Ashitaka: an audiovisual instrument

Niall Moody
Centre for Music Technology
University of Glasgow
UK
niallmoody@yahoo.co.uk

Dr. Nick Fells
Music Department
University of Glasgow
UK
n.fells@music.gla.ac.uk

Dr. Nicholas Bailey
Centre for Music Technology
University of Glasgow
UK
N.Bailey@elec.gla.ac.uk

ABSTRACT

This paper describes the Ashitaka audiovisual instrument and the process used to develop it. The main idea guiding the design of the instrument is that motion can be used to connect audio and visuals, and the first part of the paper consists of an exploration of this idea. The issue of mappings is raised, discussing both audio-visual mappings and the mappings between the interface and synthesis methods. The paper concludes with a detailed look at the instrument itself, including the interface, synthesis methods, and mappings used.

Keywords

audiovisual, instrument, synchresis, mappings, X3D

1. INTRODUCTION

Over the years there have been many artworks produced which seek to combine sound and visuals[1]. This paper presents what the author believes is a novel approach to this field of audiovisual art, focussed specifically on the design of an audiovisual instrument. The approach is based on Michel Chion's notion of synchresis in film, with the hypothesis that such a phenomenon is based primarily on the motion present in both the visuals and the audio. The paper describes a mapping methodology derived from this hypothesis, which aims to create an instrument where the output is perceived as an audiovisual whole, with the audio and visual streams not easily separated in the mind of the audience and performer. The paper concludes with a discussion of the Ashitaka instrument, which is currently in development and which has been designed according to the aforementioned principles.

[1] [8] presents a good overview of the field. See also music video for an extremely wide range of audiovisual combinations.

NIME07, June 7-9, New York, NY

2. MOTION AS THE CONNECTION

2.1 Synchresis

"the spontaneous and irresistible weld produced between a particular auditory phenomenon and visual phenomenon when they occur at the same time."[2]

Synchresis is a term coined by Michel Chion in relation to film, to describe the way in which certain audio and visual phenomena may be perceived as linked when they occur simultaneously, according to certain conditions. A concrete example of synchresis is the way in which punches are represented in film. Whereas in real life there is rarely much sound associated with someone being punched, in film we have become accustomed to hearing assorted whacks and thumps when the punch connects, to the point where it seems almost unreal, and somehow false, when a punch is depicted naturalistically. What we can see happening here is two otherwise unrelated audio and visual events being perceived by the audience as quite intimately connected, with the sound actually enhancing the image, part of what Chion terms "added value".

2.2 How Synchresis Works: a hypothesis

The longstanding presence of synchresis in film, and the way it generally goes unnoticed in the minds of the audience, suggests that it may be possible - according to certain constraints - to therefore create connections between audio and visuals that will be perceived in much the same way by virtually any audience, regardless of their cultural background. The next step is to start to look at how this may be used in an audiovisual instrument, and if possible, reduce it to simple principles that are easily put into practice.

Looking at the example of filmed footsteps (an example Chion discusses to some extent in Audio-Vision), and starting with a simplified look at the visuals, what we see is a foot in motion - it moves towards the floor, collides, and moves off. If we take a similarly simplified look at the audio and concentrate on the amplitude envelope, what we have is an envelope with a sharp attack at the point where the foot connects with the floor, followed by a short decay. From this admittedly simplistic analysis, we can see a clear connection between audio and visuals in that the motion we see in the visuals is closely related to the behaviour of the amplitude envelope. Indeed, looking at it from another perspective, the motion present in the audio (i.e. its amplitude envelope) is clearly connected to the motion we see in the

[2] [9], p.63

visuals, and this is the main hypothesis behind the design of the Ashitaka instrument. Obviously with this example we are cutting out a lot of contextual information, but the question is whether that information substantially alters the synchretic relationship. The author is of the opinion that it can be mostly discarded - if we replaced the visuals with a simple stick figure animation, and the audio with a simple enveloped sine tone, the synchresis would still be perceived. Indeed, we could go further still and replace the stick figure with simple geometric shapes colliding in a similar fashion, without the synchresis being adversely affected. This ability of synchresis to still function even with the most abstract of materials will be extremely important when we start to build up our own audiovisual connections.

