AN INTEGRATED APPROACH TO COMPUTER-BASED MUSICAL INSTRUMENT DESIGN

Niall Moody
University of Glasgow
Centre for Music Technology

ABSTRACT

In recent years a large number of new computer-based musical instruments have been created. While a number of these have implemented a degree of non-audio feedback as a performance aid (which generally takes the form of either haptics or visuals), the majority have tended to focus specifically (and perhaps understandably) on the audio output, and have neglected other forms of feedback. This paper will examine recent developments in computer-based instruments and propose that an 'integrated' approach, which combines audio, visual and haptic feedback, may be more beneficial in the construction of new musical instruments. Further, a preliminary design of such an instrument will be stated based on this approach.
Keywords – instrument, audiovisual, haptics, physical modelling

1. INTRODUCTION

Traditionally, computer-based musical instruments have tended to focus on the instrument's audio output, with any form of alternative feedback, be it haptic or visual, regarded as an afterthought. While this is understandable, given the often highly processor-intensive operations required by audio synthesis algorithms, it risks ignoring the importance of secondary senses such as sight and touch as aids to the performer. In recent years, a number of projects have attempted to redress this imbalance to a certain degree by making use of haptic feedback, or creating visual interfaces with which to play the instrument. The tendency remains, however, to limit the secondary feedback to a single dimension, or sense. This paper will first examine the more recent developments in computer-based instruments, and propose that a more 'integrated' approach, combining audio, visuals, and haptic feedback, may be particularly beneficial in the construction of new musical instruments.

2. A BRIEF EXAMINATION OF EXISTING INSTRUMENTS

Before continuing, it is perhaps necessary to define what is meant by the term 'instrument' for this paper. Essentially, the term is intended to refer to systems which were designed specifically as instruments, as opposed to a piece of software like Max/MSP, which is far more general in scope, and regarded (for our purposes) as a music programming language.

2.1. Instruments With a Haptic Focus

One of the most relevant projects with respect to this paper is the Cymatic project [5]. This project combines a physical modelling synthesis engine with a haptic interface (a joystick and a mouse, both force-feedback equipped), and the authors make a number of interesting points in the papers about it released to date. In the aforementioned paper, the authors state that the reason physical modelling was chosen as the audio synthesis method was that it counters the claim made by musicians that computer-based instruments tend to sound "'cold' or 'lifeless'" [1]. The idea being that, because these musicians are used to playing physical instruments, physical modelling techniques provide a more familiar sound and timbre. The haptic interface is intended to fulfil a similar function, and bring the musician closer to the instrument by reproducing haptic sensations which the musician will be familiar with, and to a certain degree, expect from the instrument. The interface of a joystick and a mouse could be considered problematic, however, as it is debatable whether these devices are the best choice for an interface to a musical instrument. In the paper, 'StickMusic: Using haptic feedback with a phase vocoder' [6], for example, which uses the same two devices to control a phase vocoder, the author notes that using devices for which the performer has already developed gestures relating to their primary function results in the performer often 'falling back' on these gestures. With the mouse, for instance, typical actions involve small movements (e.g. clicking-and-dragging), which do not necessarily translate into musically interesting gestures.

An intriguing approach to interface design was taken in the 'Pebblebox and Crumblebag' project [4]. Here, two interfaces were designed to control a granular synthesis audio engine, one essentially being a box of pebbles which could be manipulated (knocked into each other, etc.) by the user to control parameters of the synthesis system, the other being a bag filled with a granular material (e.g. cornflakes, polystyrene beads), again manipulated by the user

