834>J With the more complex...yeah. Without telling me any more about it. And I'll concentrate more this time on the visual also. Cause I'm starting to understand the idea. I am noticing the cutoff point, not the specific cutoff point in the range. It responds far more to the upper end to bottom end. **<limitation>** It's not a criticism I'm just saying I'm noticing the instrument.

⌘<2461583>A Yeah, that's the kind of stuff we want to hear. Around what note?

⌘<2467124>J Oh, I just say at mid range it responds, at high range it really responds. It seems to be really responsive to the higher **<exaggerates>** but when I go down to pedals and stuff like that it doesn't seem to respond hardly at all.

Clip Keywords:

behaviour : exaggerates/sensitivity

feature : audio

response : how it works

response : limitation

Clip: what it responded to and what it didn't

Collection: j2

Time: 0:42:20.4 - 0:43:53.2 (**Length:** 0:01:32.8)

Episode Transcript: aj-notes

Clip Transcript:

J *Plays, breath sounds only. V. soft attack. A interrupts.*

⌘<2592947>A Sorry James. I realise I didn't actually change it as I thought.

⌘<2602700>J Yeah.. It was interesting as to what it responded to and what it didn't. **<consistency>** And when it forgot the earlier thing and sequenced on to the next bit. **<bug?>** *Laughs.*

Clip Keywords:

behaviour : playing

feature : bug

response : consistency

Clip: The visuals come to rest, but the audio wasn't done

Collection: j2

Time: 0:43:53.2 - 0:48:49.7 (**Length:** 0:04:56.5)

Episode Transcript: aj-notes

Clip Transcript:

J *Plays again- breath sounds. Plays note, crescendos, leans into mic.*
⌘<2911885>J The visuals come to rest, but the audio wasn't done. So there's less of a relationship. Doesn't help the trust kind of element.

<control?>

Clip Keywords:

behaviour : playing
behaviour : sonic exploration
behaviour : visual exploration
feature : audio
feature : mapping
feature : visuals
response : control
response : link

Clip: idea of a decay

Collection: j2

Time: 0:48:49.7 - 0:52:16.7 (**Length:** 0:03:27.0)

Episode Transcript: aj-notes

Clip Transcript:

A You know, just let me have a quick look. It's conceivable that I can make it so it does come to rest without too much [*trouble*].
Lots of clicking, etc...A makes changes, trys a few things. It works but cuts off very suddenly.

⌘<3088650>A Well it does it but it's too sudden. You'd want it to gradually fade out rather than...

⌘<3093144>J Let's try that. Unless you've got something in mind that would be easy for you to do now?

⌘<3098906>A Um... I think that's... It might take ages. It's crude but maybe we can try it.

⌘<3113750>J Because it's quite substantial that idea of a decay - in the natural world.

Clip Keywords:

behaviour : request
feature : audio
response : suggestion

Clip: the dominance of it

Collection: j2

Time: 0:52:33.4 - 1:02:51.9 (**Length:** 0:10:18.5)

Episode Transcript: aj-notes

Clip Transcript:

J *Plays with modified spinner. Breath sounds, attack only. Short bursts of sound. Audio playback cutoff is very sudden. (Not the gradual fadeout I think he wanted.)*

⌘<3297851>J I think I know what it is. It's the dominance of it **<power balance>** (*points at screen*).

⌘<3309186>A In volume terms you mean?

⌘<3310730>J No. Um, the other ones seemed very interactive in a way I could understand. **<natural>** I can hear that this goes - in simplest terms - you hit something and it goes (*sings software response*) and then slowly...right. Alright, here you go. The pulsing thing, I don't like. It feels inhuman. **<(un)natural>** It immediately feels like something from one of the sci-fi films that I grew up with. **<metaphor>** Immediately. So maybe that's association, ok. I can't know that. Cool. The predictability of that contiunally rising thing is a huge thematic idea. Very dominating. It's very dominating that program. It's the main player in a sense. **<power balance>**

⌘<3371868>A Ok, I see what you're saying. It's too constraining, setting

too many limits [J Parameters] on you somehow. <flexibility>
¤<3378432>J Yeah absolutely. Its response is... It's like a machine. It feels like a machine. <metaphor><(un) natural>
¤<3383303>A Yeah, whereas the other ones felt like creatures.
<natural><effect vs. conversation>
¤<3385233>J Absolutely.
¤<3386459>L So it's driving you?
¤<3388510>J It's ignoring me. It's responding to me but in a very limited fashion. <control> It doesn't under... I'm talking to it and it doesn't understand. It's just coming back with the same old shit. Yeah. So when I said the word trust. It's interesting. Alright coming back to that other thing- my prejudice before (*against computer music*). That displays it.
¤<3424567>A It's like what you imagined.
¤<3425851>J It's a machine with very limited understanding and very limited range of response. <limitation> The other one - two, the other two - but particularly the middle one. I don't know what it is about it, seems to actually be more conversational. <effect vs. conversation> So instead of being this predictable. When I say predictable - I couldn't change the effect of the first one (*Spinner*). I change intensity, I change range, I change... It didn't seem to recognise just sounds without pitch. Alright, it didn't appear to. No matter what I did, its response was basically the same. <complexity> I know it wasn't but in so many ways it was. And the other ones may have been but they didn't appear to be.
¤<3486013>A Ok, I understand I think. With the spinning one [J That one. It feels alien.] Is there anything about it. Maybe could you think...I don't know. There's the visual look of it, spinning around. If you could imagine something about it. If you could imagine transforming it into something you would like to play with.
¤<3511044>J Alright, cool. If the audio wasn't this pulsing thing. If you've got the pulsing thing in combination with that spinning thing like that. I think the pulsing thing is the most dominating, alienating thing for me; the pulse. If you got rid of - and I'm not saying to do that - I'm just trying to delineate between the effect of that pulsing audio compared with the visual. If it was a different audio the visual may take on a different. The fact that it had lines - I'm just trying to think what it stimulates the brain. The lines... It feels alien! It feels mechanical! <(un) natural> The whole feel of it - audio and visual feels mechanical and inhuman. And unnatural actually, bringing us back to the visual of the hanging line with a sense of gravity, with a sense of direct response to weight you know. So it brings down to the weight business is very important to me.
¤<3588498>A You talking about the effect of the weight of the sound?
¤<3591369>J Yeah, my dynamics. It seemed to be very direct and related to the natural world. Yeah! Very much so. That one more than anything visually related to the natural world to me. It was a natural construct. <natural> You can see it. In fact I could make it physically and do it. You know, I'd have to find balls with the right density and size, but I could actually make it do that. (*ie. physically build something similar.*) You know trombone's interesting because it actually does put out sound and air and all of that kind of stuff and I could make the balls move in that way. I could play with them in that way. Whereas the other one, is so mechanical.
¤<3637336>L So are you saying that, in the spinner one, that the visual reinforces the [J Audio?] audio? <mapping>
¤<3654941>J Yeah I would say that. The audio was the most dominant alienating effect. <mapping>
¤<3658869>L Right, but if you took away the visual and just had...
¤<3662028>J I'd still find it bloody annoying.
¤<3664481>L Is that because it is (?) sending and has no variation, no kind of change of pace?
¤<3671862>J Yeah, and the pulse. The pulse just going (*sings pulse*). [L And that doesn't change the pace?] Well actually it has clicking points

where it changes so that created a little bit of interest, but it was random and I could have no effect on it.
⌘<3688755>L But it doesn't change pace does it? I mean if that pulsing changed - paused and changed pace...?
⌘<3693480>J If it changed pace in direct response in some way to my stimulus then I'd feel like it was listening to me. <consistency><control> I don't know what is in the programs right, but I had the sensation of... It's like talking to someone who's not listening to you. They've just got what they've got to say and it's one subject. It's like talking to a politician - you can ask him any question and he'll actually answer what he wants to answer. You know it seems to be one mind that..
⌘<3732921>A Well politicians always agree with you actually.
⌘<3734779>J They agree with you in order to be get what they want to say across. It's all - sorry we're getting off the thing. This feels - it's got what it wants to say and that's basically it. <effect vs. conversation> And its response - it may be complex but it feels really simple. The other ones, give the...supported by the visual - particularly the vertical one give you a distinct... They make me feel comfortable immediately.

Clip Keywords:

behaviour : playing
feature : audio
feature : mapping
feature : visuals
response : complexity
response : consistency
response : control
response : effect vs. conversation
response : flexibility
response : like/dislike/preference
response : limitation
response : metaphor
response : natural
response : power balance
response : suggestion

Clip: I control with my stimulus the speed of

Collection: j2

Time: 1:02:51.9 - 1:05:28.7 (**Length:** 0:02:36.8)

Episode Transcript: aj-notes

Clip Transcript:

A I wonder - with the spinner one, this one - if it was somehow so the tempo was completely variable and controllable by you. So it was still something spinning but it started off very slowly.
⌘<3789162>J If the audio was related to it. If the audio was related to it so if... if I could create with dynamics... For instance if I hit it hard, if it sped up (*sings digadigadiga*) and with the decrease in dynamics (*sings slower; dong, dong...*).
⌘<3811406>A So you could spin it and it would gradually return to a slower pace?
⌘<3813471>J No, no no. I control with my stimulus the speed of... right. We're talking about possibly... It's maybe one line - I don't know the logic right - but it's maybe one line could become multiple lines so I can layer. It can remember for a period of time the previous stimulus. But then that with a decay. Possibly the decay of a louder thing being longer. The natural kind of thing. So then I could hit something and it (*sings high, fast dagagagagagagaga*) and then put in another stimulus (*sings low, slow daw*) and it goes (*sings mmmmm*) while the other one's still happening so I can create layers.
⌘<3861657>A Ok, so you could play a D and a D thingo would spin really fast...?
⌘<3866683>J Yeah, if I hit it hard there is a hard response with a quick...right? And then I can play another thing while it's going right and then create another layer with a totally different density, speed of...

⌘<3882011>A So you're kind of talking about multiple spinning objects rather than...

⌘<3884801>J Yeah, you're still talking visual, I'm talking audio and then how would we create that visually where maybe in the center you could have the spinning thing and then maybe concentric possibly, you could have a slower moving thing. But my main problem is the audio. I'm still...I'm a musician so my emphasis is the audio but I do find the visual of the hanging thing totally related. It felt related absolutely to the soundscape that I liked.

