PETECUBE: a Multimodal Feedback Interface

Peter Bennett
University of Brighton
Faculty of Arts and Architecture
Grand Parade, Brighton, BN2 0JY
peterdavidbennett@gmail.com

ABSTRACT
The PETECUBE project consists of a series of musical interfaces designed to explore multi-modal feedback. This paper will briefly describe the definition of multimodal feedback, the aim of the project, the development of the first PETECUBE and proposed further work.

Keywords
Multi-modal Feedback. Haptics. Musical Instrument.

1. INTRODUCTION
A multimodal system can be defined as a system that "supports communication with the user through different modalities such as voice, gesture and typing." [1]. The 'mode' term of multimodality can be used to refer to both mode and modality. Mode refers to "a state that determines the way information is interpreted to extract or convey meaning" [1] whereas modality refers to "the type of communication channel used to convey or acquire information."[1]. 'Feedback' can be defined as "the return of part of the output of an electronic circuit, device, or mechanical system to its input, so modifying its characteristics" [2]. Hence, a multimodal feedback interface can be defined as an interface with multiple communication channels that returns a portion of its output to the input of the system. The output of the system in the case of an instrument is the sound produced, and the input can be seen as the user playing the instrument.

Many instruments have been developed that use various forms of feedback, however it is felt by the author that the instruments are normally biased towards one of the particular senses and that other sensory feedback is somewhat neglected. The aim of this project is to create a series of instruments in which all forms of feedback are equally considered, and more importantly are used together in a coherent whole. Of the five Aristotelian senses (sight, hearing, touch, smell and taste), it has been decided to concentrate upon the three that are most pertinent to playing a musical instrument; sight, hearing and touch. All musical instruments already incorporate *passive* feedback of all of these senses (i.e. you can see, hear and feel a piano or guitar). However, the interest of this project is in *active* feedback, so that the designer of the instrument can specify how an instrument will react within each of those modalities.

Successful research that explores this area is the PHASE project [3]. The PHASE group have implemented a multimodal installation that offers haptic, visual and audio feedback operating on a model of a turntable like device with both a 'writing' and a 'playing' head. The PETECUBE project differs from this in several key ways. Firstly, the PETECUBE aims to embody the multimodal feedback within a single object. Secondly, the PETECUBE is designed for live performance, not an installation, so the size and complexity of the setup is restricted by the need for portability. Thirdly, and most importantly, the PETECUBE is designed as part of a series in which each cube is limited in its modalities, and as such, individual cubes should not be considered complete instruments, but as investigations into particular combinations of sensory feedback.

NIME 06, June 4-8, 2006, Paris, France.

2. DESIGN
To impose some limits on the design of the interface, it has been decided to limit the physical form to the shape of a cube (see figure 1). Although arbitrarily chosen, the cube was found to be a useful design; it is an ideal shape on which to mount sensors and actuators, a 2D representation of a cube is easily seen as a 3D cube (figure 3), it is a robust shape ideal for rough handling and it is easily grasped by the hand (figure 4). Another consideration is that a cube is not an imitation of a conventional instrument, so that users should approach it without any preconceptions on how to play it.

Figure 1. A Prototype PETECUBE.

The system diagram below (figure 2) shows how the PETECUBE is organised. The three levels depict the user interface level at the top, the hardware level in the middle, and the software level at the bottom. At the top level the user can manipulate the PETECUBE whilst also receiving three forms of sensory feedback; vibration from the cube, sound from the speakers and visualization from a monitor or projection. The middle level consists of hardware to communicate between the user interface level and the software level. The bottom level consists of three separate software programs to handle each of the feedback modalities.

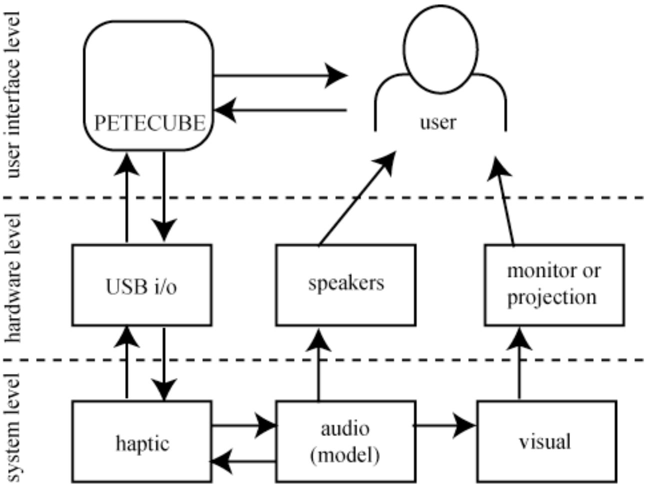

Figure 2. System Diagram.