With the footsteps example, our experience of seeing and hearing footsteps in real life clearly has a part to play - we have certain expectations about what's going to happen, what we're going to hear and see. And yet, synchresis still seems to function if we completely remove the context, and make everything as abstract as possible. It is the author's opinion that our experience of the world leads us to expect certain things from perceived motion. If we see something is moving, our brain expects to hear an accompanying sound, and vice versa, if we hear a sound which has some kind of motion present, we expect to see something moving accordingly. Obviously this rule does not apply to every situation in real life (one only has to imagine listening to a CD, or seeing someone wave), but our experience with objects which do emit sound when we can see they're in motion seems to make our brain much more receptive to linking otherwise unconnected audio and visual phenomena when it perceives a certain similarity in the temporal information of the two streams.

While it is essentially derived from a particular branch of film theory, this idea that motion can act as a connection between audio and visuals does appear to have some precedent in recent psychological research. There have been a number of perceptual effects identified which seem to point towards an integration of the audio and visual streams in the brain. For example, the McGurk effect[3] notes that certain visual stimuli can affect speech perception. If the sound of someone saying /ba/ is played to the image of someone saying /ga/, it is actually perceived by the audience as /da/. Another effect - known as the 'illusory flash' effect[17] - notes that if a single visual flash is accompanied by two rapid auditory beeps, the audience will actually perceive two flashes. The 'ventriloquist effect'[4] could also be included in this discussion. While these effects do not necessarily explain synchresis, which is surely a far more general principle than the specific effects examined here, they do seem to agree that there is a certain integration of audio and visual information in the brain. Indeed, in their investigation of another audiovisual illusion, the authors of the "Multisensory Integration of Dynamic Information" chapter in the Handbook of Multisensory Processes[18] note the following:

"Considered together, the data point to the conclusion that the experience of motion is critical for cross-modal dynamic capture to occur, and therefore this illusion reflects the integration of dynamic information."[5]

[3]First noted in [14].
[4]Originally discussed in [12], also noted in [9].
[5][18], p.57

These various findings would therefore appear to back up the author's hypothesis that synchresis is based on motion.

3. MAPPINGS

Given then that it may be possible to create an audiovisual connection where otherwise unconnected sound and visuals are perceived as one, the next step is to try and put this information to use. This is where the issue of mappings comes in, both in relation to the audio to visual mappings (and vice versa), and those of the performer's gestures to the instrument's output.

3.1 An Initial Approach

The initial approach taken with regard to audiovisual mappings was to attempt to separate and categorise the various kinds of motion which may be put to use in such mappings. This involved separating motion into *forms of motion* and *domains in which motion may occur.* Forms of motion refers to the way in which something moves, while the domain refers to the 'something' that is moving. So, a form of motion could be the collision-based motion referred to in the footsteps example, while a domain could be the position of a visual object.

Table 1: Forms of Motion

Constant Velocity
Collision-Based Motion
Periodic Motion
Gravity-Based Motion
Discontinuous Motion

Table 1 shows some example forms of motion. Generally these should be self-explanatory. Gravity-based motion refers to a kind of motion governed by forces of attraction and repulsion, like the motion of planets. Discontinuous motion refers to motion which is primarily made up of sudden jumps, similar to the rapid cutting often seen in music videos.

Table 2: Domains in Which Motion May Occur

Visual	Aural
Position (of an object)	Instantaneous Amplitude
Size (of an object)	Pitch
Rotation	Brightness
Smoothness	Energy Content
Articulation (of an object)	Spatial Position
Pattern	

Table 2 shows some example domains in which motion may occur. Again these should be fairly self-explanatory. Visual smoothness refers to the continuum of coarse, jagged shapes to smooth, rounded shapes. Articulation refers to the way in which animals may articulate their limbs. Aural brightness refers to the perceptual 'brightness' of a sound. Instantaneous amplitude and energy content are closely related, both essentially referring to the amplitude of the audio, where energy content represents a more long-term (perhaps running average) view of the signal's amplitude.

