Designing and Evaluating Virtual Musical Instruments: Facilitating Conversational User Interaction

Andrew Johnston, Linda Candy, Ernest Edmonds

Creativity and Cognition Studios

Australasian CRC for Interaction Design (ACID)

Faculty of Information Technology

University of Technology Sydney

PO Box 123, Broadway, NSW 2007, Australia

aj@it.uts.edu.au, linda@lindacandy.com, ernest@ernestedmonds.com

Corresponding author: aj@it.uts.edu.au

Keywords

Creativity, evaluation, interface design, user behaviour, music

Abstract

This paper is concerned with the design of interactive virtual musical instruments. An interaction design strategy which uses on-screen objects that respond to user actions in physically realistic ways is described. This approach allows musicians to 'play' the virtual instruments using the sound of their familiar acoustic instruments. An investigation of user experience identified three modes of interaction that characterise the musicians' approach to the virtual instruments: instrumental, ornamental and conversational. When using the virtual instruments in instrumental mode, musicians prioritise detailed control; in ornamental mode, they surrender detailed control to the software and allow it to transform their sound; in conversational mode, the musicians allow the virtual instrument to 'talk back', helping to shape the musical direction of performance much as a human playing partner might. Finding a balance between controllability and complexity emerged as a key issue in facilitating 'conversational' interaction.

In this paper, we describe an approach to the design of interactive ‘virtual musical instruments’ which respond to live acoustic sounds produced by traditional musical instruments. Our technique uses simulated physical models to map between the live sounds produced by musicians and computer generated sound and visuals.

Our aim has been to explore how real-time audio-visual software might facilitate musical expression and exploration in live performance. Our approach has been to create interactive performance works in which a musician playing an acoustic instrument (in this case a trombone) interacts with computer software live to produce a mix of acoustic and computer generated sounds and associated visuals. While the software we discuss here was developed to accompany specific compositions, it is flexible enough so as to be usable in other musical contexts including improvised performances. For this reason the software itself can be seen as a kind of musical instrument. In this paper, the software is referred to as a 'virtual instrument' in order to distinguish it from 'acoustic instruments' played by musicians.

The virtual instruments were developed as part of an artistic collaboration between two musicians, both trombonists, with professional experience and tertiary qualifications in music. One of the musicians is a recognised composer who has received a number of commissions and grants. The other has programming design skills and is currently a lecturer in a faculty of information technology.

The main artistic outcomes to date are the two-movement work *Partial Reflections,* for solo trombone and virtual musical instruments and the interactive artwork *Spheres of Influence.* Designed for use by the general public rather than expert musicians, *Spheres of Influence* is a virtual musical instrument which uses a simplified interaction scheme.

In order to inform further refinement of the software design and improve our understanding of the nature of this particular type of interaction, a qualitative study was conducted which examined the experiences of expert musicians as they used the virtual instruments. From this we have identified three modes of interaction that characterise the musicians’ approach to the virtual instruments: instrumental, ornamental and conversational.

1. Initial design criteria

The initial goal was to create compositions and virtual instruments for live performance which encouraged musical expression and exploration. We identified a number of design criteria which guided the development of the virtual instruments.

- The virtual instruments should respond in real-time to live audio via a microphone and there should be an audio and visual response to the live input.
- There should be no additional buttons, mice or sensors attached to the acoustic instrument. This means that the software will work with any musical instrument, including vocals.
- The virtual instrument should be intuitively controllable by expert musicians.
- The sounds and visuals produced by the virtual instrument should be complex and engaging.
- The relationship between live sound, the behaviour of the virtual instrument and the resulting sounds should be apparent to musicians and observers (eg. audience members).

These align with well-known criteria in this domain such as those proposed by Wessel and Wright (2002) and Fels, et al (2002).

2. Simulated Physical Models

Early on in this project we decided to base the virtual instruments on simulated physical models as a way to provide complex audio-visual responses to live sound while retaining intuitive controllability.

