3DAudioScape

v 1.4

Reference Manual

Dave Hunt
23/01/2015
T: +44 (0)208743 3055
M: +44 (0)7798 844 091
E: davehuntaudio@btinternet.com

Contents

1 Introduction

This manual describes the operation of 3DAudioScape v 1.4. This is an update of previous versions, which includes the ability to move the listener position, the implementation of OSC messages to control the position of sound sources, and various bug fixes and programming improvements.

The software, using Core Audio and Core Midi on Mac OSX, is designed to use affordable multi-channel audio interfaces such as those from MOTU, Digidesign, Focusrite and others with appropriate drivers. These interfaces generally have eight analogue audio inputs and outputs and a further one or more 8-channel ADAT I/O interface. The latter can be used as digital inputs and outputs, to a mixing desk say, or as a further eight analogue inputs and outputs by using an ADAT/analogue converter such as those from Behringer or Focusrite.

It has long been apparent that the user interface of previous versions of 3DAudioScape, designed to control ambisonic sound processing, would also be appropriate to control other spatial audio algorithms. This development of 3DAS is an attempt to arrive at a series of programs with different spatial audio engines but the same set of controls and storage files, so that extensive programming using one algorithm can then be applied to another algorithm with no, or only minor, modification being required.

There are currently four versions: 3DAudioscape, “3DAudioscape Bin”, “3DAudioscape AEP”, and “3DAudioscape Delta”. They implement different spatial audio algorithms than the first-order ambisonics used in 3DAudioscape.

3DAudioScape can run in standalone mode, using analogue or digital inputs, and outputting signals to be fed to amplifiers and then loudspeakers. It can use up to sixteen inputs. There are sixteen mono/stereo panners, allowing for 16 mono panners, 8 stereo panners, or any combination of mono or stereo panners up to the input channel limit of sixteen. Inputs can be fed analogue or digital signals from an external source, or from another audio program on the same computer.

It provides outputs to up to sixteen loudspeakers.

The limiting of input and output numbers to sixteen is fairly arbitrary, being based on the facilities provided by affordable multi-channel audio interfaces, and the amount of audio processing power available in commonly used computers, rather than any inherent limitation. Few users will have access to a larger number of speakers, and the means to mount them in suitable positions.

It can be used with an Audio/Midi “sequencing” package (Digital Audio Workstation or DAW), on the same computer (a further reason to minimise the audio processing demands of the program). 3DAudioScape has been tested with ProTools 9 to 11, Logic Pro/Express, Ableton Live, Digital Performer and Cubase. These DAWs can also be run on a separate computer with its own Audio and Midi interfaces, sending analogue or digital audio signals and Midi to 3DAudioScape running on a second computer. ProTools versions prior to version 9 have proprietary limitations, due to Digidesign”s copy-protection strategy, and have to run on a separate computer.

3DAudioScape accepts Midi Controller and Program Change data from other packages and OSC messages. Thus manual control from a variety of MIDI control surfaces, WiiMotes, Apps running on iPhones or iPads, or other programs is possible

2 Co-ordinates and Units

A mixture of Cartesian (xyz) and Polar (horizontal and vertical angles and distance) co-ordinate systems is used.

x, y and z are distances from the centre point of the speaker rig. Standard ambisonic convention is somewhat counter-intuitive and out of line with that generally used in computer graphics, so it was decided to use the latter convention, while ensuring compatibility with existing ambisonic material “under the hood”.

x (called L/R) is left-right, with positive values to the right, and negative values to the left. y (called B/F) is back-front, with positive values to the front, and negative values to the rear. z (called D/U) is down-up, with positive values upwards, and negative values downwards.

Azimuth (Azi) is a horizontal angle in degrees. Zero is centre front, and angles increase clockwise, opposite to standard ambisonic convention. Plus or minus 360 degrees is full circle, i.e. identical to zero.

Zenith (Zen) is the vertical angle in degrees. Zero is the horizontal plane, and angles increase in the up direction to 90 degrees and are negative below horizontal, -90 degrees being directly below.

Gain amounts are normally in dB. –90dB is effectively off, 0dB is unity gain. In some cases gain is expressed as a multiplying number. 0 (zero) is off, 1 is unity gain, and negative numbers reverse the polarity of the signal.

3 Installation and Set Up

Computer Requirements

Minimum: Apple Intel computer, 2GHz dual core, 2 GB RAM. Mac OSX 10.6.8
Recommended: At least Apple 2.4Gz 4Gb RAM (Intel). Mac OSX 10.6.8 to 10.9.

All the digital signal processing within 3DAudioscape is “native” (i.e. performed by the computer) so the faster, the more powerful, the more RAM, the better, particularly if it desired to run other audio programs at the same time.