Some brief examples of mappings based on these categories are shown in figure 1. The first one refers to our simplified, abstract version of the footsteps example. The

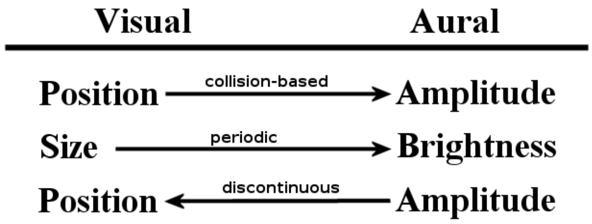

Figure 1: Some example audiovisual mappings

second one has the size of a visual object controlling the brightness of the audio (say by altering the cutoff of a low pass filter) in a periodic fashion, while the third example has transients detected in the audio jumping an object to a random new position.

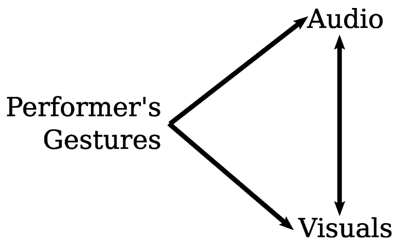

Figure 2: Simple performer-audio-visual mapping scheme

This approach works reasonably well in terms of describing our earlier, film-based audiovisual examples, however the introduction of a performer into the equation would seem to complicate things. For instance, if we look at existing mappings research with respect to musical instruments, forms of motion are rarely if ever discussed, presumably because this is something that is directly under the performer's control, and generally works at a higher level to the mappings involved in an instrument. The fact that the performer is in control of the forms of motion then surely negates, or at least reduces the impact of any audiovisual mappings set up with the previous scheme. Another problematic area is the question of how we fit the performer's input into the scheme. If we take a simplistic approach, and just implement it as in figure 2, we still have a problem as to how to create an output that's perceived as an audiovisual whole. This scheme doesn't really offer any suggestions as to how that would be achieved, so it seems we need a more structured approach.

3.2 Relevant Prior Work

There has of course been a lot of research conducted into mappings for musical instruments in recent years. A brief recap of some of the significant ideas from this research is included here as background to the following discussion of the mapping methods used in Ashitaka.

There are three main strategies which describe how one set of parameters may be mapped to another: *one-to-one*, *one-to-many*, and *many-to-one* mappings (the last two are also sometimes known as *divergent* and *convergent* mappings respectively). These strategies appear in much of the research, and tend to form the basis for most existing mappings. A number of experiments carried out by Andy Hunt[13] seem to suggest that often, one-to-one mappings may actually be less useful than more complex mappings such as one-to-many or many-to-one. The experiments demonstrated that the subjects found it easier to use interfaces with complex mappings (generally few-to-many), and that they also enjoyed these interfaces more. The reason given being that multiple one-to-one parameters require a greater cognitive load than fewer one-to-many parameters. Two different directions are also present in the existing research, with some researchers advocating mappings explicitly defined by the instrument builders, and others instead advocating neural network- or artifical intelligence-based solutions. Generally, the neural network approach is intended to create instruments which respond and adapt to the performer's gestures, and their mental image of how the instrument works, rather than the other way round (which tends to be the case when learning existing instruments).

Among the various mapping approaches that have been developed, there are two that the author feels are particularly interesting. The first is outlined in the paper *'Mapping transparency through metaphor: towards more expressive musical instruments'*[11], where the authors describe various instruments which were all designed around particular metaphors. *'Sound Sculpting'*, for example, uses a clay-based metaphor, where the performer controls the instrument with the kind of gestures you might use with a block of clay. *'MetaMuse'*, on the other hand, uses a rainfall metaphor (the performer controls it with a sensor-equipped watering can) together with granular synthesis. This approach is interesting because in a way it sidesteps a lot of the issues involved in mapping (what's the most appropriate parameter to map to this input? etc.) by basing the mappings on existing, well-known phenomena. As such, the development of the mappings is focused on emulating the 'source' of the metaphor, as the performer would understand it. The second approach makes use of *'perceptual spaces'*[7], where low level sound synthesis parameters are first mapped to more meaningful, perceptual parameters or spaces (i.e. the perceptual 'brightness' of a sound rarely has a parameter of its own in sound generating algorithms), before they are mapped to an input, the parameters of which have also been mapped to a perceptual space. This approach is particularly interesting because it makes use of how we perceive our gestures and the effects of our actions, rather than a more trial and error approach of connecting low level parameters together in the hope that a satisfying connection is made.