Clip Keywords:

feature : audio
feature : mapping
feature : visuals
response : musical strategies
response : natural
response : suggestion

Clip: isn't bonding

Collection: j2

Time: 1:05:28.7 - 1:07:26.0 (**Length:** 0:01:57.3)

Episode Transcript: aj-notes

Clip Transcript:

A That's interesting because - to reveal a bit more of the technology - as I said they're all based on physical models and the reason - you said before I'm talking visually when I'm talking about how this thing works - is I'm trying to think of it as an object as a physical object. So what the software is doing is taking your sound and mapping your sound into forces which act on this physical object. As the physical object moves, it also makes sounds and those sounds could be like the string one which is sort of atmospheric, sine wavey kind of sound or it could be like the spinning one where it is your sampled sound.

⌘<3968587>J However, I know it's my sampled sound because I heard at one point the timbre and I heard when you went over and spoke into it very clearly the timbre so that gives a relationship right? However I never sensed it because of the pulsing thing. It wasn't a strong thing. So something that should be bonding which is hearing your own timbre back - you would imagine that would be bonding - isn't bonding. So, it feels inhuman and alienating whereas the other one felt interactive, encompassing and encouraged, you know. I was sitting there with that one then going, 'fuck, that didn't work. I'll try this! Oh Jesus it's still doing that kind of thing.' You know what I mean.

⌘<4021976>A Just not enough variety...

⌘<4023133>J Um...It doesn't feel intuitive. The other one felt intuitive. I don't know if these are the correct words.

⌘<4031437>A I know what you mean, so they are the right words.

⌘<4033451>J Yeah, it felt mechanical. I know it's much more complex, but the most dominating feature is that delay effect. The pulsing effect.

Clip Keywords:

feature : audio
response : complexity
response : natural

Clip: I'd like to be able to create layers

Collection: j2

Time: 1:07:26.0 - 1:12:59.3 (**Length:** 0:05:33.4)

Episode Transcript: aj-notes

Clip Transcript:

A Just brainstorming a little bit - one of the reasons we did that spinner one kind of with the style it had is because we wanted a complete contrast with the string one.

⌘<4053300>J Well it certainly is. Now if we're talking movement...

⌘<4055487>A Not so much good and bad.

⌘<4056560>J No it's a total contrast absolutely. And from that point of view it could be useful. But if was to be working with you I'd like to be able to create layers and then I could work with that. And if you could put that thing on each of the layers that with my intensity the speed of the pulse changed so I could have numerous speeds of pulses going. Then I'd really enjoy it. But I imagine that's quite complex.

⌘<4102090>A I think it's doable. [J Of course it's doable. Anything is doable!] One thing is with all this stuff. The devil is in the detail so you can come up with this idea that in your mind is going to be great and then you play with it... There's a lot of trial and error. And it's amazing what difference what I think is a tiny little change can have a huge impact artistically.

⌘<4134412>J Absolutely. This layers thing - it doesn't appear. Like the pulse seems to be set through the whole soundscape is my point. If you can layer it so there's numerous layers of pulses within the soundscape. [A That you have control over.] Yeah, absolutely. Even if it was three; even if it was two it'd be far more enjoyable.

⌘<4155429>A And what would you imagine... Like from the playing point of view how would you switch your attention from one layer to another. I mean like you could have a pedal or something where you say...

⌘<4165380>J No no no no. I'd like it all to be there at the same time.

⌘<4169753>A With sound. [J Yeah.] But I'm thinking so if you've got a layer here, a layer there and another layer there, what... Is this a mid range, this is low range, this is high range or something?

⌘<4179315>J More to do with I would say dynamic.

⌘<4183359>A Ok, so louder dynamics would make this layer...

⌘<4188992>J Ok, talking logic. Let me think logically for a moment - seriously. Alright I put in a new input - yeah because I'm not playing constantly. Oh that's what it would do. I would use it more as an instrument and it'd be more interactive because I would play an idea, say an idea- strong dynamic, bam - which would signal a train of response possibly with quick, quicker. Right? Then I would play another idea with a different dynamic which would stimulate now a parallel response. Not related to the first one, so it'd multi...right? And if we're thinking polyphony it would be nice that I could get maybe four layers or something happening. If I was aware that there was only four layers I would allow the first one to decay out before getting to the fourth one, so we wouldn't get (*claps hands*) cutoff.

⌘<4252232>A Yeah I think I know what you're saying. And the musical ideas you're talking about, these little phrases or whatever, are they like a certain length or that would be just whatever you felt like.

⌘<4263579>J I would...If I was happy with the response right, then I could play quite short ideas. And then I can use it more as an instrument, in composition.

⌘<4287759>A So we need to make the whole thing richer somehow, more controllable, richer.

⌘<4291588>J Yeah, so maybe having four of them (*spinners*). You know what I mean, four things in parallel, but not necessarily working. One triggers with your first idea. If it's engaged then the second one's waiting for an idea and then responds.

⌘<4311956>A I wonder... I can't help it- I'm thinking technically again - but you could have say four mics or something.

⌘<4317116>J Oh shit yeah. That would be the easiest way. Or as you said a pedal.

⌘<4323979>A Yeah. I'm like you, I like the idea of not having switches and things - just sound.

⌘<4331152>J From a performance point of view, the visual for the audience is a big part of it. So it'd be great to actually have four microphones across a stage, so they're actually seeing you walk so they're actually and then they start to understand the logic. The audience starts to understand the logic now they're expecting something. Then if for a period of time we have each mic is the same logic for want of a better word and then through the piece you change the logic so you're actually surprising the audience

again. That kind of idea might be fun to play with.

Clip Keywords:

response : control
response : musical strategies
response : suggestion

Clip: weather vane?

Collection: j2

Time: 1:12:59.3 - 1:14:53.1 (**Length:** 0:01:53.8)

Episode Transcript: aj-notes

Clip Transcript:

A Would you imagine, if you were performing with it, that you would be looking at the screen and responding or do you imagine that the screen's behind you so you're just in the music?

⌘<4388436>J Yeah I was thinking about that. With the ball one, with the line (*string*) - I want to see it. It's really supportive. <motivation?> I'd have to see the new visuals of the first program to know if I want to see it. I would tend not to want to see it with the first one. [A Certainly as it is now anyway.] It didn't enhance it. Sonically, I'm - see how you get the layers there. I haven't got my glasses on - is there five layers there? Those stacks. That kind of idea...

⌘<4435174>A It's really interesting you say that - it's funny... When we first started working on this one Ben was expressing a similar kind of idea. I think he was asking for three of them on the screen and he was saying that if you play like in the high range then the top one would respond... He was certainly talking layers.

⌘<4455827>J Yep, mine's more the intensity. So with my dynamic range if you want to continue the pulse concept, for the pulse to be central to the compositional theme, to have a direct control of the speed of the pulse. Instead of using range I'd like that to be with dynamic, because that's quite a natural thing.

⌘<4479970>A You could push the thing to start it going.

⌘<4482960>J Absolutely. If you think about one of those, what are they called, weather vane? You know, so you blow hard - zing! - you know what I mean.

Clip Keywords:

feature : mapping
feature : visuals
response : control
response : metaphor
response : motivation
response : suggestion

Clip: excite a situation or calm a situation musically

Collection: j2

Time: 1:14:53.1 - 1:17:51.7 (**Length:** 0:02:58.6)

Episode Transcript: aj-notes

Clip Transcript:

A Another dream I kind of had of one day doing is somehow using so you could have these things - like a weather vane - on the screen and you could somehow load it up with music by playing but then you could reach out with your hand and (*swipe it around*).

⌘<4508689>J Yeah great. The beautiful idea for an interactive kind of thing. For people off the street that's brilliant. If it's a performance - concentrating on what my area of skills are and what my development... Cause I know how to excite a situation or calm a situation musically.

⌘<4542336>A I see what you mean. So this (*suggestion for gestural playing*) would be a kind of crude gimmick.

⌘<4548164>J For me yeah. Let's get away from criticism of that idea and get back to the other one where if there's a more natural relationship - what I would call a more natural relationship towards the kind of thing I'm doing and the response of the program. If you want a feeling of domination

and alienation, that's certainly there with that one. I'm not being sarcastic. If you want the feeling that the machine actually is the dominant thing then that creates it quite strongly. All of that. It's very strong, the feeling of alienation makes me uneasy. And if it's in a different section of a long piece then it certainly creates tension. And a feeling of alienation. I did 'Kraanerg' - you know Iannis Xenakis. You know?

⌘<4610320>A I don't know the work but I know Xenakis.

⌘<4610963>J Cool. Man, that piece. I got asked to do it by Roger Woodward that guy. I didn't have any time so at the last minute I was checking out the tape of the one recording that there was and I nearly smashed the car! Seriously. I tried putting it on at home - I had young kids at the time - and the house nearly exploded. Like I understand there's a purpose - you know we don't always want to be expressive. But this very strongly is more in that vein. It's inhuman. Inhuman is the wrong word but you know what I mean? I kind of is.

⌘<4656716>L Unrelenting?

⌘<4658642>J Yeah, absolutely. So if you want to create alienation - cause that's what it does - whereas the other one drew me in. It felt like a being I was playing with.

Clip Keywords:

feature : mapping
response : effect vs. conversation
response : power balance

Clip: If you had four of them running

Collection: j2

Time: 1:17:51.7 - 1:20:30.1 (**Length:** 0:02:38.4)

Episode Transcript: aj-notes

Clip Transcript:

L Could you imagine something like the string one in a situation where you were playing with other musicians - human musicians - and that was one of the instruments?

⌘<4683836>J Yes I could and I'm immediately thinking, 'Do all the musicians affect it?' That's a question, not a statement.

⌘<4698995>A It'd certainly be a lot more complicated and might not be as easy to see...

⌘<4703032>J ...the relationship absolutely.

⌘<4706461>A It depends on the kind of music you're playing.

⌘<4707365>L A number of different people playing to it. Like when a jazz band plays sometimes you get people to step up.

⌘<4714302>J Yeah- oh... So we've got one mic which is the stimulus in this case. One person plays with the logic and the other ones are playing, and then they move up. Yeah that's one way. Another way is, only one person plays with the logic and the other ones play off it. Yeah easily, no problem with that. It's more complex... But if you had four mics all going to the same program it would break down the relationship between each individual the program. In my opinion.

⌘<4752572>A Yeah I think so too.

⌘<4753774>J Cool. If you had four of them running each with a stimulus from each player then that's a different thing again.

⌘<4763141>A I suppose you could have soprano, alto, tenor, bass or something like that. Different size... Cause if they were all the same I imagine that might be a bit boring, but if they were somehow similar but different that might be more interesting.

⌘<4770650>J It would be interesting. The other idea is, you could have a central mic that was right open and the players circle around it and then it brings in...They know they have to create space and layout to allow...And they have to make a choice to how dominant the logic is compared to them as individuals. It's an interesting problem, because you're considering...it definitely brings up the thing that the logic is an entity. It's not an effect it's actually an entity. I'm surprised I'm even thinking like this, but I found it so enjoyable that I'm going yeah it

felt like an entity. I had a duet happening.