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. 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 and amplitude 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 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. When we hear these sounds, we discern in them traces of the physical actions that caused them. 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. 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.

The implication is that instruments that facilitate a more direct connection between the physical actions of performers and generated sounds are more likely to produce sounds which people can engage with at this fundamental level. The musician using such an instrument can draw on their knowledge of how things work in the world to inform their interaction. 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.

In a seminal paper presented at the 1992 International Computer Music Conference, Joel Ryan argues that making musical ideas more concrete by using software models (including physical models) 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).

3. Physical models as a mapping layer

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. We argue that it is not necessary to use physical modelling to model the *source* of the sound in order to maintain these links. An interaction strategy that we developed during this project was to use physical models to mediate between the live audio generated by the musician and the audio-visuals generated by the virtual instrument.

Figure 1 shows a high-level view of how this works.

Insert Figure 1 here

The structure of the interaction is as follows:

- The musician plays into a microphone attached to the computer. This audio is passed through a software module (Puckette et al. 1998) which quantifies various features of the sound such as pitch and volume.
- The musical feature parameters are mapped to forces which act on a simulated physical model (a ‘mass-spring’ model). Because the model simulates 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 visual 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 trombone.

In contrast to the more traditional use of physical models in sound synthesis, the physical models that are used in this approach do not resonate at acoustic frequencies. Rather, they move at slower, “haptic” rates (up to around 20Hz). The movements of these models are used to control, rather than directly generate, the synthesized sounds created by the computer. This is a variation on a technique described by Momeni and Henry (2006).

We chose to use physical models in this way for a number of reasons. Firstly, we felt that the link between the live sound and the movement of the on-screen physical model was easily perceivable and would be intuitively understood by musicians and audience members. Secondly, because the timbre (or sound quality) of the generated audio is directly linked to the movement of the physical model, we felt that the sounds produced were interesting, complex and retained a physical character. In addition, because each mass is linked, directly or indirectly, to other masses in the structure, moving one mass (by playing its associated note) will cause other masses to move. By playing a number of carefully chosen notes in succession, the movement and resulting audio output can become quite complex. Thus, a simple physical structure and simple mapping rules can generate complex and sometimes surprising results without compromising high-level controllability.

Finally, when the structure of the physical model was carefully chosen, we felt that the 'physicality' of the interaction encouraged a playful approach that would engage musicians and encourage free-flowing exploration.

3.1 Example: Partial Reflections I

To illustrate the technique outlined above, we now describe the virtual instrument created for the first movement of *Partial Reflections*. This has been described in detail elsewhere (Johnston et al. 2005, Johnston et al. 2007). We provide a summary below to help clarify the way we have used physical models.

The virtual physical at the core of this instrument is a simple structure comprising 12 masses linked together by elastic 'springs' (figure 2).

Insert Figure 2 about here

Each mass in the model is associated with a particular pitch-class (or note) and, by playing into the microphone connected to the computer, the musician can exert forces on the model. If the player plays a G for example, then the mass at the bottom of the model is pushed with a force that is proportional to the volume of that note (i.e. the louder the note, the more force is applied). When no notes sound, masses in the model return to their resting positions, hanging down from the top of the screen. Movements of the masses cause audio to be output by the computer. The software stores the frequency of each note played by the musician and this frequency is associated with the appropriate mass. For example, if the musician plays an A with a frequency of 440Hz, then the A mass is assigned that frequency. If they subsequently play a lower A with a frequency of 220Hz then this replaces the value of 440Hz previously assigned to the A mass. As the masses move around, associated oscillators generate pitches (sine waves) at their assigned frequency. The faster they move, the louder their pitch sounds.

In order that the generated sounds are more interesting, the partials of the player's sound are treated the same way. That is, they are 'stored' by the masses. We use pitch-recognition software (Puckette et al. 1998) to identify the two strongest partials in the sound and these are associated with the appropriate mass. For example, if the sounded A has partials with a pitch-class of E and G, then the E and G masses will be associated with the frequencies of those partials. When those masses move sine waves at those frequencies will sound at a volume proportional to how fast they move. The effect of this is that by carefully selecting pitches and volumes the musician/composer can influence the timbre (or tone) of the computer output. The performer is effectively controlling the parameters for additive synthesis by manipulating the physical model with their sound.