It was decided to break up the software level in this manner to create a modular system. Communication between the programs is achieved using Open Sound Control [4] over a network. There are two main benefits to this approach. Firstly, the programs can be created in different languages depending on the requirements, and secondly the programs can be run on separate computers to avoid slowing down the response of the feedback.

2.1 Haptic Design

The role of haptic feedback is to allow the player to feel the result of their playing. An example of this in a traditional instrument would be feeling the skin of a drum vibrating after it has been struck. Benefits of incorporating haptic feedback in an interface include improving the players accuracy [5], and allowing the user to have less reliance on visual and audio feedback.

Haptic feedback can be achieved most simply with vibration, such as using a motor with an off-balance weight attached. This is the method employed in many games controllers and mobile phones. Although simplistic, the vibration motor has been deemed a good starting point before exploring more advanced methods such as force-feedback.

2.2 Audio Design

The problem with designing an electronic instrument is that any sound imaginable can be potentially used, thus giving an overwhelming choice to the designer. One way to simplify the problem is to divide the generation method of electronic sounds into sampled and synthesised. Sampled sounds allow the user to play any sound they like, thus making the instrument more versatile, whereas synthesized sound has the potential for greater expressivity but a more limited range of sounds. It has been decided to use sampled sounds to start with, so that the cubes can be used in a variety of musical situations. However, when wider varieties of PETECUBE's have been produced, investigation into synthesized sound, especially physical modeling synthesis, will be made.

The audio module is important because it acts as the central model of the system. The model holds the instruments current state, which is continuously updated from the output of the PETECUBE's sensors. In the case of using sampled sounds, the parameter being updated is the position of the tape-head within each sample. When the audio model has been updated, the state is then translated into the different output modalities to be fed-back to the user.

2.3 Visual Design

The visual design of traditional non-electronic instruments is generally directed by the mechanical constraints presented when building the instrument, without these constraints it is problematic in deciding how to visualize the virtual instrument. To narrow the design possibilities it has been decided to represent the virtual cube in a relatively realistic manner, so that it is easy to see the link between the virtual cube and the real cube. This virtual cube can then be visually augmented in a manner that would be impossible with a real cube. This augmentation currently takes the form of sound samples being projected perpendicularly from the six faces of the cube, so that the user can see the sample that is being played.

A particular importance of visual feedback is not only in informing the user, but also in displaying the instrument to an audience. Ideally, the visual depiction is clear enough so that the audience can gather what is going on, but at the same time dynamic and exciting enough so that they don't lose interest.

An addition to the visualization of the instrument is anaglyphic 3D-glasses. This allows users (and the audience) to see the virtual cube as three-dimensional. Using anaglyphic 3D glasses is just a temporary stage though, as ultimately the visualization should be located on the cube using augmented reality techniques.

3. PROTOTYPE

A fully functional prototype has been made, as outlined below.

3.1 Prototype Hardware

The prototype uses six light-dependent resistors (LDR's) to sense the users movement. To optimise their sensitivity, LDR's on opposite faces are linked together in a half-bridge, so that the signal generated is the difference between the two sensors readings, rather than the absolute value from each sensor. This has the advantage of negating the ambient light conditions, allowing the cube to be used in nearly all lighting conditions. A less obvious advantage is that by arranging the sensors in the half-bridge, the six sensor outputs are reduced to three, using less ports of the USB i/o.

The i/o hardware is the National Instruments USB-6008. This was chosen because it offers 8 analog inputs, 2 analog outputs and 12 assignable digital i/o lines, potentially allowing two PETECUBE's to be run simultaneously. Another advantage is the scalability when using its device independent C++ library, as a device with more i/o lines, or a higher sampling rate could be used at a later date, with minimal change in the code.

To provide the vibration, two motors with unbalanced loads were appropriated from a Playstation dual-shock controller. Because the loads have different weights, varying intensities of vibration can be achieved. Both motors are placed in the centre of the PETECUBE and secured firmly.

3.2 Software Design

The software is broken up into the three modules of haptics, audio and visualisation as outlined above. The three modules use Open Sound Control [4] to communicate over a network connection, allowing the flexibility of running the programs on separate computers if needed. The use of this is not only to spread the processing load, but can be used in a performance where one laptop could be positioned on stage connected to the USB i/o whilst a second laptop could be positioned at the back of the room and plugged into the mixing desk and projector. In this situation, a wireless network can be used to avoid the use of long cables.