Clip Keywords:

feature : audio
feature : visuals
response : effect vs. conversation
response : enjoyment
response : musical strategies
response : use/purpose

Clip: they're playing each other

Collection: j2

Time: 1:20:30.1 - 1:26:58.8 (Length: 0:06:28.7)

Episode Transcript: aj-notes

Clip Transcript:

L There's something about the particular response of the spinner one, which is very negative because it's not listening to you. In the case of the other one - the string effect - there's an issue about predictability - what you can control and what you can't (yet?). Did you find that each time you did something that you were getting the response...?

⌘<4858780>J There was a good balance. Because, in discussion, I don't know what you're going to say next.

⌘<4865952>L Well I just wondering what was unexpected and what was like expected?

⌘<4869450>J I liked that harmonically it was unexpected at times. It created new ideas. My ears are really good - my pitch is not perfect but it's pretty good. So, it was generating ideas just as other musicians suggest harmonic ideas to me and I go 'ah...' and I flow off it. So that's what I loved about it. It was unpredictable in what it generated harmonically at times. But not so unpredictable that it was like, 'What?!' You know, it was unpredictable...

⌘<4905091>L So you were getting unexpected effects which you enjoyed...

⌘<4908979>J Yeah absolutely. It felt like playing with someone. It felt like playing free with a keyboard player or a guitarist.

⌘<4915483>L Can I ask that question because it's quite interesting that you're looking for accuracy a lot of the time and sometimes that things that you thought of as being not quite right have from the musicians point of view been quite interesting. I mean that's a non-computer thing to do, to be unpredictable. What I'm saying is that we always think of the logic, the predictability but we also are quite engaged by the idea that it'll do something not quite...It's consistent but not unsurprising.

⌘<4956428>A There's a balance there between control and...

⌘<4962092>J It's what we call 'reasonable' in the community. We don't want everyone to be the same. But, we, you know...say people need medical treatment when their response is not within socially recognised boundaries. So, there's unpredictability but it's within the...

⌘<4992103>A Like a good conversation in other words.

⌘<4995265>J Absolutely. That nails it. It felt like a conversation. I was having a musical conversation with an entity, in the second one. I preferred it to the third one (*circle spheres*) actually. I don't know why. I enjoyed the third one. It'd be interesting to see it.

⌘<5012142>A I could ask you... I might be because the second one has sort of responded to your harmonics more. Well, it responded to harmonics whereas the third one actually didn't. It was purely just the fundamental of what you were playing.

⌘<5029027>J Yeah so it was less complex so it didn't suggest as much.

⌘<5032296>A Yeah exactly, the sounds that were coming back were less complex.

⌘<5035549>J Yeah, so I really enjoyed the second one. That's what felt like conversational.

⌘<5039127>A That's the interesting thing for us is to get this balance, like we were saying, between control [J Absolutely.]. You don't want absolute control because then it would just be like talking to a brick wall.

«<5048748>J Absolutely, it's no different from humanity. It's not different from...It's like made by human beings for human beings so it has the same goals that a human being has, hopefully. Unless, you really want to alienate. If you want to alienate then don't follow any of the ideas that human beings aspire to in their best sense. So for instance that kind of thing (*spinner*) brings up alienation which then stimulates the feelings of lack of confidence, you know, vulnerability, anger.

«<5092261>A You said that would stimulate that feeling in you as a musician or in the audience or...?

«<5096386>J Both! Both, one hopes. Maybe one hopes. One imagines, that the feelings are shared.

«<5109181>A I guess I was thinking...I could imagine if I was playing a piece I maybe playing a piece which is strident as intended to be uncomfortable for people to listen to in a way but I'm not feeling that as a musician as such. You know what I mean maybe?

«<5125860>J Possibly. But when I play. Sorry - it sounds pompous... When I play, I'm in the audience.

«<5137159>L But do you try and make them comfortable?

«<5138416>J No. But I'm in the audience watching myself play. Alright? I'm trying not to be here. I'm out there watching a guy playing or listening to a guy playing.

«<5153942>A So in other words are you saying you want to feel the music as you're playing it? It's not...

«<5160562>J Yeah...It's the sound. I play the sound. I'm not the trombone player, I'm out there I'm hearing the sound and I have an effect on the sound. It just happens to be a guy with a trombone over there. And there's that feeling of being able to get inside - well actually get outside of yourself to get inside of it. Certain things you know, trust - I said the word trust, that's it. You can have that alienating feeling with human beings you know. Vast, huge amounts of bands have that. They're individuals playing you know. They're all quite good but they're still individual. When you're move, when you're actually moved, when your heart is moved is when you know there's absolute...They're playing for each other and they're playing each other. You know, there's real communication.

Clip Keywords:

feature : audio
response : complexity
response : consistency
response : control
response : effect vs. conversation
response : surprise

Clip: cool conversation!

Collection: j2

Time: 1:26:58.8 - 1:33:19.6 (**Length:** 0:06:20.8)

Episode Transcript: aj-notes

Clip Transcript:

L The audience has to be already with you in a sense, don't they? I mean, if it's a... If any new music has come out of - (hundreds of years of it?) - and the audience has not been ready for that music...And they just don't respond.

«<5233007>J I have no argument about that, no argument about that.

«<5235920>A I wonder about that...

«<5239191>J But -possibly what you're saying... This concept I'm talking about - I have faith that it actually overcomes all boundaries. For me when you've got people on stage and there's that kind of intimacy and trust, alright? Unless the people are pissed and aren't there for the purpose of... If they're there for the purpose of actually engaging with the music right, then when there's intimacy they are drawn in despite themselves. No matter what (note?)

«<5277563>L I still think they've got to be ready.

«<5279116>J Well, the ready is they've come with the desire to interact. I have no argument with that. And there's various ways. Like, you know,

we're on a huge rave here. But you know, down there the opera house and everything - they get the funding - no they get the people turning up, bums on seats, because of the predictability. One would argue they can't put too much new material in because then there's not predictability, because often when people go there they're tired, they want...yeah, you know what I mean. This is where education comes in and all the rest of it where people are excited about things that are new. I have no argument with you, I totally agree with the idea, but this concept of intimacy with the musicians and a desire to be intimate with the audience. For me this is what I've worked on my whole life is building this oneness, alright? And then once you've got that trust, you can take them further than you normally would have been able to. If you alienate them you can't take them anywhere.

α<5349256>A You can still... The new music you're talking about, Stravinsky and stuff like this...

α<5352485>L It's still not acceptable. I've been to the opera house when they've done Shoenberg and Stravinsky and they still don't love it - it's quite obvious. Some of the musicians will love it.

α<5364315>A I'm not sure about that, because there has always been an audience for that stuff. Maybe it's...

α<5368969>L Yeah but they're the kind of people who know music.

α<5371443>J There's an education, there's aesthetics.

α<5374701>A It's exposure...

α<5376796>L Yeah, it's exposure and learning to listen.

α<5379324>A/J Exactly.

α<5381751>L I mean I took some jazz friends, who love jazz - old jazz, new jazz, etc. We went to the basement to see the Necks and they were just - they couldn't...they'd never heard jazz like this before. They didn't think it was jazz.

α<5394577>J Yeah. They're my friends - all the people in the Necks are my friends. But I go to it sometimes. [L I thought it was fantastic.] It is! Because it is! But you've got to be ready - you said you've got to be ready and you're absolutely right. You've got to be ready. You've got to be in the right frame of mind. It's like even at home, I can't listen to all my music all of the time. I have to be very specific what I listen to cause it's what I need. You know it takes a lot of courage to be an audience to go to... Consequently, there's groups that I'm in that don't play as much as other groups. ooh.. the very end of the expression is: it's McDonalds. It's the most popular thing because it's always the same. It may be shit, but it's always the same shit alright. And it's sweet, it's fatty, it's got the fundamental things. And so you get people, without education, in food, without a desire to know any more, that they'll go there. There's a lot of people... You know there's that classic thing with the Americans going overseas. They want a find this, the McDonalds in each town because they know they won't get food poisoning and they'll know what the food is. And then there's other people who go overseas so they can experience (other food). But that's courage. So when I used the word trust before - the trust thing - and you can build the trust and it builds peoples courage and then their limits open up and then they can actually hear some... The Art Orchestra doing that more successfully. We're doing these collaborations with Ruby Hunter and Archie Roach right. It's broadened the base of listeners you know. Because there's more songs for a start. They're amazing human beings that sing amazingly. Ruby's voice is challenging at times, you know it's not the standard kind of voice. And then we play... Some people think we're a support band for Archie, and we're not down and doing what we should be doing which is just always experimenting but maybe there's...it's always balance. That's what I found with the second one. There's balance, there's rules - not rules but there's understandings and relationships. To build trust and confidence, you know, if you alienate... I don't know what Shoenberg and Stravinsky's - I don't know what their philosophy was. I haven't studied enough to know, what... But I understand that period of music was a response to the wars.

⌘<5568376>L Well Shoenberg just changed the music (for ever?).
⌘<5570896>A Yeah, I'm not a guru on Shoenberg either...quite intellectual.
⌘<5578340>J But, that's an intellectual response - that's a desire to smash and start again or ignore and start again. If you're going to get people alienated because of that you shouldn't be surprised! You've just changed the rules! You know what I mean - I'm sick of people whinging about this...

Clip: you have to have one element that people can hold onto still

Collection: j2

Time: 1:33:19.6 - 1:35:05.4 (**Length:** 0:01:45.8)

Episode Transcript: aj-notes

Clip Transcript:

L What you're saying about the string is that it's kind of bridges that boundary between [J It helps.] having a framework so that people can understand and having something you can extend. And you've got a possibility of going into...

⌘<5616502>J Yeah absolutely. You've got to build the trust first. Or, people know they're coming, knowing they're not gonna know what they're gonna get. And knowing they're going to be confronted. If they have a desire to be confronted then it's cool. But that's the mission thing or desire...

⌘<5636921>A It's interesting with a lot of that stuff like the Shoenberg. They threw out one side of it - the tonal aspect of it - but it's what they keep you know. The overall shape and the big form, things like that.

⌘<5652158>J In free music - free improvisation, I've always argued that you have to have one element that people can hold onto still. And then there are...then they can deal with everything else. As long as they've got something to hold on to.

⌘<5667891>A Those fundamental things, like contrast and opposition...

⌘<5672386>J The Necks! The Necks build a trance-like thing, you know, and it's a slow build. That's the general theme of how they work. It's a very overly-simplistic way but they...there is a common thing always through there.