The first thing to note about the approaches mentioned is that they are solely concerned with *musical* instruments. Although visuals are involved in some of the instruments based on the aforementioned metaphor-based and perceptual spaces ideas, these are solely to aid the performer in understanding the instrument's operation, and are clearly subservient to the audio. Indeed it seems that very little research has been conducted into mappings for use in audiovisual instruments. The Computer Music article *'Dynamic Independent Mapping Layers for Concurrent Control of Audio and Video Synthesis'*[15] does examine this area somewhat, but the resultant mapping strategy seems to rely on an implicit assumption that if the audio and video synthesis engines are controlled by the same input signals, there will be a clear connection made between the two dimensions. While this approach may prove useful, we are interested in a perhaps more intimate connection between audio and visuals, where there are explicit audio-visual mappings in addition to those of the performer's input gestures. The perceptual spaces approach is perhaps the most relevant of the mapping strategies mentioned here, as it ties back to our discussion of synchresis being based on our perception and experience of the physical world.

3.3 'Audiovisual Parameters'

Following on from this, we come to the mapping strategy devised for the Ashitaka instrument, which itself relies somewhat on the use of perceptual parameters. The aim, as stated earlier, is to create an instrument which outputs an audiovisual whole, as opposed to apparently separate audio and visual streams. As such, this mapping scheme proposes the use of 'audiovisual parameters' - high level parameters which control individual 'audiovisual wholes' within the wider instrument. Figures 3 and 4 demonstrate how these parameters would fit into such an instrument.

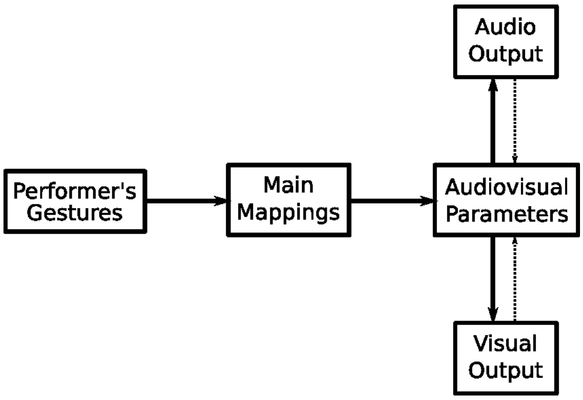

Figure 3: Basic mapping scheme

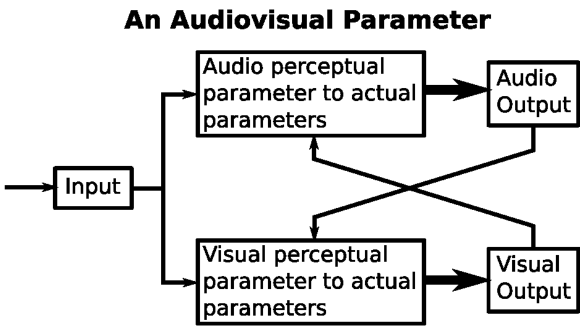

Figure 4: Audiovisual parameter scheme

As can be seen from figure 4, an audiovisual parameter consists of a single perceptual audio parameter, and a single perceptual visual parameter, both controlled by the same input, but also connected in some way to the output of the opposing domain. A perceptual parameter in this case refers to a high level description of some part of the (audio or visual) output, which is easily perceived by the audience or performer. An example could be the perceived brightness of the audio, something which is generally affected by multiple parameters in the synthesis engine (as such, a perceptual parameter is itself a mapping layer). The perceptual parameters also accept input from the opposing output in order to create a more intimate audio-visual connection. To understand why this is necessary, if one considers the (aural) amplitude envelope of a physical model as one perceptual parameter, the model will traditionally continue to resonate after the performer has stopped their input. If the visuals do not receive any feedback from the audio output, they will have no way of knowing that the audio is still effectively in motion, leading to a break in the audiovisual connection, a conflict between what the audio is doing, and the (presumably static) visuals.