Email: niallmoody@yahoo.co.uk

[1] [5], p.1

to control granular synthesis parameters. What is interesting about these instruments is that they do not use any kind of 'active' haptic feedback - there is no haptic signal output from the computer used to drive motors, for example. Instead the haptic experience is entirely derived from the physical construction of the interfaces, which is intuitively linked to the method of sound generation (in that the physical properties of the interfaces would typically produce a 'grainy' sound when manipulated). The idea of an 'internal haptic feedback' is also used in the 'Sillytone Squish Factory' instrument [3], where the interface is a molded rubber shape embedded with pressure sensors. To play the instrument, the performer squeezes and manipulates parts of the shape. The rubber itself provides a certain resistance to these actions, and always returns to it's original shape. What these two projects demonstrate is that haptic feedback in a computer-based instrument does not necessarily have to make use of complex electronic devices and electrical signals to function successfully. If the interface is designed well it should be possible for the haptic feedback (which is essentially providing the performer with some form of information about the instrument's state) to be built into the interface in a way that allows the performer to make an intuitive connection between the instrument and the interface. The author will return to the idea of an 'intuitive connection' later on, but essentially what is meant is that the haptic feedback from the interface should correspond to the sensations and forces that the musician would naturally expect to be the result of their actions.

2.2. Instruments With a Visual Focus

Musical instruments with a specific visual element are comparatively thin on the ground with respect to those which incorporate a haptic focus. Of those, Golan Levin's Audiovisual Environment Suite (and particularly it's accompanying thesis [2]) is perhaps the most significant resource currently available with respect to instruments which combine audio and visuals. The suite consists of five programs which all make use of so-called 'painterly', visual interfaces. One of the primary points Levin makes in his thesis is his belief that audiovisual instruments should permit the user/performer complete control over the placing of visual 'elements', or material, on the 'canvas'. This point, in keeping with Levin's proposal of a 'painterly' interface, tends to tip the balance of the instruments (certainly with regard to the interface) more towards the visual, than the audio realm. Indeed, the interface used for the instruments often reflects this slight prejudice, with a number of the instruments designed to be controlled from a wacom tablet. In his thesis, Levin explains this tendency as being necessary in order to provide the same degree of control over the spatial placement of visual material on the visual canvas as is available for audio material on the aural 'canvas'. He goes on to make the interesting claim that without this particular degree of visual control, the instrument would "seem more like a musical instrument accompanied by a sound-responsive graphic, than a true audiovisual construction system in which the visual dimensions are as malleable as the sonic ones" [2]. This presents an interesting challenge in the design of musical instruments which have some kind of visual dimension - do we go for an interface which allows for a particularly specific level of visual control, and risk losing out on the benefits of a more musically-oriented controller, or focus on creating a *musical instrument*, and risk the visuals becoming subordinated to the instrument's musical needs? In particular, the use of a wacom tablet, or similar stylus-based interface, presents certain problems (related to those of using a mouse as mentioned earlier) when used as the interface of a musical instrument, in that it brings with it a number of prior gestural associations, which may or may not be useful in the control of a musical instrument.

3. DISCUSSION

With regard to Levin's claim that without such a specific degree of control over the spatial aspect of the visuals, the visuals will become subordinated to the audio, it would seem (from Levin's thesis, anyway) that this is essentially an either/or case. Further, as this paper is focused on the design of a musical instrument, it follows that the only option is for the visuals of the instrument to merely present a visualisation of the instrument's state. The author would suggest, however, that visualising the instrument's state can, in itself, be a particularly valuable tool for the musician, and that this description, of 'merely' visualising the instrument's state, perhaps understates the value of having such a visualisation. Ultimately, one of the primary aims of this paper is to present an approach to instrument design that will try to bring the musician closer to the sound generation mechanism of the instrument. One of the aspects which most distinguishes computer-based instruments from physical, acoustic ones, is the fact that with a computer-based instrument, the sound generation mechanism is separate from the interface with which the musician controls it. Computer-based instruments are always limited by the fact that the sound generation mechanism is essentially an abstract process, occurring within a 'black box', which is entirely separate from the interface used to perform with it. In contrast, the interface of a physical instrument is often tightly connected to the (physical) sound generation mechanism (e.g. the string of a violin is both an integral part of the sound generation mechanism, and the main part of the instrument with which the musician interacts). There is significant evidence that the direct physical or haptic feedback this connection provides is a particularly important input for the musician [3], in that it provides data concerning the mechanical, as opposed to auditory, state of the instrument. This can then be used by the musician to determine the potential results of their actions, which will clearly often be affected by the mechanical state of the instrument. While it may never be possible with a computer-based instrument to approach the degree of integration between in-