⌘<5690432>L Have you ever listened to people like Terry Riley? [J No.] I think the Necks are the next Riley and Philip Glass (the link between their music?) seemed very strong to me.

Clip Keywords:

response : consistency

response : control

response : surprise

Clip: trying to find the balance between structure and creativity

Collection: j2

Time: 1:36:21.3 - 1:37:25.7 (**Length:** 0:01:04.4)

Episode Transcript: aj-notes

Clip Transcript:

J Not only that- I'll just argue this. Creativity and structure. We're constantly as an improvising musician trying to find the balance between structure and creativity. Which funnily enough from my experience relates absolutely to the thing of being a human being where you have to find a balance between structure and creativity. If it's just total creativity it's anarchy and there's...

⌘<5807661>A Well creativity has no meaning does it in that sense...

⌘<5810062>J There's nothing.

⌘<5811015>L I've just written a paper on the subject in Leonardo Transactions, only I put 'constraints', 'creativity and constraints'. [J Right.] Constraints not in the negative sense of the word but just (provide) the structure.

⌘<5824248>J No I understand. As soon as you say constraint it brings up in

me the immediate perception of negativity.

⌘<5831606>L Yeah, that's right, it has that. It doesn't actually mean that.

⌘<5834607>J No. But in common practice it does. I'm fascinated by that too. The original meaning and then the practice of it...

Clip Keywords:

response : complexity

Clip: HAH!

Collection: j2

Time: 1:37:46.5 - 1:38:00.4 (Length: 0:00:13.9)

Episode Transcript: aj-notes

Clip Transcript:

J (*Reading from questionnaire*) The instrument responded to music in a way I would describe as 'natural' - HAH!

⌘<5873235>A There's no 'Hah!' on the form!

⌘<5875553>J No? That'd be just off the page here to the left.

Clip Keywords:

response : natural

Clip: actually not simple

Collection: j2

Time: 1:38:00.4 - 1:39:12.7 (Length: 0:01:12.2)

Episode Transcript: aj-notes

Clip Transcript:

J 'The instruments behaviour was consistent. It responded to identical notes in the same way.' Ah- pardon me. 'The instrument's behaviour was consistent.' This is actually a simple question but it's actually not simple you know what I mean?

⌘<5892391>L Did you get what you thought you were going to get? Is that what it means?

⌘<5895039>A Well it's more like, if you played two B flats in the same way...

⌘<5899544>J I didn't notice that. I didn't go for that. It's not something that I checked out really. I'm not meaning to be critical. Well I am but not in a negative sense.

⌘<5912274>A Well, did you find it somewhat predictable in the way that it responded so that you felt you could... Maybe it's like the trust but a different kind of aspect of it. If you did something it would respond...

⌘<5931230>J For me it was consistent but not in a pleasurable way.

⌘<5936074>A/L That's ok.

⌘<5943772>L So the stimulus response - you got the response you expected.

Clip Keywords:

response : consistency

response : enjoyment

response : like/dislike/preference

Clip: machine would have to be more dominant

Collection: j2

Time: 1:39:12.7 - 1:41:50.2 (Length: 0:02:37.6)

Episode Transcript: aj-notes

Clip Transcript:

The people you're getting in - are they coming from a wide background? Are you getting non-improvisers in?

⌘<6011005>A Predominantly so far, because of my background...you're the first improviser we've had. But we've only had four people so far.

⌘<6022049>J Have my responses been very different from them?

⌘<6028131>L Your playing has. It's been fantastically different - really good.

⌘<6030560>A Yeah, you got some amazing sounds out.

⌘<6034730>L You got things that I have never seen out of this thing.

⌘<6039805>J Well, you only get out of it what you put in to it. That's the

point, if you've got non-improvisers the machine would have to be more dominant one could argue.

⌘<6053769>A Yeah, they need more structure.

⌘<6058620>J If you've got an improviser... I actually listen to everything I play and I see this... When I'm playing with other people I don't just listen to what I'm playing. Like I said I'm out the front. So this kind of program or instrument is a very different relationship if you've got a really developed improviser. You know what I'm saying. I'll keep exploring possibilities. You know if I played with this over a period of time I'd have then suggestions that we could probably... I could go, "oh I hit a wall here. I couldn't get any more."

Clip Keywords:

response : power balance

Clip: perfect balance between consistency and inconsistency

Collection: j2

Time: 1:42:15.7 - 1:43:58.8 (Length: 0:01:43.1)

Episode Transcript: aj-notes

Clip Transcript:

J You've got me confused now. You've got me excited and confused.

<motivation>

⌘<6149466>J Now this is... See this is why I find this question confusing. Cause there's no sense of satisfaction or dissatisfaction. So I find that even though it's a simple question it actually doesn't satisfy me at all. It doesn't give me a chance to say what I wanna say. It was pleasantly inconsistent or its balance of consistence and inconsistency.

<consistency><surprise> You know what I'm saying- it's too limiting.

⌘<6180950>A Yeah I know what your saying. I mean, that's why we've got the video camera here... This (*questionnaire*) is very much a blunt instrument. [J It's a very blunt instrument.] So if you feel or disagree that it was consistent but you liked that...

⌘<6200537>J I don't... It was the perfect balance between consistency and inconsistency.

⌘<6204982>A Yeah I see what your saying, so to say 'neutral' (*on questionnaire*) is probably the wrong thing isn't it.

⌘<6207849>J Totally wrong because - 'neutral' - how bloody positive is that? It's not positive at all. Anyway, moving on, I'll get this other one...'The instrument was conceptually simple'. No, it was complex- it was beautiful.

Clip Keywords:

response : complexity

response : like/dislike/preference

response : motivation

response : surprise

Clip: I was prejudiced against it and now I'm not

Collection: j2

Time: 1:43:58.8 - 1:46:06.2 (Length: 0:02:07.4)

Episode Transcript: aj-notes

Clip Transcript:

A If you want, you could write... If you feel there is a better word...

⌘<6240505>J A new statement? Right.

⌘<6242269>A Because we're only getting a small number of people in anyway so I don't think we're going to make a statistical analysis of these kinds of responses. I'm really more interested in learning about what we're doing.

⌘<6270498>J If I can suggest... I'm already suggesting. I'm being far to dominant and I apologise. However, what's important - what I would find important. You know I said I was prejudiced against it and now I'm not. Cool. If it's a machine, I don't wanna hear it. If it's a human being's creation which...creates the opportunity. **<expression>**

⌘<6315517>A Like being an instrument builder.

⌘<6316694>J Yes. Then I totally support it. If it's just a bloody machine spurning out stuff, I don't wanna hear it. It's the dehumanisation of the society. I hate it. It's an example of that. I don't wanna hear a random note generator. If something can enhance human beings' creativity then it's positive and the second one felt like that. <expression> It was a totally enjoyable experience. So, if that was to put that in here (*the questionnaire*) I'll come back and fill it out.

Clip Keywords:

response : enjoyment
response : expression
response : like/dislike/preference
response : limitation

Clip: I felt I was playing with the instrument

Collection: j2

Time: 1:46:06.2 - 1:46:38.7 (**Length:** 0:00:32.5)

Episode Transcript: aj-notes

Clip Transcript:

J (*Reading*) 'I felt in control of the instrument.' I felt I was playing **with** the instrument, so... Sorry to be such a pain in the arse. I feel very strongly about this stuff. Alright circle spheres which is the last one.

Clip Keywords:

response : control
response : effect vs. conversation

Clip: Bachs you have to work

Collection: j2

Time: 1:47:10.7 - 1:49:12.2 (**Length:** 0:02:01.4)

Episode Transcript: aj-notes

Clip Transcript:

J That's an interesting example. The pieces I write, one of them is 14 bars long, 2 part, that's it. Yet it's 10 minutes. So I go for the simplest vehicle to allow the people...

⌘<6457665>A Yeah ok. In the same way that the simplest instruments can sometimes...

⌘<6462331>J So there's that.

⌘<6463834>A It part of the simple complexity somehow.

⌘<6467066>J Like that trombone that I play.

⌘<6472096>A A King?

⌘<6473880>J It's a Bach 16G with dual bore slide. So, for some people it's limiting. For me...and, like you can pick up a Yamaha and get a good sound immediately but Bachs you have to work but because of that it gives you this range.

⌘<6497024>A It was always the same with the bass trombones. You'd pick up a Yamaha and go 'Oh my god I'm a guru' but that was it. What you got after 5 minutes was what you got.

Clip Keywords:

response : complexity
response : flexibility
response : simplicity

Clip: positive discomfort

Collection: j2

Time: 1:49:17.0 - 1:57:57.3 (**Length:** 0:08:40.3)

Episode Transcript: aj-notes

Clip Transcript:

L It'd be interesting to know.. The circle one is the one you've said the least about.

⌘<6567213>J The third one. Yeah I'm having trouble. God it's annoying.. I hate the questions...

(*Muttering...*)

⌘<6587472>L Well you played.

⌘<6590643>J A long time? Great agree then. It's indicative that I enjoyed myself.

⌘<6599812>L Well you played 3 and then you played 4 (minutes) remember you played twice with the circles.

⌘<6608575>J This is with the heads right.

⌘<6613377>L And when you started focusing on the music, which is the first one you said you were (?) and then you went back and focused on the visuals.

⌘<6622716>J I played differently. [L You said it was great.] I still prefer the second one.

⌘<6629543>L But relative to the first one, how does this one relate?

⌘<6636246>J Oh heaps better, heaps better.

⌘<6640547>L Because you described it as being like hanging lanterns. It almost seemed as if it was more intuitive.

⌘<6644862>J Yeah without question the string spheres was absolutely enjoyable. The circle spheres was interesting. The spinner I found annoying most of the time, cool, and I was working - which is not necessarily bad. It's like I had to fight. It felt like I was fighting it consistantly trying to find, 'well how can I make something happen?'

⌘<6671519>L How about circles (??)? I mean I don't want to condition your response but [J I went into other areas...] I thought you were experimenting.

⌘<6680550>J Absolutely, particularly when I was doing with the visual you know. The second one (*string*) is calming. It's very natural and calming. (Pause) Sorry can I write on the back of here (*questionnaire sheet*)? So what is it: spinners - constraining, limiting, same thing I know but..., unnatural. Good, what do you call the second one then? [A String] String - natural, stimulating, and I'll put in brackets 'suggesting new harmonies', in fact I really enjoyed that alright. If you've got a person with good ears, they're going to love that, because it's like you're playing with someone - it's suggesting ideas. I don't know how you're coming up with them and I don't want to, because otherwise it'll affect it too much. What do you call the third one? [A/L Circle] Circle - less stimulating, right, harmonically. However, um... Actually it felt quite natural, alright? Because it was less stimulating harmonically it made me work harder.