Encapsulating audiovisual parameters in this way also allows us to make use of some of the other mapping strategies mentioned previously, as we can treat the audiovisual parameters as equivalent to the audio synthesis parameters in a more conventional musical instrument. This is where the *'Main Mappings'* stage in figure 3 comes in, and it is intended that this will be where the main focus on the instrument's 'playability' will be (in the form of one-to-many or many-to-one mappings), with the audiovisual parameters acting as fixed audiovisual wholes, perceived in the same way regardless of cultural background.

4. EXPERIMENTAL 'MINI-INSTRUMENTS'

As part of the development of the preceeding theory, a number of experimental 'mini-instruments' were developed in order to investigate audio-visual mappings in a performance context. All of them are designed to be performed via a MIDI fader box. A selection are discussed here.

4.1 Yakul

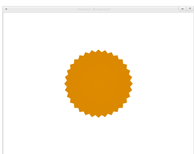

This was the first instrument developed, and is essentially designed around the rate of change of a single input slider. When no motion is detected, there is no sound output, and the visuals only show the white background. When the slider is moved slowly, an orange circle appears in the visuals, while the amplitude envelope of an additive synthesizer rises in the audio domain. When the slider is moved fast, the orange circle becomes a star shape, and the additive synthesizer's higher partials become prominent. The idea was to link sudden movement on the input to harsh, jagged shapes and sounds on the output. As such, the audiovisual connection is relatively strong, though as an instrument it is extremely limited and not particularly successful.

4.2 Moro

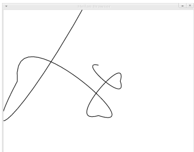

Moro is one of three instruments with monochrome bezier curve visuals and a simple string-based physical model for the audio. The other two instruments investigated audio-visual mappings where visuals were solely controlled by the audio, and the audio was solely controlled by the visuals (as opposed to being controlled directly by the performer's gestures), respectively. Moro, by contrast, has the performer directly controlling both sound and visuals (with sound and visuals also mapped to each other). While it was closer (on paper) to the author's original idea of an audiovisual instrument, Moro's audiovisual connection is not strong, and as an instrument it is not particularly successful, as it is lacking in expressive range and is not much fun to play.

4.3 Jiko

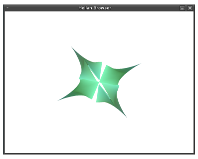

The comparative failure of Moro led to the development of the 'Audiovisual Parameters' approach to audio-visual mappings, of which Jiko is the most complete example. The visuals of this instrument consist of four 3d NURBS surfaces, potentially allowing the performer high level control over the surfaces' shapes, while the audio again consists of a string-based physical model. Jiko has four audiovisual parameters, with the audio and visual perceptual parameters involved being; Audio Energy Content and Visual Size; Audio Pitch and Visual Angularity; Audio (instantaneous) Amplitude and Visual Rotation; Audio Decay Time and Visual Colour 'Temperature'. Jiko is definitely the most successful of the mini instruments when considered as an instrument, though the audio-visual connection is not a strong as would be desired. It is thought that this may be down to the relative lack of expressive range in both audio and visuals.

5. ASHITAKA

Having described the mapping strategy used, we can now move onto a discussion of the Ashitaka instrument itself. Though not complete yet, the instrument is essentially a 3d object within a wider 3d environment in the computer, running a piece of software developed specifically for this purpose, called Heilan.

5.1 Heilan X3D Browser

Heilan is an X3D browser[6], meaning that it displays and allows interaction with 3d worlds stored in the X3D file format[6]. As such, the Ashitaka instrument is a single node within a wider scene, allowing for interaction with other objects in the scene, or indeed other instruments. The use of an established 3d (or virtual world) specification makes it trivial to create complex visual and aural spaces within which the instrument can play. Though Heilan does not support the entire X3D specification[7], it is already capable of displaying complex scenes.

Audio in Heilan is spatialised in 1st order Ambisonic B-format[10]. This allows for full 3d sound output, provided enough speakers are used. Heilan is the first X3D browser which focuses on audio performance, using a low latency audio engine (courtesy of PortAudio[3]) which, though common in music software, is something of a rarity in the X3D world.