The haptic module's main task is to connect with the USB i/o and to map incoming OSC to voltages out, and incoming voltages to OSC out. It also scales the data, so that the OSC messages are kept within a universal range. The current method of generating the haptic output is to map the amplitude of the sound to the amplitude of the vibration. Although crude, this gives a relatively coherent experience. Currently, both motors are used in the same manner, although improvements are to be made so that the two different intensities of motor are used to greater effect.

The audio module has the central task of not only generating sound, but also to send information on the audio models current state to the haptic and visualisation modules. Currently the audio module is being prototyped in Max/MSP [6], although it is planned to use the Synthesis Tool Kit [7] in C++ for further work. The sound model currently used is a simple one-second long sound file (selected by the user) that can be scrubbed back and forth by the input from the cube. As there is a sample on each side of the cube, six samples can be loaded at any one time. As the playback position within each sample is changed, a ramp is generated between the old and new position to ensure a relatively smooth transition. Because this ramp time is fixed, it becomes possible to control the speed (hence pitch) of the sample by scrubbing faster or slower.

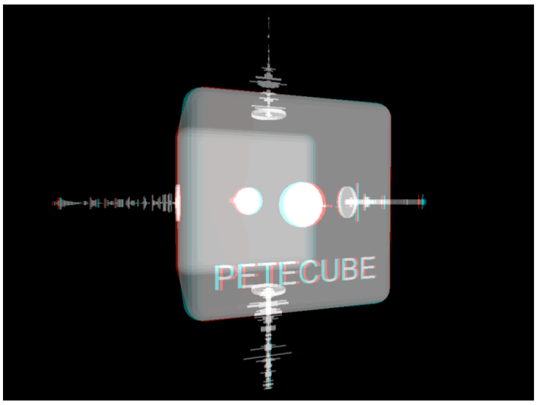

Figure 3. Functioning Prototype Visualisation.

The visualisation module is written in C++ using OpenGL [8]. An anaglyphic library [9] is used to display the model in three dimensions, suitable to be seen with red/cyan glasses (see figure 3 above). The model is designed in a 3D modeling package and then imported into the visualisation module so that an accurate representation is used. To display the waveforms that protrude from the surfaces of the cube it was decided to use a dynamically updating approach. This avoids having to send the entire waveform from the audio module to the visualisation module. To achieve this it is necessary for the audio module to send not only the current position of the 'tape-head' in the sample but also the amplitude of the wave at this point. This allows the visualisation module to build up an image of the waveform after a couple of sweeps of the 'tape-head'. This has two distinct advantages. Firstly, because only the position and amplitude are being sent there is only minimal increase in network traffic, as compared to the alternative of sending a whole waveform over the network. Secondly, because it is a real-time update, if the sample is changed in the audio module, the visualisation will update to reflect this.

3.3 Results

Informal testing has given positive results, especially in the way that people use the visual feedback to determine more accurately what they are playing. The only confusion seems to be in the slight rotation of the virtual cube that accompanies the movement of the play-heads. This misleads people to think that the rotation of the real cube controls rotation of the virtual cube. To remedy this the rotation can be removed, then reinstated when some form of rotational tracking is added to the cube.

The hand-held cube design has been beneficial, as people are not as intimidated as they may be if presented with a traditional instrument. This has proven useful in a gallery situation, where the cubes' design needs to invite people to pick them up and interact with them.

It has been found that the sample-scrubbing model works well for abstract sounds and expression, however it proves to be not particularly suitable for more controlled or measured performance, especially if a pitched sound is required. Rhythmic sounds can be convincingly used, and the gestures involved in rapidly moving towards and away from the cube are a successful method of playing the PETECUBE. Another method of playing that has been found is using a strong uni-directional light-source (such as a desk-lamp) and rotating the cube without deliberately covering the LDR's with the hand. This allows rotational gestures to be used for easily repeatable sound generation. For finer control of samples, a method of cupping the hands over opposing faces of the cube, and using the palms to block out light allows subtle movements to be captured.

A current problem lies in the haptic feedback. The mapping between the audio model and the two motors is underdeveloped compared to the audio and visual feedback, resulting in slight incoherency between the sound and the vibrations. Although the mapping is being developed further, it is felt by the author that a more sophisticated haptic system needs to be explored in future PETECUBE's to catch up with the development of the audio and visual feedback.