⌘<6805803>L Could you get complex effects?

⌘<6808788>J Yeah, but I had to... I was leading, I had to lead. This one was suggesting ideas.

⌘<6819902>A Complex effects - complex musical effects? Not the visual effects?

⌘<6824155>J oooohh...yeah. I was less comfortable with the circle, therefore it was stimulating in a sense. I chose to engage with it and go 'well, shit alright, what can I do with this?' It's more... Maybe this sounds dismissive - I don't mean it to... More game-like.

⌘<6850163>L Was it more challenging?

⌘<6853518>J It was engaging because... Yeah it challenged me because I didn't feel as comfortable. So it's that positive discomfort possibly. Yeah, positive discomfort that's (what I'm gonna call it?). I can see this in your thesis.

⌘<6877297>A This is right. That's it - it's the title. Absolute discomfort.

⌘<6883021>J Well, check this out. One of my favourite phrases is on Bernie M... I don't know the origin of it. But um, he's got an album and I play on it. I think I play on that album - I think I've got a song on that album. But it's called 'Ugly beauty' and I think that's just fantastic. I just think that's so wonderful, if you're talking Shoenberg... Like Kraanerg, after... you know this Xenakis piece. After a period of time I loved it. I loved... I was just waiting for this next majestic phrase to come out you know. Lloyd Swanton played in that as well - a number of jazz players did and freelancers, because it was so non-inside the ball, you know. And he sings me back my solo, he sings me back my lines 15 years

later! (*Sings...*) I'm sorry I'm on a rave. Positive discomfort...stimulating. So...drew...me...in...to... What do you call it? It was...poking me kind of thing...you know like, 'What are you gonna do?'

⌘<6965891>L Provocative?

⌘<6967194>J Provocative! Thankyou! Because it was less harmonically complex, I had to find other ways to create complexity. Is that cool? (*muttering while writing...*) I don't know how to spell by the way...

⌘<7014646>A While you're thinking about that circle spheres by the way. One thing I haven't really mentioned yet is that that one had colour in it. Or at least more colour than the string spheres did. [J That's true, it did too.] And I noticed you playing with different sound colours to get different visual colours.

⌘<7032304>J Ahhh...That's interesting - I hadn't... I wasn't...

⌘<7035496>A Maybe you weren't doing that consciously but...

⌘<7036942>J Yeah yeah, that's why I'm fascinated in this. This is what I think... You know I don't know what's going on all the time. That's fascinating- very interesting. I hadn't thought about that. It's funny because with the string thing I could see colours. They may not have been there.

⌘<7054476>A There is some colour there but...

⌘<7056192>J No but I could see lanterns funnily enough. Because it reminded me of a physical thing I remembered lanterns. Yeah, you know, hanging like this. Or foam balls, like density foam balls. I was trying to think what it would be that would be acting like that. Anyway it's interesting.

Clip Keywords:

behaviour : can I make it...?
behaviour : what does it do if I...?
feature : audio
feature : visuals
response : complexity
response : enjoyment
response : interest
response : like/dislike/preference
response : limitation
response : metaphor
response : natural

Clip: opened my mind up

Collection: j2

Time: 2:02:52.7 - 2:03:26.6 (**Length:** 0:00:33.9)

Episode Transcript: aj-notes

Clip Transcript:

J If you actually want to get something performed - call me. Particularly if it's with the second or third one. If you do something with the first one you can stick it in the middle between the second and third one and we've got three movements. I'm really busy but I'm really quite excited about this. Particularly cause I was so closed minded about it before and this has actually opened my mind up quite a lot. Help me clarify a few things.

Clip Keywords:

response : interest
response : like/dislike/preference

Summary

behaviour : can I make it...?	1
0:08:40.3	
behaviour : exaggerates/sensitivity	3
0:04:09.9	
behaviour : playing	7
0:39:36.1	

behaviour : request	1
0:03:27.0	
behaviour : software affects music	4
0:20:05.9	
behaviour : sonic exploration	2
0:11:39.0	
behaviour : visual exploration	2
0:14:04.0	
behaviour : what does it do if I...?	1
0:08:40.3	
feature : audio	13
0:56:21.7	
feature : bug	1
0:01:32.8	
feature : mapping	5
0:22:44.1	
feature : visuals	9
0:47:44.1	
response : complexity	9
0:35:19.4	
response : consistency	6
0:30:25.4	
response : control	8
0:40:36.6	
response : distraction	1
0:06:42.5	
response : effect vs. conversation	8
0:30:29.8	
response : enjoyment	9
0:35:40.2	
response : expression	2
0:04:08.4	
response : flexibility	2
0:12:19.9	
response : how it works	3
0:03:55.9	
response : interest	2
0:09:14.2	
response : like/dislike/preference	9
0:34:22.7	
response : limitation	5
0:23:57.3	
response : link	2
0:11:39.0	
response : metaphor	4
0:30:00.0	
response : motivation	2
0:03:36.9	
response : musical strategies	8
0:33:01.5	
response : natural	5
0:23:46.7	
response : power balance	3
0:15:54.6	
response : simplicity	1
0:02:01.4	
response : suggestion	8
0:33:03.9	
response : surprise	4
0:12:12.6	
response : use/purpose	3
0:06:12.8	

Clips:
1:46:38.3

34

Total Time:

Appendix D

Interview with Ben Marks

This interview with composer/trombonist Ben Marks was conducted via iChat (a text-based internet chat tool) on January 9, 2007. It provides additional perspective on the nature of the virtual instruments and musicians' experiences with them. I refer to this interview in chapter 5, section 5.3.3.

Interview with Ben Marks

January 9, 2007

AIM IM with Ben

11:15 AM

Ben: maaate

Andy Johnston: mate!

Ben: so

Andy Johnston: just let me make sure i can save these transcripts...

Andy Johnston: ok seems to work

Ben: fire away

Andy Johnston: well.. the main purpose of the interview is to get feedback on the nature of the 'virtual instruments'...

Andy Johnston: First up: what makes the partial reflection toys successful/not successful as instruments

for "musical expression"

Andy Johnston: Maybe start with spheres..?

Ben: Partial Reflections 1

Ben: ?

Andy Johnston: yeah we can start with the first movement if u like..

Ben: PR1

11:20 AM

Andy Johnston: yep

Ben: PR1

Ben: sorry I'm doing the wrong thing

11:25 AM

Andy Johnston: mate r u still there?

Ben: PR1. Successful elements: The movement of the chain of spheres had a natural, flowing quality which made clear connections with the sounds played into it - pitch to ball mapping. This mapping was highlighted by change in colour as well as movement, further reinforcing the relationship between sound and image. The sounds emitted by the virtual sculpture, while derived from the sounds going in, were interesting and sustained enough to create a complex auditory environment. The otherworldly nature of this environment was important in feeling engaged with the sculpture and its processes.

11:30 AM

Ben: PR1. Not successful: The limited variety of action of the spheres made them predictable after a short time so the length of time you might choose to spend with them is shorter. The lack of change through time is part of this same limitation - time spent with the sculpture is more about working out what is going on and once that is done you are left with the "play" element. You enjoy playing with them but it can seem without purpose.

Ben: PR2 now?

Andy Johnston: um.. can we look at some of what u just said in a little more detail?

Ben: sure

11:35 AM

Andy Johnston: First up, the 'natural, flowing' quality and clear links between sound and behaviour of the virtual instrument. can u say why this was so important?

Andy Johnston: (mate could I ask you to hit return after each sentence maybe?)

Ben: Connection. There is a cognitive aspect (if that is the right word) and an intuitive aspect. You can recognize your actions in the way they push the balls and you recognize something visually (supported by the nature of the sounds) that is more about water or air.

11:40 AM

Ben: the recognition can be something intuitive, or subconscious

Ben: the flowing nature brings about this intuitive response I think

Andy Johnston: any comments about this in relation to the audience's experience of the work?

Ben: I guess I assume it is much the same as mine - there is a sense of identification with the sculpture

Ben: having me there performing adds a different element in the sense that there is a kind of dance going on or that people are witnessing a conversation

Andy Johnston: yep i see..

Ben: kind of erotic, kind of like stroking a cat.....

11:45 AM

Ben: or dog

Andy Johnston: but not a porcupine

Ben: PR2

Ben: PR2 is more porupine

Ben: porcupine

Andy Johnston: anyway...I was also interested in the 'otherworldly' term. It's interesting that the response

was 'natural' but still otherworldly. any comments?

Ben: otherworldly in that the sound, coming from the sculpture, creates a sonic landscape

Ben: it is unique to the sculpture

Andy Johnston: was there some characteristic of the sound that made it otherworldly do you think?

Andy Johnston: a number of people have mentioned that the sounds are 'eerie', etc.

Ben: the overtones certainly gave a sense of space - their length of playback and natural decay perhaps made you think of a larger space

11:50 AM

Ben: perhaps the eerie quality comes from the complexity of the sound - there is quite a lot of sonic information coming out (an illusion to a kind of 'intelligence')

Andy Johnston: yeah i see what you mean..

Ben: the spheres kind of sing to you but we can't be sure what about. This lack of recognition (to contradict myself earlier) makes people feel a bit eerie maybe

Andy Johnston: yeah i see...

Andy Johnston: they recognise what is happening but it's still strange?

Andy Johnston: unusual language maybe..

11:55 AM

Ben: perhaps there is a contradiction that is good - you feel connected to the work, in its world, the flow, it's behaviour, it's direct response BUT you can never really know how it 'feels' or what it 'thinks'

Andy Johnston: yeah nice. a bit senstery..

Ben: it looks like it might be trying to say soemthing

Ben: yeah

Andy Johnston: ok, moving on to the limitations...

Andy Johnston: I understand you mean there is a kind of 'exploratory' phase where you are trying to figure out what's happening, but once you get it the interactions is less satsifying?

Andy Johnston: (are)

Ben: in a sense yes - there is also the fun of playing with it without this goal of 'working it out'. A legacy of computer gaming?

Ben: Having a performer helps avoid the working out thing a bit perhaps

Ben: the work is about relationship rather than in generating one directly with the audience - a step removed perhas

Andy Johnston: hmm. you mean the relationship between performer and virtual instrument?

12:00 PM

Ben: yes

Andy Johnston: so you felt less directly connected with the audience perhaps?

Ben: watching someone interact is important somehow

Ben: I think the conversation between me and the sculpture is kind of overheard by the audience

Andy Johnston: is there a difference between watching you interact with the trombone and instrument vs.

you interacting with trombone alone?