Insert Figure 3 about here

4. User experience study

While pleased with the artistic results of our work, we wanted to further explore musicians' experience with the software. The questions we wanted to pose were:

- How do the musicians approach using sound controlled instruments such as these?
- What impact does using the software have on the musicians' music making?

With these broad questions in mind, a series of in-depth investigations of the experiences of expert musicians was carried out. The musicians played three different virtual instruments, each of which used a different physical model in the mapping layer to map between live sound and computer-generated sounds and visuals. The reader is referred to Johnston, et al (2007) for a fuller description of the workings of the virtual instruments. The instruments were:

- *Partial Reflections I* (PR1): the instrument used in the first movement of *Partial Reflections* (described above).

- *Partial Reflections II* (PR2): the instrument consists of 12 fast-moving spheres orbiting around a fixed central point. Only the very beginning of each note (the articulation, or 'attack') exerts force on the masses. The instrument records the attacks of notes played by the musician, and plays them back in a driving, rhythmical pulse controlled by the speed of the orbiting spheres. The spheres spin further out from the central point as they receive more force and as they spin further out the pitch of the replayed attacks is lowered.
- *Spheres of Influence* (SI): a virtual instrument that uses a simplified interaction scheme derived from that of PR1. The masses are arranged in a circle and move independently of the other spheres (i.e. they are not linked together). The masses 'store' the frequency of live notes in a similar manner to PR1. However, SI does not store or replay the frequency of partials.

'Traditional' Human Computer Interaction approaches have focused on measuring user performance when carrying out various well-defined tasks such as navigating a web site or entering figures into a spreadsheet. Software that is designed to facilitate musical expression presents a problem in this context as it is difficult to formulate tasks to assign to users that are measurable, but also meaningful (Wanderley and Orio 2002). If the aim had been to produce a general-purpose musical instrument for performing music in a well-established tradition, then this would be simpler. Tasks such as playing a scale, trilling, etc. could be assigned and measurements made to ascertain how successfully users are able to execute them. The benefit of this approach is that it would be possible to objectively compare two different virtual musical instruments in terms of their playability. However, where the instrument is intended to create new and unusual sounds, in effect, to explore new languages of composition and performance, this approach is problematic. Part of the rationale for creating these instruments is that they disrupt habitual ways of thinking about music so that musicians are stimulated to try new ways of playing and composing. Measuring how effective they are at facilitating performance of current styles of music might be interesting, but it would not necessarily help us learn more about designing to encourage divergent thinking.

As this was a new style of interactive music software, we were not sure how the musicians would structure their interactions or how they would conceive of the relationship between their playing and the computer-generated sounds and visuals as they attempted to make music.

4.1 Research Methods

A series of studies of musicians playing and commenting on the virtual instruments was carried out. The intention was that by observing and conversing with experienced professional musicians as they used the virtual instruments, we would gain insight into the impact that using the instruments had on their music making. It is important to stress that the musicians selected for the study were highly trained and at the top of their profession, including principal players from major professional orchestras and leading improvisers.

Seven musicians participated in the study with each using the virtual instruments for approximately two hours in total. Because the software was designed to work with instruments that are predominantly monophonic (such as the trombone), the musicians were all wind or brass players (trumpet, trombone and clarinet). We deliberately chose musicians who had an interest in contemporary music, especially those who also compose and/or improvise, and who had more than 10 years professional experience in a range of musical contexts.