The other significant feature of Heilan, not common to X3D browsers, is Open Sound Control[19] support. Heilan runs as an OSC server, and its OSC implementation allows for any object in an X3D scene to have its attributes exposed to an OSC client, by simply assigning the relevant node an OSC address. This gives the software substantial flexibility, particularly for realtime performance. It is through OSC that the Ashitaka instrument's physical interface will control the instrument.

A relatively simplistic audiovisual piece has already been produced with Heilan, called *'origins'*. The piece was presented as part of the Musica Electronica concert series at Glasgow University, and features audiovisual connections derived from the synchresis/'motion as the connection' ideas discussed previously[8]

Heilan is open source (licensed under the GPL) and written in C++ using the SDL[4], TinyXML[5] and PortAudio libraries. It is currently available for Linux and Windows (an OSX version is also planned), and can be downloaded from: `http://www.niallmoody.com/heilan/index.htm`

5.2 The Interface

Following an examination of currently-available musical (and general purpose) interfaces, it was felt that Ashitaka would be best served by a custom interface. The interface designed for the instrument is derived from a clay-based metaphor (inspired by the Sound Sculpting instrument discussed earlier). It is therefore intended to present the performer with the ability to use the same gestures they would use with a block of clay. As such, it can be stretched, twisted, and has 4 force sensors to enable more 'sculptural' gestures. It also has a 3-axis accelerometer to provide positional data. The interface is wireless, connecting to a PC via bluetooth. The Free2Move F2M03AC2[1] bluetooth module is used for this, running its wireless UART firmware so it appears as a serial connection on the PC. The data from the sensors is first converted to a digital format, then passed to this module by a PIC16F874A[2] microcontroller. Figure 5 shows the current (working) prototype of the interface:

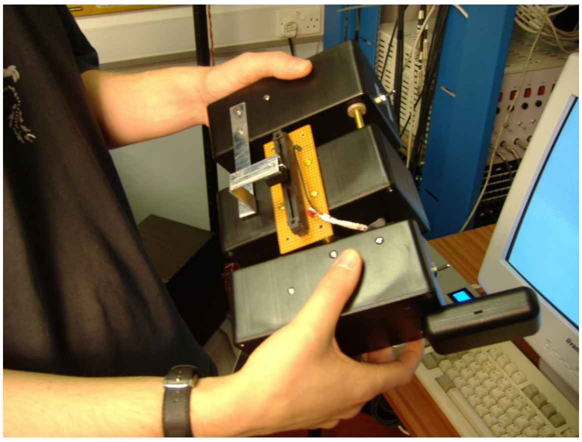

Figure 5: The physical interface to Ashitaka

5.3 Audio

Ashitaka's audio is handled by a modified version of the Tao physical modelling language[16]. Tao uses a particularly processor-intensive method of physical modelling, so in order to run in realtime, Ashitaka is restricted to a basic string model. This model has, however, been supplemented with various methods of excitation, and various DSP algo-

[6]X3D is the (primarily XML-based) successor to VRML.

[7]It currently supports the X3D Interchange profile, with some additions - most notably the Sound component.

[8]A video of the piece can be downloaded from: `http://www.niallmoody.com/heilan/videos.htm`

rithms, to expand the range of expression possible with the instrument.

5.4 Visuals

The visual component of the Ashitaka instrument is still very much in flux at the time of writing. A number of different approaches have been attempted, one of the earliest being a 3d object made up of NURBS surfaces (similar to the Jiko instrument described earlier), discarded for not providing a large enough expressive range. The next approach taken was a sphere whose vertices could be individually deformed, though this resulted in a clumsy and not particularly pleasing visual aesthetic. The current approach (shown in Figure 6 below) is to map a 2d Tao surface (not sonified) to a sphere. The performer then can excite the surface as they excite the underlying Tao string in the audio synth. The sphere can also be twisted, shown in the third image. The motion of the Tao surface makes for a more dynamic and interesting visual output to the approaches tried previously, but it is still very much lacking in expressive range, and the performer doesn't have nearly as much control over it as they do over the audio. As such, the final visual output will probably look very different to the images below.