Ben: it is quite a public sculpture perhaps. The trombone isn't so much an issue in live performance (as an object) it's more that instead of it being the spheres and me it's the audience and me

Ben: I'm communicating with the spheres rather than the audience

Andy Johnston: v. interesting- the way that the natural movement of the spheres seems to give the impression of intelligence.. something you can converse with. would they seem as intelligent/interesting without sound?

Ben: not as interesting but still interesting I would say - like a zoo exhibit behind glass

Andy Johnston: yeah- I can see you don't really communicate with the trombone (as a thing) in the same

way as you communicate with the spheres.

12:05 PM

Andy Johnston: nice metaphor!

Ben: yep

Ben: not the metaphor though

Ben:

Andy Johnston: so maybe- you use the trombone more as a medium and the spheres as a kind of.. well I

was going to say muse but that's probably not the right word..

Andy Johnston: stimulant? partner?

Ben: beast

Andy Johnston: amen

Ben: oracle idea comes from here you see

Andy Johnston: yeah i can see that..

Andy Johnston: ok, this is an important point i think. You don't really conceive of the spheres as an instrument in the same way as the trombone then?

Ben: no. It is more another musician I'm jamming with, although predictable and only really responding to me. The sound and look of this other musician though is enough to want to jam though

12:10 PM

Andy Johnston: i see. Do you conceive of PR2 in the same way?

Ben: the composed piece is kind of like an agreement

Ben: PR2 is a bit different I think

Andy Johnston: it's more of an instrument?

Ben: I feel more like it supports and empahsizes what I'm doing - it aids my communication with the audience

Ben: I never thought of that before...shit.....

Andy Johnston: can you think of any analogies from the trombone world? like a mute maybe? or some

kind of electronic effect?

Ben: a kind of light show. the sound was interactive though but more like a sampler that manipulates than it's own intelligence

12:15 PM

Ben: hard one to think of. Not a mute but a periphery add-on device of sorts

Andy Johnston: but not something you 'converse' with a la spheres?

Ben: there was still a type of conversation but the sculpture was more fixed in its ways

Andy Johnston: yes I see.. Can you say what the successful/unsuccessful characteristics were?

Andy Johnston: (back in 1 min)

Andy Johnston: (back(

Ben: Just one more thing - if PR1 was stroking a cat, PR2 was more like picking up the cat and dropping it again and again - to see what it does. The continuous nature of PR1 was different to the tableaux

nature of PR2

12:20 PM

Ben: PR2. Successful: strong clear relationships between sound and image (sound burst = image "burst" or stretch), sound feeding back had a strong independent nature (glissing up) , could take a

large variety of sounds and create varying textures, strong visually

12:25 PM

Ben: PR2. Less successful: intuitively not as strong as PR1, less sense of the sculpture trying to communicate back (more direct nature of the sound perhaps), how to stop in performance, lack of evolution

through time (you learn it quite quickly as well), less variety/complexity in sound (more rhythmic activity though),

Ben: its a start

Ben: One thing that does occur to me is that the sounds I made for PR2 were more complex than PR1 and perhaps needed to be

Andy Johnston: yeah I agree. um, Can I clarify the continuous vs. tableaux thing? What makes PR2 more

continuous? More movement/action?

Ben: I was thinking about the composed stuff. I wrote like that for each sculpture

Ben: suited their natures I think

Andy Johnston: i see

Ben: PR1 is definitely more human I think - tall, wavy etc.

Ben: side issue

Andy Johnston: yep- can you say why PR2 was less intuitively understandable?

12:30 PM

Ben: maybe because of that - less recognizable. We are less familiar with this type of visual event

Andy Johnston: which type? the spinning?

Ben: Yep the spinnign. don't really see many windmills and particles etc.They are more in the imagination than real visual/experiential worlds

Andy Johnston: but the spheres are in real world?

Ben: spinners were like an engine

Andy Johnston: ah i see.. more mechanical- less natural?

Ben: spheres evoked more real world references I think

Ben: references to things natural and flowing - water mostly for me

Andy Johnston: yep i get it..

Ben: spinners maybe a bit conceptual perhaps - the idea of hot particles is good but it just don't evoke as much perhaps

Andy Johnston: you mentioned the sound for PR2 being more independent than PR1.. Can you expand on that?

Ben: PR2 sampled sounds with less manipulation, other than shortening them and variously transposing them

12:35 PM

Ben: Unlike PR1 it didn't have it's own sound world really - with out an input you feel it would be defunct but you sense that PR1 exists despite

Andy Johnston: ok so to me that sounds like pr2 is **less** independent in a way?

Ben: visually both are quite independent

Ben: more to do with sound and a more inuitive response to image. Part of it is also that the sounds in PR1 created a sense of space with it's sounds (resonance in a virtual space) where was PR2 was

direct and very much in the space with you - less mysterious, more obvious that the sounds came from the input

Andy Johnston: ok- i think i see what you mean. are you happy to keep going or do you want to finish up soon?

Andy Johnston: (happy to come back to this later if you've got stuff to do...)

12:40 PM

Ben: might just do a few things and come back. Shall we say 10 minutes?

Andy Johnston: ok see you at 12.50ish

Ben: 11:50 in 1987 for me!

Andy Johnston: no prob- see you then. say hi to jason & kylie for me.

12:50 PM

Direct Instant Message session started

12:55 PM

Ben: okay I'm ready to go. Jason and Kylie have gone home

Andy Johnston: now..my next question is: what do you think are the critical characteristics of these 'instrument extensions' (or whatever) are?

Andy Johnston: (and I'm not referring to santa's instrument)

Ben: PR1 - intuitive visual appeal (it's "natural" quality), it's own sound world, responsiveness, tweakability is also important

Andy Johnston: i think by critical characteristics i mean of both..

Andy Johnston: so if/when we make more, what characteristics would you expect them all to have..

1:05 PM

You left the chat.

Reconnecting to Ben...

Ben: you're gone?

Andy Johnston: don't know what happened?

Ben: could be WORD!!!!!!!

Andy Johnston: I just got the thought bubble from you and then ... nothing!

1:10 PM

Andy Johnston: bloody microsoft

Ben:Microsoft...

Ben: did you get my previous stuff?

Andy Johnston: last i got was "tweakability is also important"

Ben: of both - well maybe much the same really, responsiveness, intuitive visual mapping of sound, notions of independent activity (sound worlds - visuals will always be independent anyway), conversational qualities to various degrees, sense of character in the type of relationships generated (placating, arousing, punching, poking etc. with associated emotional responses) probably many more..... time to generate and explore is also crucial to the project
co-express is a word we used in apps. this might be a characteristic - that there is a sense of co-expression BUT I'm no sure any more. Even as we chat I'm getting different ideas about where to go.....
Imagine if we worked our visuals to more human like shapes - could become stronger this way and thus BECOME (rather than be) a defining characteristic Part of this is that I still we are in a state of flux with the work rather than finished. We are still learning this sound/visual language and attempts to define characteristics of our work reflectively might not be useful or might not indicate future directions
Am I getting a bit ahead of myself here?

Andy Johnston: (u can cut and paste if necessary)

Andy Johnston: ah- got it

Ben: I did mate.

Ben: I might be a yokel computer guy but mate I can cut and paste!

Andy Johnston: (sorry)

Andy Johnston: sorry- reading

Ben: should we get together this month by the way? Might be difficult (Mum is up and down alot with her health) but would be nice

Andy Johnston: I'm around but have Z to look after. She is in day care Mon-Tues though...

Ben: will think on this.....

Andy Johnston: mate- i heard a chat note come through but didn't see anything?

Ben: responses to text?

Ben: can you see me now?

1:15 PM

Andy Johnston: yep but may have missed something. did u respond to my "I'm around.." text?

Ben: yeah

Andy Johnston: ah- did u say will think on this.....

Ben: yes

Andy Johnston: ok- then all is well now I think (but something a little strange may be happening)

Ben:bill gates

Andy Johnston: so this question (bloody bill) is really going to my PhD.. Trying to figure out what you think are the key things that make this type of software work..

Andy Johnston: natural response seems important

Andy Johnston: consistency?

Ben: pd seems to operate well like this - natural laws etc. so this would be important'

Ben: I think flexibility is important

Andy Johnston: flexibility in the way the software can be used or flexible approach to designing it?

Ben: Flexibility in design

1:20 PM

Ben: the sculptures themselves also need to be flexible even if this is just volume control, or the various sliders you use

Andy Johnston: yep

Ben: the work has to be adaptable to environment

Ben: strong core of functionality gives it this adaptability

Ben: functionality

Ben: bach chorales played on piano, sung, organ, trombone quaret - strong core

Andy Johnston: by functionality you mean what the software does. a strong core idea?

Andy Johnston: (still there?)

Ben: yep, but also what the music does - can't write strong core music and expect the visuals to fit (no room) or have a wild sculpture on its own

Andy Johnston: (ok)

Ben: core duo

Ben: core interaction

Ben: wish I had a macbook

Andy Johnston: i know mate...let it go...

1:25 PM

Andy Johnston: they will be passe soon..

Andy Johnston: i want a nokia n800

Ben: the core I think is about the music/visual relationship

Ben: not necessarily how closely linked or how similar they are or what they share (see Xenakis quote on my blog) but more how clearly defined their relationships are to each other. Out of this the work can

grow on its own

Andy Johnston: the mapping between sound and image/sculpture behaviour...

Andy Johnston: maybe it's not just how clearly defined they are but how appropriate the mapping is?

1:30 PM

Ben: I think the mapping grows out from something in the core - same with the music

Andy Johnston: the core of the sculpture?

Andy Johnston: or the core of the relationship b/w sound & visuals?

Ben: core I see as something quite stable, a vague idea even, concept, something that needs articulation (which is the work)

Andy Johnston: there is a goethe quote on my blog which might be also in the same area you are getting at..

Andy Johnston: but does this core idea exist before we start to design/write something?

Andy Johnston: "Colour and sound do not admit to being compared together in any way, but both are referable to a higher formula, both are derivable, though each for itself, from this higher law. They are like two rivers which have their source in one and the same mountain, but consequently pursue their way under totally different conditions in two totally different regions, so that throughout the whole course of both no two points can be compared."

Andy Johnston: (Goethe)

Ben: yes and no I think - the core is preceded by the desire to do this in the first place. Each project core is a version of this basic desire

Andy Johnston: maybe let me put the question another way..

1:35 PM

Andy Johnston: What we did for PR1 & 2 in a technical sense was design an interactive sound-sculpture

thing that could be controlled/played by sound...