[2] [2], p.106

[3] [1], chapters 18 and 19 (p.229-260)

strument and musician possible with physical instruments, it is the author's belief that it should be possible to at least minimise the disadvantages inherent in computer-based instruments from the way they separate the musician from the sound generation mechanism. By providing the musician with visual and haptic feedback of the instrument's state, the author aims to present the instrument as an actual object with it's own (observable) properties and behaviours, rather than as a collection of abstract, sound generating algorithms.

Throughout this paper the issue of the interface, or controller, used to interact with the instrument has been key. It is hopefully apparent from the brief examination of existing instruments previously that using general purpose controllers such as joysticks and mice in the control of musical instruments presents certain problems. In using general purpose controllers, the instrument runs the risk of the musician falling back on gestures previously learned with those controllers. Taking the mouse (again) as the most pronounced example of this, the gestures typically performed with it (i.e. in it's intended use) are small, and mostly discrete. This clearly presents problems for instruments which require or benefit from large gestures, as the musician will have to break free of their prior, conditioned responses to the device. The author would therefore like to propose that, in designing a new computer-based musical instrument, it is vital that the interface or controller is designed specifically to suit the instrument. The range of gestures it affords should be intimately related to the instrument's method of sound generation. Prior gestural knowledge or experience is not necessarily a disadvantage (indeed, it may aid the musician in learning the instrument), but combined with general purpose devices such as those discussed previously, there is the possibility that it will tend to work against the musician, and end up prematurely limiting the gestures they perform with the device.

In designing an interface which is intimately related to the instrument, we come up against the problem of mapping gestures to synthesis parameters. The problem is far from trivial, but the author would like to propose a solution that will hopefully side-step some of the more difficult issues involved. Rather than attempt to determine what kinds of gestures naturally relate to particular timbres or sonic characteristics, perhaps a simpler solution is to make use of gestures and interactions familiar to all musicians (indeed, to all human beings) from their experience of the physical world. All physical objects have certain properties and associated actions in common - they have a certain mass and texture, they can be picked up, moved around, they can be stretched and squashed to a greater or lesser degree, and so on. From experience therefore, it is common to be able to look at an object and instantly determine what kind of actions are possible with it (and in particular, what kind of sound may be emitted as a result of such actions). It is logical to conclude that co-opting these physical properties may prove to be particularly valuable in the design of a computer-based musical instrument. In a sense, this is the reasoning behind the use of physical modelling techniques with respect to sound generation. The author would like to suggest, however, that for a physically-modelled instrument (referring to a 'new' instrument, not a model based on an existing physical instrument), this approach should be extended to the interface. Although this is not necessarily possible with other synthesis techniques, physical modelling provides the scope necessary for the instrument builder to define a particular, musical model, and then extrapolate that model out to the interface. As the sound generation model relies on physical properties and processes to produce it's output, the model will conform to the same physical processes familiar to the musician from their experience in the physical world. By tapping into this familiarity, there is an opportunity to side-step some of the more complicated issues involved in mapping gestures to the instrument's audio generation mechanisms, in that the interface just needs to be able to replicate the kind of gestures which would be available for the physical model were it to exist as an object in the physical world.

4. PRELIMINARY INSTRUMENT DESIGN

It is hopefully clear from the preceding sections that the author favours physical modelling synthesis as the basis for an 'integrated' instrument. The design the author is currently working on envisages the instrument as essentially a musical block of clay - an amorphous blob, modelled in the computer as a network of mass-spring nodes. To play it, the musician will interact with it in the same way they would a block of clay - by stretching it, squashing it, twisting it, and making indentations, striking it to excite the model - and it will respond in turn, with it's shape changing according to the state of the physical model. Ideally, the interface to such an instrument would simply be a block of clay which would dynamically change it's shape according to the state of the physical model (for example pulsating and vibrating as the model reacts to the musician exciting it). According to the principles outlined previously, of the controller being intimately connected to the sound generation model, such a controller would be perfect, because it's physical construction is essentially the same as that of the synthetic physical model at the heart of the instrument. To the best of the author's knowledge, however, such a material does not exist, so an alternative controller has to be designed. In doing so, the aim (as stated previously) is for the controller to be able to replicate all the gestures which would be possible with a musical block of clay, in a way that makes it possible for the musician to quickly grasp how their gestures with the controller relate to the sound output from the instrument.