Our approach, drawing on that described by Suchman and Trigg (1991), was to give the musicians freedom to use the software in any way they wished and to make music with it in order to explore its potential. We used the concurrent think-aloud technique (Ericsson and Simon 1993) in order to gain insight into their experience: that is, we asked the musicians to 'think aloud' as they interacted with the software. Because brass and woodwind players are obviously unable to speak and play their instrument at the same time and, because we did not wish to interrupt the flow of performance, we did not ask musicians to interrupt their music-making to make comments. Instead we simply asked them to verbally report what they were thinking and perceiving as frequently as they were able during their time playing the instruments. This meant that they played for some time, commented on what was happening, played some more, made further comments and so on.

We then asked a series of open questions and asked them to complete a short questionnaire. In addition to the first author and the musician, an observer attended each session and took notes to provide an additional perspective. Each session was video recorded to facilitate full data capture and further analysis.

4.2 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. We chose to use techniques from grounded theory (Glaser and Strauss 1967, Glaser 1978) to code and analyse the data gathered. Grounded theory was a good fit for our needs because it facilitates the generation of theory closely tied to the evidence from rich qualitative data. At a high level, the basic steps of the doing grounded theory analysis process as we applied it in this study are:

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.

5. Findings

5.1 Modes of Interaction

The musicians in our study demonstrated three modes of interaction. At different times they approached the virtual instruments as:

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

The musicians were observed and described, moving between these modes as they played with a particular instrument. Thus, the modes 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 4). As such, a musician may for example begin in 'instrumental' mode, move to 'ornamental' mode for a time, and then eventually end up in a 'conversational' interaction.

Insert Figure 4 about here

We will now describe the three interaction styles, illustrating each with examples from the user experience sessions.

5.2 Instrumental interaction

When initially designing the virtual instruments, we anticipated that this 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 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. Key words 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 is, 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.

"I just feel that this [Spheres of Influence] actually feels much more like an instrument. You know, it's not telling you what to do. You're affecting it." (Musician K)

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 the potential of carefully designed virtual instruments of this type as tools to develop technical and aural skills in traditional music education.

"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." (Musician P)

5.3 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 the musicians' sound. 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.

"That [the rhythmic pulse generated by Partial Reflections II] has quite a dictatorial effect…" "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." (Musician D)

Other musicians found the pre-determined nature of the ornament overly dominating.

"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." (Musician J)

Partial Reflections II was the instrument most often approached as an ornament by the musicians. While some of the musicians liked the effect and saw it as a potential 'aid' for establishing a connection with the audience, 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 PR2's design which seem most responsible for encouraging this ornamental approach. First, the force which animates the *Partial Reflections II* virtual 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. Second, 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 was 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 comes at the cost of reduced intimacy and controllability.

5.4 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.

"[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." (Musician J)

A key characteristic of this mode of engagement is that musicians are moving beyond simply controlling the virtual instrument and are allowing it to 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'. This implies that the response was not surprising as such, but rather that the sounds and visuals had 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 ‘conversation partner’ mode of use 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, reflecting and transforming the sonic input in ways that move the performance in new musical directions.

“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].” (Musician P)

6. Implications for Interaction Design

The issue of control is the single biggest factor that differentiates the three modes of interaction. In purely instrumental mode, the musician is aiming for complete control over every aspect of the virtual instrument. In purely ornamental mode, the musician surrenders control to the virtual instrument as it transforms their sound in ways that they affect only at a high level. Conversational mode involves sharing of control, a shifting balance of power between musician and virtual instrument. At times the musician will lead and direct the virtual instrument, moving towards an instrumental approach, and at other times will give the virtual instrument more autonomy, allowing it to suggest new musical directions.

Whilst the conversational mode of interaction is the most interesting, it is also the most difficult to design for. To support this mode, the virtual instruments need to retain controllability while simultaneously being able to introduce new, occasionally surprising, musical material. It was interesting that the simple physical models that we used in our virtual instruments (particularly *Partial Reflections I*) were able to provide this balance in many cases.