4. FURTHER WORK

Now that the basic system is set up and functional, it is possible to continue research into further variations of PETECUBE (for current progress see [10]). The aim is to create a series of cubes, each with different combinations of sensors, actuators and control models, which can be used as the basis for investigation into feedback in musical interfaces. Examples of potential cubes are listed below:

- Record-Cube. A cube that can record and play back data from all of its active modalities. Can this be used with a series of cubes for a form of multimodal sequencing?
- Twist-Cubes. Two cubes joined by a motor and encoder. This would allow more advanced force feedback in a rotational manner.
- Shock-Cube. Can unpleasant feedback (such as electric shocks) be used in a multi-modal interface?
- Tele-Cube. Cubes that are connected at a distance over a network (or the internet). Is directing feedback from one cube to another remote cube useful in collaborative music making? Can feedback from two remote cubes be simultaneously displayed in a single cube?
- Tracking-Cube. The use of an inertial or gyroscopic tracking system would allow the visualisation to accurately follow the movement of the cubes, while also adding an extra input modality.
- Push-Pull-Cubes. Two cubes connected by a linear damper (such as a Magneto-Rheological Fluid Damper). This would allow the user to use the cubes in an accordion-like manner, with controllable linear damping. Effects could be explored such as making it harder to 'push' through a louder sample.

Another aspect to be explored in further work is the embodiment of all the feedback modalities within the cube. The haptic feedback is already localised within the cube, however the audio and visual modalities rely upon speakers and monitors respectively. The audio speakers are likely to be the easiest to locate within the cube, whereas the visual augmentation will require advanced Augmented Reality techniques.

Other intended further work will involve standardising the OSC message system so that the software modules will become interchangeable. This will then lead to creating templates for each module in various languages (C++, Java, MaxMSP) so that it is a straightforward task for other people to develop their own modules. Due to the relatively cheap parts and very simple design, it is then hoped that people will experiment with building their own PETECUBE to accelerate the research in multimodal feedback.

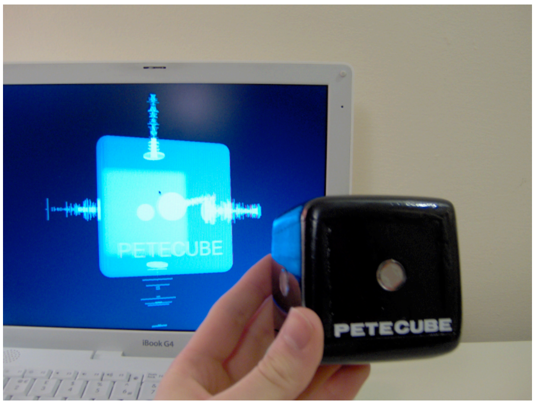

Figure 4. The PETECUBE in use.

5. CONCLUSION

This paper has discussed the definition of a multimodal feedback interface, discussed design concerns in its realisation, given an account of the current state of the PETECUBE, and outlined possible further work. As the project progresses, it is hoped that the PETECUBE will become the basis for many experiments into multimodal-feedback instruments.

6. ACKNOWLEDGMENTS

The author would like to thank the NIME reviewers for their feedback.

7. REFERENCES

[1] Schomaker, L., Nijstmans, J., Camurri, A. Lavagetto, F., et al., *A Taxonomy of Multimodal Interaction in the Human Information Processing System*, Esprit Basic Research Action 8579, Multimodal Integration for Advanced Multimedia Interfaces (MIAMI) 1995.

[2] Collins English Dictionary. http://www.collins.co.uk/wordexchange/, 2006

[3] Rodet, X. Gosselin, F. Mobuchon, P. *Study of haptic and visual interaction for sound and music control in the PHASE project.* In Proceedings of the 2005 International Conference on New Interfaces for Musical Expression (NIME05), Vancouver, BC, Canada.

[4] Wright, M. Freed, A. Momeni, A. *OpenSound Control: State of the art 2003.* In Proceedings of the 2003 conference on New Interfaces for Musical Expression (NIME 03), Montreal, Canada.

[5] O'Modhrain, S. *Playing by Feel: Incorporating Haptic Feedback into Computer–Based Musical Instruments.* PhD thesis, Stanford University, Palo Alto, California, November 2000.

[6] Max/MSP, http://www.cycling74.com/, 2006

[7] Scavone, G. Cook, P. *RTMIDI, RTAudio and a Synthesis Toolkit Update* In Proceedings of the 2005 International Computer Music Conference, Barcelona, Spain.

[8] OpenGL.org. http://www.opengl.org/, 2006

[9] Bourke, P. *Creating Anaglyphs Using OpenGL.* http://astronomy.swin.edu.au/~pbourke/opengl/redblue/, 2006

[10] PETECUBE, http://www.petecube.com/, 2006