Audio and Midi Hardware

A PCI card based, Firewire or USB2 audio interface. This should have as many outputs as the number of speaker you wish to use, and as many analogue or digital inputs as you require.

A Midi interface. Most multi-channel audio interfaces contain a Midi interface, but these are also available as separate items usually using a USB connection. Many keyboards and controllers now also have a USB connection which carries Midi data.

There are a number of hardware controllers, which output Midi data, available. Most Audio/Midi sequencers (DAWs) can send Midi data to external devices. Thus a program running on one computer can control 3DAudioscape running on a second computer, spreading processing power and allowing the use of slower computers.

A USB Midi controller can interface easily with 3DAudioscape to provide “hands-on” control of many parameters.

Installation

Unzip the downloaded file and copy the 3DAudioScape folder to a suitable place on your hard drive, usually in the Applications folder. It may be useful to make an alias of the application only by dragging it to the Dock.

The program is copy protected.

When attempting to run the application for the first time you will be presented with a Password dialogue, which gives you a #Password number, which is unique to that computer. If you email this to me I will send back a Password, which should be copied and pasted into the Password box when you next try to run the application.

E: davehuntaudio@btinternet.com
dave@3d-audioscape.com

The demo version does not require a password, but will mute for four seconds after two minutes of operation, and again every two minutes.

NB
Proper initialisation requires that the program can read two files, one called "3DAS.xml" and the second called “3DAS-Panners.xml”, which should reside in the same folder as the application, and not in a sub-folder. These ensure that parameters are set to suitable values rather than to zero, which would render the program mute. These files can be modified and saved with the same names to suit the user’s system and preferences.

Other Software

3DAudioScape can run in standalone mode, taking audio from analogue or digital inputs, processing them, and outputting signals to be fed to amplifiers and then loudspeakers. Many parameters can be controlled by Midi, allowing the use of Midi controllers.
Soundflower (free) software can be used to route the audio signals from a DAW running on the same computer to 3DAudioScape
www.cycling74.com
Jack, also free, is an alternative way to route the audio signals between audio applications and the hardware audio inputs and outputs of an audio interface.
www.jackosx.com

Useful additional software

Audacity is a useful free multi-track editor that can read many audio formats and convert between them.
http://audacity.sourceforge.net/

De-Interleaver is a useful Mac OSX application that can interleave mono audio files into a multi-channel file, and split a multi-channle interleaved file into separate mono files.
http://www.birmingham.ac.uk/facilities/BEAST/research/mulch.aspx

SoundFilesMerger has similar functions.
http://www.e--j.com/?page_id=197

Junxion software can be used to convert joystick data into Midi data to allow manual control of spatialisation in either 3DAudioScape, or a DAW.
http://www.steim.org/steim/junxion_v4.html

OSC commands can be managed and converted to and from Midi with “Osculator”

4 3DAudioScape

Previous versions of 3DAudioScape used an ambisonic algorithm for sound spatialisation, with a binaural encoding of virtual speakers providing a binaural output for headphone listening. This is continued with 3DAudioScape 1.4, though versions using alternative algorithms are now available.

3DAudioScape Bin is a version in which the sources are directly binaurally encoded. This produces a better binaural simulation than that in the ambisonic version. To provide a usable 3D reverberation model, sources are also encoded and sent to the same ambisonic reverb engine used in the ambisonic version of 3DAudioScape. This is then binaurally recoded from the output of virtual loudspeakers as before. As reverberation is by its very nature “diffuse”, first order ambisonics, and the slightly “fuzzy” nature of the binaural encoding process this is considered acceptable. The reasoning behind this approach is that there is no obvious way to make a three-dimensional binaural reverb. It also allows the same reverberation model to be used in other versions using different algorithms, simplifying the achievement of consistency between them.

3DAudioScape AEP uses an “Ambisonic Equivalent Panning” algorithm. This was formulated by Martin Neukom and Jan Schacher in 2007/08. It combines encoding and decoding into one process for each sound source. It allows the use of higher order ambisonics, and this order can be different for each sound source. Higher order ambisonics gives greater directional precision than first order, though it does require using a greater number of loudspeakers to avoid spatial images being pulled towards the loudspeakers. The reverberation model is identical to that in the ambisonic version.

3DAudioscape Delta is based on the ideas of “Delta Stereophony”. This was formulated and used in large-scale theatre sound systems in East Germany in the 1970s, when digital audio electronics and computers were not readily available behind the Iron Curtain. Sound can be presented to audiences in large spaces over a large number of loudspeakers by altering the delay times and levels (and possibly also the equalisation) of mono signals sent to each speaker. Thus each signal to be transmitted would be subject to an amplitude/delay matrix of processing to derive speaker feeds. Mixing of sources takes place at this speaker-feed level. There have been a number of systems from Richmond Sound Design, Level Control Systems and TiMax from Outboard Electronics that use similar ideas. Pursued to its logical conclusion it becomes Wave Field Synthesis, which requires a very large number of loudspeakers and considerable digital signal processing power. The reverberation model is identical to that in the ambisonic version.