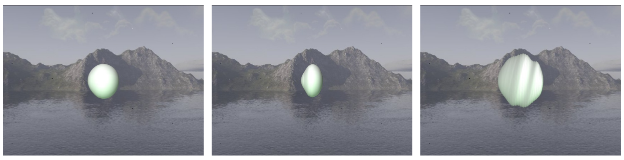

Figure 6: Ashitaka's visual output

5.5 Mappings

As mentioned previously, the instrument uses a clay-based metaphor for its interface. This influences the mapping significantly, as this metaphor primarily relates to visual (and haptic) sensation. As such, certain parameters are essentially mapped in a one-to-one fashion (for example, the position data from the interface corresponds directly to Ashitaka's position in the X3D environment, and thus the spatial position of the audio, and where the visual object appears on screen). A number of the audiovisual parameters have also been initially derived from the visual side, attempting to fit an aural perceptual parameter with a visual one such as the length of the object. More advanced mappings have also been introduced however, derived from high level determinations of how the performer is, say, shaking the interface, or the detection of sudden, sharp movements.

6. CONCLUSION

Based on Michel Chion's notion of Synchresis, a strategy of creating audiovisual mappings was introduced, followed by a discussion of how these mappings may then be used in an audiovisual instrument. Based on these principles, the Ashitaka audiovisual instrument was then outlined, along with the hardware and software built to accommodate it. This includes the physical interface and the audio and visual synthesis methods used, as well as the initial mappings developed for the instrument, though some work remains to be done in order to tie the various elements together to form a cohesive and satisfying instrument.

7. REFERENCES

[1] Free2Move Wireless products. Website: http://www.free2move.se/. Accessed: 25/4/06.

[2] Pic16f874a microcontroller. Wesbite: `http://www.microchip.com/stellent/idcplg?IdcService=SS_GET_PAGE&nodeId=%1335&dDocName=en010238` Accessed: 23/01/07.

[3] PortAudio. Website: `http://www.portaudio.com/`. Accessed: 25/4/06.

[4] Simple Direct Media Layer Home Page. website: `http://www.libsdl.org/`.

[5] TinyXML. Website: http://www.grinninglizard.com/tinyxml/. Accessed: 25/4/06.

[6] Web3D Consortium. website: `http://www.web3d.org/`.

[7] D. Arfib, J. Couturier, L. Kessous, and V. Verfaille. Strategies of mapping between gesture data and synthesis model parameters using perceptual spaces. *Organised Sound*, 7(2):127–144, 2002.

[8] K. Brougher, J. Strick, A. Wiseman, and J. Zilczer. *Visual Music: Synaesthesia in Art and Music Since 1900*. Thames and Hudson, Apr. 2005.

[9] M. Chion. *Audio-Vision: Sound on Screen*. Columbia University Press, 1994.

[10] D.G. Malham. Spatial Hearing Mechanisms and Sound Reproduction. website: `http://www.york.ac.uk/inst/mustech/3d_audio/ambis2.htm`, 15 Feb. 2002.

[11] S. Fels, A. Gadd, and A. Mulder. Mapping transparency through metaphor: towards more expressive musical instruments. *Organised Sound*, 2002.

[12] I. Howard and W. Templeton. *Human Spatial Orientation*. Wiley, 1966.

[13] A. Hunt and M. M. Wanderley. Mapping performer parameters to synthesis engines. *Organised Sound*, 7(2):97–108, 2002.

[14] H. McGurk and J. MacDonald. Hearing lips and seeing voices. *Nature*, 264:746–748, 1976.

[15] A. Momeni and C. Henry. Dynamic Independent Mapping Layers for Concurrent Control of Audio and Video Synthesis. *Computer Music Journal*, 2006.

[16] M. Pearson. Tao. Website: `http://taopm.sourceforge.net`, 31 Aug. 2005. Accessed 18/01/2006.

[17] L. Shams, Y. Kamitani, and S. Shimojo. What you see is what you hear. *Nature*, 408:788, 2000.

[18] S. Soto-Faraco and A. Kingstone. *The Handbook of Multisensory Processes*, chapter Multisensory Integration of Dynamic Information, pages 49–68. MIT Press, 2004.

[19] M. Wright. Open Sound Control Home Page. website: `http://www.cnmat.berkeley.edu/OpenSoundControl/`, 23 June 2004.