Ben: that's a bit of core

Andy Johnston: What I want to ask is, what characteristics of these sound sculptures are critical. ie. if they

were missing, the whole thing just wouldn't work..

Andy Johnston: I've thought that consistency is one important characteristic for example. You play two

Bbs and the sculpture responds in the same way. Do you have any other key characteristics?

Andy Johnston: (and do you think consistency is important..)

Ben: Yes. Stability is much the same I guess - actually key in any concept of core. The sounds from PR1 were kind of unstable but they were produced stably

Andy Johnston: unstable in the sense that they were unpredictable?

Ben: yes

1:40 PM

Ben: a degree of unpredictability is important though - keeps interest and ideas surrounding exploration going

Andy Johnston: and how was the production stable? do you mean consistent within itself?

Andy Johnston: ie. 2 Bb's = same response

Ben: yes consistent within itself

Andy Johnston: cool. any other key ideas?

Ben: you worked alot at this I think and the fact the music is composed is also a desire for consistency

Ben: formalizing

Andy Johnston: yep

Ben: Is formalizing our work key to it's success?

Andy Johnston: maybe for us. we are poor classically-trained musos..!

Ben: composed rather than improvised, predictable rather than random

Andy Johnston: scared to go out unprepared!

Ben: but the preparedness let our work succeed in environments outside of which it was created
- the work speaks despite poor conditions

Ben: except power failure

Andy Johnston: yeah- that would be a problem!

Andy Johnston: maybe if i could get a torch...

Andy Johnston: or 12 torches

1:45 PM

Ben: It might be worth trying to duplicate our sculpture with torches and a cazoo just to see if the idea is stronger than it's actual realization? Would it work?

Andy Johnston: sure. probably be much easier to coordinate too!

Ben: but if it did work that would say something I think.....

Andy Johnston: true... mate- maybe we can call it a day for now?

Ben: sure. i have to go to McDonalds. Lots of whisky last night

<Farewells...>

Ben has gone offline.

Ben is now online.

Appendix E

Guidelines for Observers

The following guidelines were provided to the observers who attended the user study sessions to provide additional perspective. See chapter 5, section 5.4.1 for details.

Investigation of musicians' experience with sound-controlled virtual instruments based on physical models

Guidelines for Observers

Thankyou for agreeing to help with our research. The aim of this study is to gain insight into brass musicians' experience using prototype software we have developed. The software allows the musician user to interact with a kind of 'virtual instrument' that the computer displays. By playing and/or singing into a microphone, the musician can push, poke and prod the virtual instrument which moves around in response. As it moves it also produces sound. This behaviour in effect allows the musician to 'play' the instrument with sound.

There are several criteria that we kept in mind when designing the prototypes. These were:

1. The instruments should respond in a way that seems natural. In the prototypes we've tried to do this by basing the instruments on virtual physical models that respond in physically plausible ways.
2. Instrument response should be consistent. ie. Two perceptually identical notes should appear to have the same effect on the virtual instrument;
3. Instruments should be simple but allow skilled musicians to create complex effects;
4. The instruments should be interesting, engaging and motivate the musician;
5. The musician should feel in control of the instrument, but the instrument should retain the ability to surprise, in the sense that the musician may discover something about their technique or gain insight into some aspect of their music making by using the instrument;
6. The instrument should encourage a playful, exploratory approach, especially in new users. They should encourage the musician to consider questions such as 'What does it do if I play...?', 'How can I make it...?'
7. The relationship between live sound, the behaviour of the virtual instrument and the resulting sounds should be apparent to observers (eg. audience members).

These criteria were all oriented towards ensuring that the virtual musical instruments facilitated musical expression and exploration. The key questions we want you as an observer to keep in mind are:

- How well do the prototypes embody the above design criteria?

- How successfully do the prototypes encourage and facilitate musical expression?
- Which specific characteristics of the prototypes help or hinder musicians in this regard?
- What impact does using the software have on the musicians' music-making?

We don't know exactly what you will observe, and as this is a qualitative investigation we don't require you to count occurrences of specific behaviours. Rather we are interested in anything you observe which will help us to further understand the musicians' experience with the prototypes and how these might relate to the design criteria.

Examples of the kinds of evidence that could be relevant:

The musician plays a note, stops, looks at the screen with a puzzled expression on his/her face and repeats the same note. The response from the model is different and the musician shrugs.

The musician plays a few notes and laughs and says, "That's interesting!"

As one of the design criteria states that the relationships between the sounds the musician makes and the behaviour of the software should be apparent to observers, we would also be interested in any comments on how apparent you found these links to be.

We want to get as rich a picture as we can of the interaction between musician and the prototypes from your perspective. Please try not to self-censor – any comments and observations that you feel might be relevant to the research questions should be noted down.

Debrief

Shortly after the evaluation session we will conduct a debriefing session with you to go over your notes and discuss your observations.

Again, thank you for your time. If you require clarification of any points please ask.

Appendix F

Consent Forms

The consent forms completed by participants in the user and audience studies are included on the following pages.

UTS : IT : CREATIVITY & COGNITION STUDIOS

INVESTIGATION THE AUDIENCE'S EXPERIENCE OF THE AUDIO-VISUAL WORK "PARTIAL REFLECTIONS": UTS HREC 2004-011P

CONSENT FORM

I _____ (*participant's name*) agree to participate in the research project "Investigation of the audience's experience of the audio-visual work *Partial Reflections*" being conducted by Andrew Johnston of the Creativity and Cognition Studios at the University of Technology, Sydney.

I understand that my participation in this research will involve being interviewed about my experience during a concert which includes the audio-visual work *Partial Reflections*. I understand that the whole procedure will take approximately 5-10 minutes of my time. I agree that the audio of the interview may be recorded.

I understand that the purpose of this research procedure is to investigate my experience of the audio-visual work *Partial Reflections*. I also understand that my views as expressed in the interview may be quoted or referred to by the researchers within their PHD theses and/or in published journal articles and/or conference papers.

I agree that the research data gathered from this project may be published in a form that does not identify me in any way. I understand that the data will be stored securely and confidentially at UTS for at least 5 years.

I am aware that I can contact the researchers or their supervisor Prof. Ernest Edmonds if I have any concerns about the research. I also understand that I am free to withdraw my participation from this research project at any time I wish and without giving a reason.

I agree that the researchers have answered all my questions fully and clearly.

Signed by _____/_____/_____

Witnessed by _____/_____/_____

NOTE:

This study has been approved by the University of Technology, Sydney Human Research Ethics Committee. If you have any complaints or reservations about any aspect of your participation in this research which you cannot resolve with the researcher, you may contact the Ethics Committee through the Research Ethics Officer, Ms Hadiza Yunusa (ph: + 61 2 9514 9615, Research.Ethics@uts.edu.au) and quote the UTS HREC reference number. Any complaint you make will be treated in confidence and investigated fully and you will be informed of the outcome.

UTS : IT : CREATIVITY & COGNITION STUDIOS

INVESTIGATION OF MUSICIANS' EXPERIENCE WITH SOUND-CONTROLLED VIRTUAL INSTRUMENTS
BASED ON PHYSICAL MODELS: UTS HREC 2004-011P**CONSENT FORM**

I _____ (*participant's name*) agree to participate in the research project "Investigation of musicians' experience with sound-controlled virtual instruments based on physical models" being conducted by Andrew Johnston of the Creativity and Cognition Studios at the University of Technology, Sydney.

I understand that my participation in this research will involve using interactive audio-visual software and being interviewed about my experience. I understand that the whole procedure will take approximately 1-2 hours of my time. I agree that my experience of the software and any post experience interview may be recorded on video and audio.

I understand that the purpose of this research procedure is to record my experience of the interactive software. I also understand that the record of my experience and my views as expressed in the interview may be quoted or referred to by the researchers within their PHD theses and/or in published journal articles and/or conference papers.

I agree that the research data gathered from this project may be published in a form that does not identify me in any way. I understand that the data will be stored securely and confidentially at UTS for at least 5 years.

I am aware that I can contact the researchers or their supervisor Prof. Ernest Edmonds if I have any concerns about the research. I also understand that I am free to withdraw my participation from this research project at any time I wish and without giving a reason.

I agree that the researchers have answered all my questions fully and clearly.

Signed by _____/_____/_____

Witnessed by _____/_____/_____

NOTE:

This study has been approved by the University of Technology, Sydney Human Research Ethics Committee. If you have any complaints or reservations about any aspect of your participation in this research which you cannot resolve with the researcher, you may contact the Ethics Committee through the Research Ethics Officer, Ms Hadiza Yunusa (ph: + 61 2 9514 9615, Research.Ethics@uts.edu.au) and quote the UTS HREC reference number. Any complaint you make will be treated in confidence and investigated fully and you will be informed of the outcome.

Appendix G

Questionnaire

The questionnaire administered during the user study is shown below. To avoid confusion, simplified names for each of the virtual instruments were used. On the questionnaire, 'Circle Spheres' refers to *Spheres of Influence*, 'String Spheres' to *Partial Reflections I* and 'Spinner' to *Partial Reflections II*.

Investigation of musicians' experience with sound-controlled virtual instruments based on physical models

Questionnaire

Please complete the following questionnaire by circling the response that best reflects your level of agreement with the following statements.

Name: _____ **Date:** _____

Virtual Instrument (circle): Circle spheres String spheres Spinner

The instrument responded to the music in a way that I would describe as 'natural'.

Strongly disagree Disagree Neutral Agree Strongly agree

The instrument's behaviour was consistent. (ie. It responded to identical notes in the same way.)

Strongly disagree Disagree Neutral Agree Strongly agree

The instrument was conceptually simple.

Strongly disagree Disagree Neutral Agree Strongly agree

The instrument allowed me to create complex musical and visual effects.

Strongly disagree Disagree Neutral Agree Strongly agree

I found the instrument interesting and engaging.

Strongly disagree Disagree Neutral Agree Strongly agree

I felt in control of the instrument.

Strongly disagree Disagree Neutral Agree Strongly agree

Appendix H

Transcript of Audience Study

This is a transcript of a series of short interviews conducted with audience members who attended a concert featuring the works for virtual instruments at the Queensland Conservatorium on May 17, 2007.