Other versions with alternative algorithms (e.g. VBAP) may be produced at a later date.

The following is based on the ambisonic version of 3DAudioScape. The other versions are basically similar, though inappropriate controls have been hidden and disabled, or appropriate ones revealed and activated. This is done to ensure that any .xml files are interchangeable between the various versions.

Main Window

This is intended to be as simple as possible. Further functions are contained on other windows accessed from here via buttons or a menu.

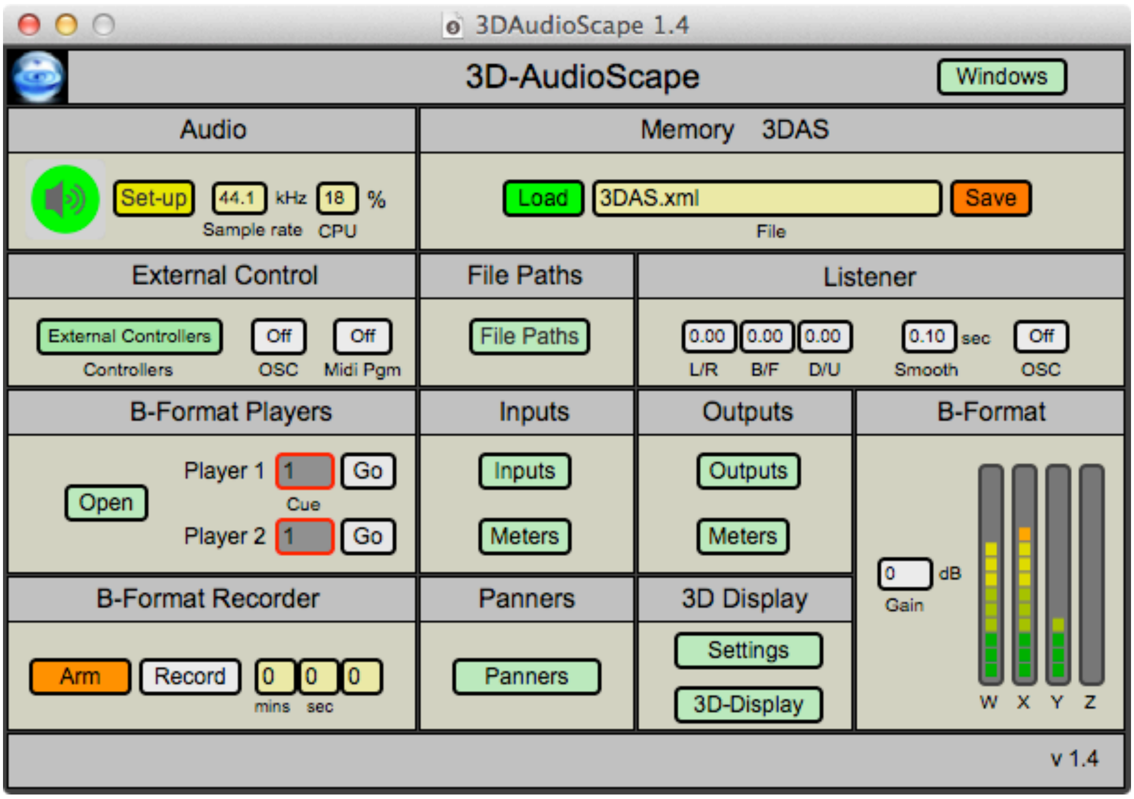

Light green buttons open windows for detailed operations and settings for the function named on the button.

Probably the most important of these other windows is the “Outputs” sub-window. Settings here should match the loudspeaker layout. If not the audible results will be plain wrong, and subsequent operations extremely confusing. Once correctly set up for the loudspeaker array to be used these settings should be saved and will then be recalled automatically. They will only require alteration if the loudspeaker array is altered, or if working on a different array.

To change the settings of any “number box” simply scroll the numbers with a vertical mouse movement, or click in a box, type in the number, and press “Return”. Number boxes with a decimal point can be coarsely changed by placing the cursor on the integer (before the decimal point), and scrolling. Fine changes are achieved by placing the cursor on the number after the decimal point and scrolling.

Some number boxes cannot be changed by the user and just provide information: these are a yellow colour, e.g. the “Sample rate” and “CPU” boxes, and the timer boxes in the “B-Format Player” and “B-Format Recorder”.

Other boxes contain drop-down menus.