The author's current design can be seen in the three pictures following.

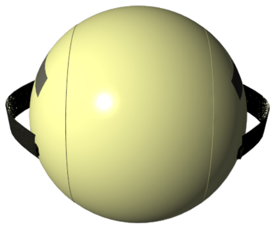

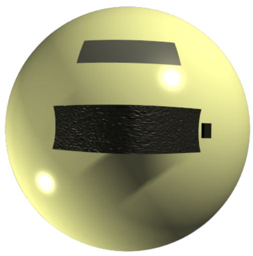

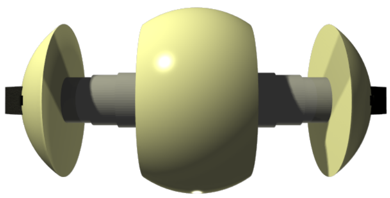

While this design does not look much like a block of clay, it should nevertheless be capable of the majority of the gestures required of it. The two hemispherical sections may be rotated and extended to twist and stretch the model, respectively. The grey sections are intended to be pressure (and position) sensitive pads, for the musician to shape the model by applying (slow) pressure, as well as excite it by quickly tapping them. Ideally the entire surface of the controller would be sensitive to pressure, but this design should prove adequate for this instrument, considering the musician's hands will tend to remain in the same position (due to the hemispherical sections). The straps on the hemispherical sections are intended to allow for large, sweeping gestures (one could imagine a 'bowling' kind of action, where the musician swings the controller (with one hand) in a downward-forward arc, with the hemispherical sections extending as a result of the momentum). With respect to the haptic feedback issues discussed earlier, the controller will make use of an internal haptic feedback related to the physical model. The hemispherical sections will be sprung in some manner, so that they always return to an equilibrium position, and it requires progressively more force to stretch them apart. In addition, the interface may incorporate some kind of active, vibrational haptic feedback to provide further information about the instrument's current ('mechanical') state.

With respect to the visual side of the instrument, this will take the form of a 3d graphical representation of the internal physical model, existing within a 3d graphical environment with a number of other graphical (as well as audio) objects with which the instrument may interact. The idea is that the instrument should not exist entirely on it's own, but within a complex environment where it may be joined by other instruments, all of which will be positioned (aurally) using Ambisonics to allow for any number of speaker configurations. To move the instrument through this 3d space, the position of the controller will be tracked (possibly using colour tracking), allowing the musician to place the sound in the synthetic environment according to the controller's position in their physical environment.

5. REFERENCES

[1] Perry R. Cook, editor. *Music, cognition, and computerized sound : an introduction to psychoacoustics*. MIT Press, 1999.

[2] Golan Levin. *Painterly Interfaces for AudioVisual Performance*. PhD thesis, Massachusetts Institute of Technology, September 2000.

[3] Geoffrey C. Morris, Sasha Leitman, and Marina Kassianidou. Sillytone squish factory. In *New Interfaces for Musical Expression Conference*, 2004.

[4] Sile O'Modhrain and Georg Essl. Pebblebox and crumblebag: Tactile interfaces for granular synthesis. In *New Interfaces for Musical Expression Conference*, 2004.

[5] Stuart Rimell, David M Howard, Andy M Tyrrell, Ross Kirk, and Andy Hunt. Cymatic. restoring the physical manifestation of digital sound using haptic interfaces to control a new computer based musical instrument. In *International Computer Music Conference*, September 2002.

[6] Hans-Christoph Steiner. Stickmusic: Using haptic feedback with a phase vocoder. In *New Interfaces for Musical Expression Conference*, 2004.