At first glance it might seem that software intended to facilitate conversational interaction with humans might need to draw on techniques of artificial intelligence in order to simulate the reactions of a human musician. Our experience, and the feedback from the musicians who participated in our study, indicates that using physical models as a mapping layer between live sound and computer generated sound and visuals, is a mechanism that can provide the balance between controllability and complexity necessary to stimulate conversational interaction. This, of course, is not to say that artificial intelligence techniques are inferior or that they do not show potential, merely that another approach can also facilitate a type of conversational interaction.

These findings have interesting implications for interaction designers, especially those interested in designing to support or encourage creative work. The first is that while controllability is important, sacrificing some controllability in order to encourage a more conversational interaction can stimulate new ideas and encourage divergent thinking. Simple dynamic systems, such as simulated mass-spring models, can help by providing high-level predictability and controllability, coupled with behaviours which are complex, rich and sometimes surprising.

Software such as Paul Haeberli's *Dynadraw* is an interesting example of using such an approach:

> "The program Dynadraw implements a dynamic drawing technique that applies a simple filter to mouse positions. Here the brush is modeled as a physical object with mass, velocity and friction. The mouse pulls on the brush with a synthetic rubber band. By changing the amount of friction and mass, various kinds of strokes can be made. This kind of dynamic filtering makes it easy to create smooth, consistent calligraphic strokes." (Haeberli 1989)

By adding a simple mass-spring model to a traditional drawing program, Haeberli enabled a more complex, conversational interaction. The user moves the mouse to draw on-screen as with any standard 'paint' program. The difference is that faster mouse movements cause the 'rubber band' attached to the on screen brush to stretch. The interplay between mouse movements and simple mass-spring systems results in drawings which are interesting and complex, while retaining a clear link to user gestures.

An additional difficulty for interaction designers designing for conversational interaction is the challenge of understanding the complex nature of interaction as it relates to creative work. Evaluating interaction designs in this context is difficult. Users can tell us whether they like playing with particular virtual instruments or not, but this, while helpful, does not necessarily tell us how to design better ones. On the other hand qualitative feedback can swamp us with detailed personal preferences. We have found that engaging in loosely structured dialog with expert creative users is effective in building understanding of the sometimes complex ways in which they interact with software while engaged in creative work. Grounded Theory techniques were useful because they helped ensure the voices of the musicians were heard and aided the discovery of patterns in the way they approached the virtual instruments.

7. Conclusions

In this paper, we have described the design of a set of virtual musical instruments that are based on physical models and a study of their use by expert musicians. The potential of physical models in this way in the music domain has been demonstrated and has promise in other areas for giving users intuitive control over complex systems, while simultaneously providing potential for a two-way conversational mode of interaction between user and system.

Interaction design for creative uses and users is challenging for a number of reasons, not least because it is difficult to unambiguously measure how successful or unsuccessful various strategies are. Our approach of giving expert musicians freedom to play with the virtual instruments and asking them to verbalise and reflect on their experiences, has provided us with a number of insights into what impact various aspects of our interaction design have had on their music making: namely, that they interact with the virtual instruments in three modes, instrumental, ornamental and conversational, which relate primarily to the balance of power between musician and software.

The creative work and research findings we describe here are a starting point for exploring how musicians interact with software instruments such as this. As we become more experienced at designing virtual instruments, we expect our understanding of the ways in which musicians experience and engage with them will become more sophisticated. Our hope is that our work to date will help by increasing the visibility of some of these issues and giving interaction designers and musicians a language for talking about them.

Acknowledgements

The composer/musician Ben Marks was co-creator of *Partial Reflections* and *Spheres of Influence*, composing the music and making a significant artistic contribution to the development of the software. Our thanks to the musicians who participated in our user experience study for their time and insightful comments. Our thanks also to the developers of Pure Data, Physical Modelling for Pure Data and Graphical Environment for Multimedia for creating and making available the software which made this work possible.

This research was partly conducted within the Australasian CRC for Interaction Design, which is established and supported under the Australian Government's Cooperative Research Centres Program. An Australian Postgraduate Award provided additional financial support.

References

Cadoz, C.; Luciani, A. & 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.