**Interviews with audience members
Concert including performance of Partial Reflections
Queensland Conservatorium
May 17, 2007**

A = Andrew Johnston (interviewer)
am1 - am7 Audience member interviewees
((...)) - Observed action or comment

§<11863>A ((Asks what he thought relationship was between what Ben did and what was on the screen.))
§<20105>am1 Not too many ideas. It seemed as if he had eyes in the back of his head and seemed to see what the screen was doing and seemed to follow what the screen was doing more than the other way around. I'm not sure how to say it, but I'm fairly sure Ben doesn't have eyes in the back of his head. It seemed as if it was really well written together.
§<47861>A So do you think... Was your perception that what was happening on screen was kind of pre-recorded and that he was playing along with it?
§<55541>am1 um, yes and no. [A You're not sure?] Yeah I wasn't quite sure. I feel some of it might have been but other elements seemed to be a bit on the spot.
§<71768>A ((Thanks, etc... signing forms...))
§<94936>A ((Approaches am2.))
§<109256>A Just wondered if you could explain to me what you thought was going on in relation to what he was playing and what was going on screen?
§<119961>am2 um I thought there was, um like, it was an echo kind of effect - not really, just kind of slightly. And the little wormy thing would move, like recede to the back and, yeah it was good.
§<137740>A So did you see a relationship between what he played and what was coming up on the screen?
§<140180>am2 Yeah
§<141919>A Ok, and do you think that was pre-recorded or do you think... Like what was your understanding or your impression? Do you think it was pre-recorded or was it happening...?
§<151415>am2 Not it looked like it was happening at the same time. Like it would just follow what he was doing, pretty much straight away.
§<159387>A ((Thanks, etc...))
§<199563>A ((Approaches am3))
§<210265>A What we're specifically interested in is... if you could explain to me what your impression was about how it responded to the sounds Ben was making? Did you see a relation between what Ben did and what happened on screen?
§<221928>am3 Oh, yeah! ok.
§<225150>A So you did... So what was your impression of what the relationship was? Do you remember?
§<229814>am3 Kind of moved... well to me it looked like it moved with him dynamically like as he moved or as his sound grew the little dots got bigger or whatever. Sorry!
§<241468>A No, that's exactly the kind of thing we're looking for! So, there was two movements. The first one was like a string one. [am3 That was like the dots and string, yeah.] and the second one was the dot in the middle... (unclear) Did you see the same kind of relationship happening with both of them or was there a difference?
§<257384>am3 Yeah. The second one. Oh I don't know, just how I saw it move-like more with tempo or intensity rather than dynamically or...so.
§<267094>A ((Thanks, etc...))

§<293702>A ((Approaches am4))
§<306005>A Specifically I'm interested in your understanding of what the relationship was between what Ben played and what happened on screen and what you heard. Did you see any relationship there?
§<316550>am4 Um, direct relationship yeah.
§<317470>A You saw a direct relationship?
§<319038>am4 oh yeah.
§<319644>A Could you give me just your brainstorming impressions of what the links might have been between what he played and what you saw?
§<325332>am4 I think there was a pitch correlation, a dynamic correlation and it was pretty obvious visually what was sort of happening with each thing. But when he played multiphonics and stuff he got like variations of everything. And it was really interesting the way he mixed it up - soft and louds with different effects.
§<343681>A ok, cool. The two movements - if you remember - did you see a similar relationship with both of them between what he did and what happened on screen or was there a difference?
§<351280>am4 oh, I thought it was different. Yeah I thought the first one was a bit more sort of obvious than the second one was pretty subtle and that whole raising of pitch all the time was sort of hypnotic in a way. So yeah, I wasn't really concerned with the visual too much in the second movement I was just really listening to what was going on.
§<373647>A ((Thanks, etc...))
§<408047>A ((Approaches am5))
§<425062>A We're interested in if you saw a relationship between what Ben played and what happened on screen and the sounds that were coming back from the computer?
§<432949>am5 Certainly some of them, yes. Definitely.
§<435981>A Can you give us your off the top of your head impressions about what the link might have been between what he played and what the computer did?
§<441861>am5 Well I think he was trying to imitate balls flying around. That was the impression I got with the second one particularly.
§<450369>A Yeah, ok and the link between what he played. Like the sound that he made.. Did you see a link between the sound that he made and what was happening?
§<457777>am5 Sometimes. I won't say all the time, but yeah sometimes.
§<460855>A Ok, sometimes. ((Thanks, etc...))
§<531895>A ((Approaches am6 and am7.))
§<567560>A Basically I'm interested in did you see a relationship between what Ben played and what happened on the screen? And what you heard coming back.
§<574911>am6 um, yeah... Yeah I did. It was quite effective I thought especially at the beginning where they had that sort of cluster that sort of went out. But a lot of the effect was the electronics that kind of enhanced that. And I thought that actually a lot of the stuff that he had to play after that kind of didn't necessarily have as much to do with it. And he was doing a lot of really large intervals that didn't seem to have so much to do with what was going on there.
§<608579>A You didn't see such a relationship when he was playing that as with the earlier stuff?
§<611323>am6 But, but there were certainly times when he was doing stuff that was kind of enhancing what was going on with the electronics. I thought the electronic noises that they had definitely had a lot to do with what was going on, but it became a bit monotonous for me after a while.
§<628683>A Ok, sure. And if you had to take like a stab in the dark, like off the top of your head how it actually works. Not technically how it works

but I guess what the mapping is between what Ben does and what happens.
Could you take a guess at how it might work?

§<640169>am6 Oh, God... Well when I saw it I was thinking of like... I was thinking of molecule diagrams and electrons and photons sorts of things and atoms.

§<655372>A So if you had to take a guess at like what his sound did to those molecules, like what the relationship was between what he played and what the molecules did, would you be able to say?

§<663899>am6 oh God. [A Like the louder he played... If he played louder what would you expect to happen say?] Well certainly when he was doing um, the section I didn't like so much, where he was doing all the large intervals, um there were a lot more of those dots and lines mixed in further out towards the edges.

§<687633>A The bigger the intervals, the more it spins out maybe?

§<694985>am6 Yeah I suppose that's true...

§<696933>A ((Thannks, etc...))

§<714388>A ((to am7)) Do you mind if I ask you the same kind of question?

§<719101>am7 Yeah sure. I actually found there was a huge relationship in the beginning between what I was seeing and what Ben was playing.

§<723731>A So that was the string one?

§<725355>am7 The string one yeah. And I was lost, thinking it was like slow motion beads or something like that. Made of something you could see through or something like that. And I found I lost that visual contact a little bit - it wasn't maybe as strong in the second one. But still certainly some relationship.

§<745770>A OK, you saw the relationship. And if you had to take a wild stab in the dark about what the relationship specifically was. Like for example if you played louder what would you expect to happen?

§<753705>am7 I actually thought maybe it was air related, like blowing it - what I was seeing visually maybe?

§<761080>A So like if he was blowing a lot of air you would see it move back further?

§<763992>am7 Yeah maybe seeing like a light object or something like that and being affected by air or something like that. More the physical aspect of the playing than what I was actually hearing.

§<778893>A So not so much the volume of audio but [am7 Yeah, more the physical, yeah.] the volume of air.

§<783355>am7 And how that was affecting visually, yeah.

§<786684>A And in the second one?

§<788164>am7 In the second one? ((laughs)) I have to think what the second one was.

§<790515>A That was the one where it was much faster and was spinning around.

§<792723>am7 Spinning around yeah. I haven't got... That hasn't left much of an impression on me as the first one. With the rest of the concert in my mind it's left me a little bit.

§<803889>A Ok, thanks very much.

§<806488>am7 Um, but did you ask like technically how they actually... Did you ask that before? How it works?

§<814008>A Yeah well I guess I was interested in - not technically how it works to the level of how to program a computer to do the same thing. But just to say like what the relationship was between I don't know the pitch of the notes he played...

§<824733>am7 And visually yeah. Ok, alright.

§<831876>A ((Thanks, etc...))

Appendix I

DVD Track Listing

The DVD which accompanies this thesis contains demonstrations of the virtual instruments and recordings of concert performances.

Track 1 Demonstrations of the virtual instruments. This video is intended to complement the technical descriptions provided in chapter 4 (Virtual Musical Instruments) by briefly demonstrating each of the virtual instruments.¹

Track 2 *Partial Reflections I & II* performance by Ben Marks, recorded at the Queensland Conservatorium on May 17, 2007.

Track 3 *Partial Reflections III* performance by Jason Noble and Diana Springford, recorded at the Sound Lounge, Sydney on November 27, 2007.

¹Unfortunately, due to limitations of video equipment on-screen colours are not clearly identifiable.

Appendix J

Software CD

The virtual instrument software, developed using Pure Data, is included on the accompanying CD. A 'README' file details how to run the software.

Appendix K

Publications

Refereed publications which feature work described in this thesis are listed here.

- Johnston, A., Amitani, S. and Edmonds, E. (2005), Amplifying reflective thinking in musical performance, in L. Candy, ed., 'Creativity and Cognition 2005', ACM Press, London, pp. 166–175.
- Johnston, A., Candy, L. and Edmonds, E. (2008), 'Designing and evaluating virtual musical instruments: facilitating conversational user interaction', *Design Studies* **29**(6), 556–571.
- Johnston, A., Candy, L. and Edmonds, E. (2009), Designing for conversational interaction, in 'Proceedings of New Interfaces for Musical Expression (NIME2009)' (in press).
- Johnston, A. and Edmonds, E. (2004), Creativity, music and computers: Guidelines for computer-based instrumental music support tools, in S. Elliot, M.-A. Williams, S. Williams and C. Pollard, eds, 'Australasian Conference of Information Systems', School of Information Systems, University of Tasmania, Hobart, Tasmania.
- Johnston, A. and Marks, B. (2007), 'Partial reflections', *Leonardo* **40**(5), 510–511.
- Johnston, A., Marks, B. and Candy, L. (2007), Sound controlled musical instruments based on physical models, in 'Proceedings of the 2007 International Computer Music Conference', pp. vol1: 232–239.
- Johnston, A., Marks, B., Candy, L. and Edmonds, E. (2006), Partial reflections: Interactive environments for musical exploration, in 'Engage: Interaction, Art and Audience Experience', Creativity and Cognition Studios Press, pp. 100–109.
- Johnston, A., Marks, B. and Edmonds, E. (2005a), An artistic approach to designing visualisations to aid instrumental music learning, in Kinshuk, G. Sampson, Demetrios and P. Isaias, eds, 'Cognition and Exploratory Learning in the Digital Age (CELDA2005)', IADIS Press, Porto, Portugal, pp. 175–182.
- Johnston, A., Marks, B. and Edmonds, E. (2005b), 'Spheres of influence' - an interactive musical work, in Y. Pisan, ed., 'Interactive Entertainment (IE2005)', Creativity and Cognition Studios Press, Sydney, Australia, pp. 97–103.
- Johnston, A., Marks, B. and Edmonds, E. (2006), Charmed circle: An interactive toy for musicians, in 'International Conference on Digital Interactive Media Entertainment & Arts'.