Ericsson, K. A. & Simon, H. A. (1993), *Protocol Analysis: Verbal Reports as Data*, MIT Press, Cambridge, MA.

Fels, S.; Gadd, A. & Mulder, A. (2002), 'Mapping transparency through metaphor: towards more expressive musical instruments', Organised Sound 7(2), 109-126.

Glaser, B. G. (1978), *Theoretical Sensitivity*, The Sociology Press.

Glaser, B. G. & Strauss, A. L. (1967), *The discovery of grounded theory: strategies for qualitative research*, Aldine de Gruyter, New York.

Haeberli, P. (1989), 'Dynadraw'. http://www.graficaobscura.com/dyna/index.html

Johnston, A.; Marks, B. & Edmonds, E. (2005), 'Spheres of Influence' - An Interactive Musical Work, *in* Yusuf Pisan, ed.,'Interactive Entertainment (IE2005)', Creativity and Cognition Studios Press, Sydney, Australia, pp. 97-103.

Johnston, A.; Marks, B. & Candy, L. (2007), Sound Controlled Musical Instruments Based On Physical Models, *in* 'Proceedings of the 2007 International Computer Music Conference', pp. vol1:

232-239.

Leman, M. (2007), *Embodied Music Cognition and Mediation Technology*, The MIT Press.

Puckette, M. S.; Apel, T. & 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.

Momeni, A. & Henry, C. (2006), 'Dynamic Independent Mapping Layers for Concurrent Control of Audio and Video Synthesis', Computer Music Journal 30(1), 49-66.

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.

Suchman, L. A. & Trigg, R. H. (1991), Understanding Practice: Video as a Medium for Reflection and Design, *in* Joan Greenbaum & Morten Kyng, ed.,'Design at Work: Cooperative Design of Computer Systems', Lawrence Erlbaum Associates, Hillsdale, NJ.

Wanderley, M. M. & Orio, N. (2002), 'Evaluation of Input Devices for Musical Expression: Borrowing Tools from HCI', *Computer Music Journal* **26**(3), 62-76.

Wessel, D. & Wright, M. (2002), 'Problems and Prospects for Intimate Musical Control of Computers', *Computer Music Journal* **26**(3), 11-22.

Woods, D. & Fassnacht, C. (2007), 'Transana v2.21', *Madison, WI: The Board of Regents of the University of Wisconsin.*

All Figures as follows:

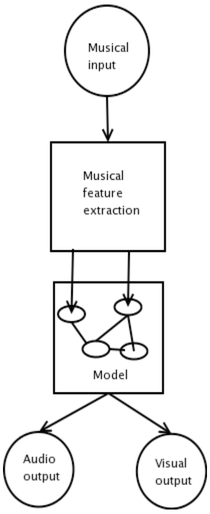

Figure 1: Block diagram showing the use of a physical model to map between musical input and audio/visual output.

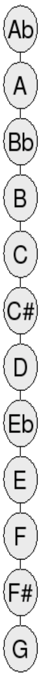

Figure 2: Mass-spring physical model for *Partial Reflections I*

Figure 3: Trombonist Ben Marks plays with *Partial Reflections I. (Photo: Ros Hodgekiss*)

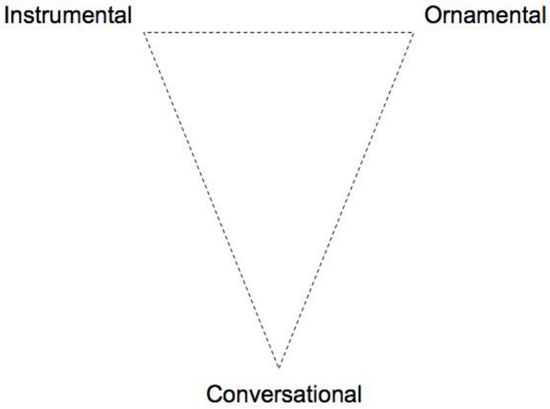

Figure 4: Three modes of interaction mark boundary points of a map of interactions with a virtual instrument.