Windows

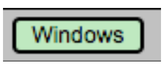

The “Windows” switch, top right, opens a window in which collections of control windows can be assembled. The windows can then be dragged around the screen to desired positions and the result stored to one of ten memory locations. This is useful if you wish to quickly switch between various combinations of windows.

Audio Panel

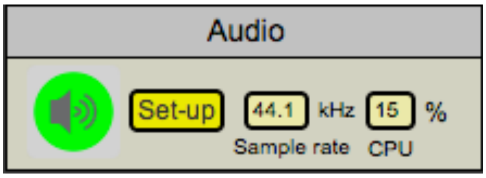

This panel contains a switch to turn the audio “On/Off” and a “Set-up” button, which opens a window allowing the selection of the audio interface and adjustment of various audio digital signal processing (DSP) parameters. Two boxes indicate the current sample rate and percentage utilization of the audio processing power of the CPU.

Audio Set-Up Window

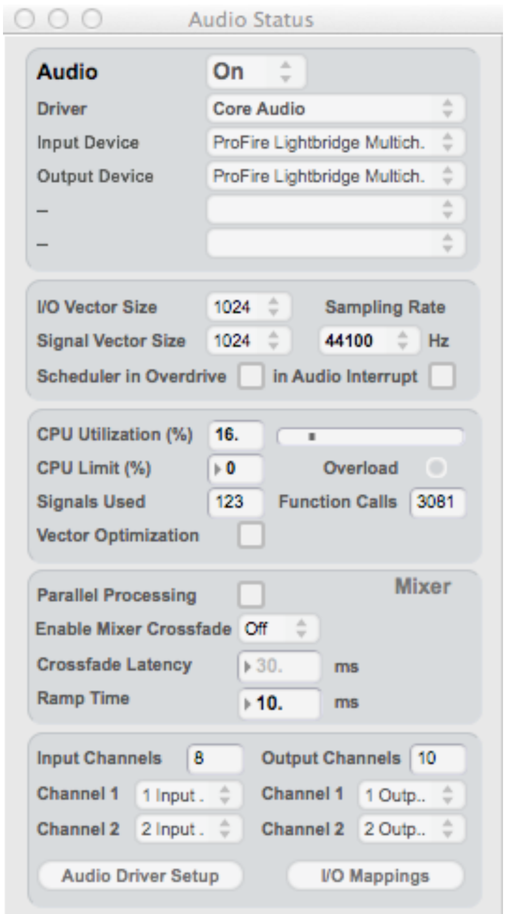

Select “Core Audio” as your “Driver” and your audio interface as the Input and Output Device.

The “Sampling Rate” should match that of any audio to be played, or to be input from a digital input or other DAW. The “I/O Vector Size” (i.e. buffer size) affects the latency and audio performance. The I/O and signal vector size boxes should be set to 512 unless audio glitches or CPU overloads occur, or latency is perceived as problematic. If glitches occur increase them as this will reduce the CPU usage.

The “I/O Mappings” button on the bottom of the “Set-up” window opens another window where the mapping of available inputs and outputs to those of 3DAudioScape can be changed. The input mappings can also be set in 3DAudioScape itself, and generally changes from the default will not be required here. Some audio interfaces (e.g. MOTU 828) will have a stereo output, labeled Left and Right, and a number of mono outputs, labeled 1-8. Left and Right will be outputs 3DAudioScape outputs 1 and 2.

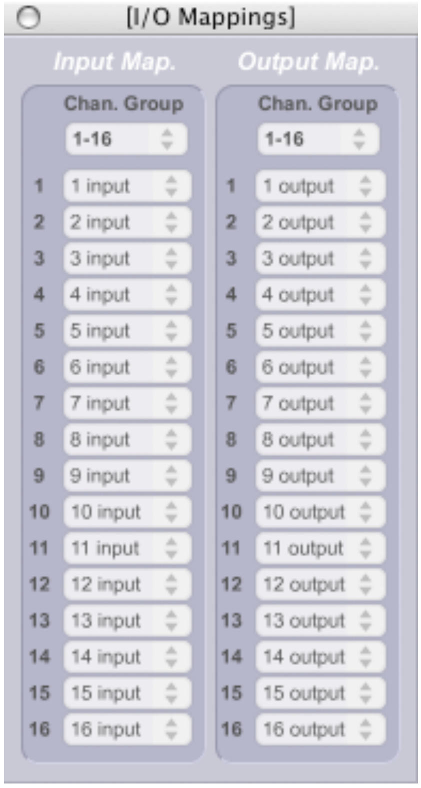

All settings are saved as a preference when the program is quit.

The “”Audio Driver Setup” button will open Apple’s “Audio MIDI Setup” window, where a global audio driver can be selected for all audio programs. This should not be necessary, as the setting in the above “DSP Status” window will override this.

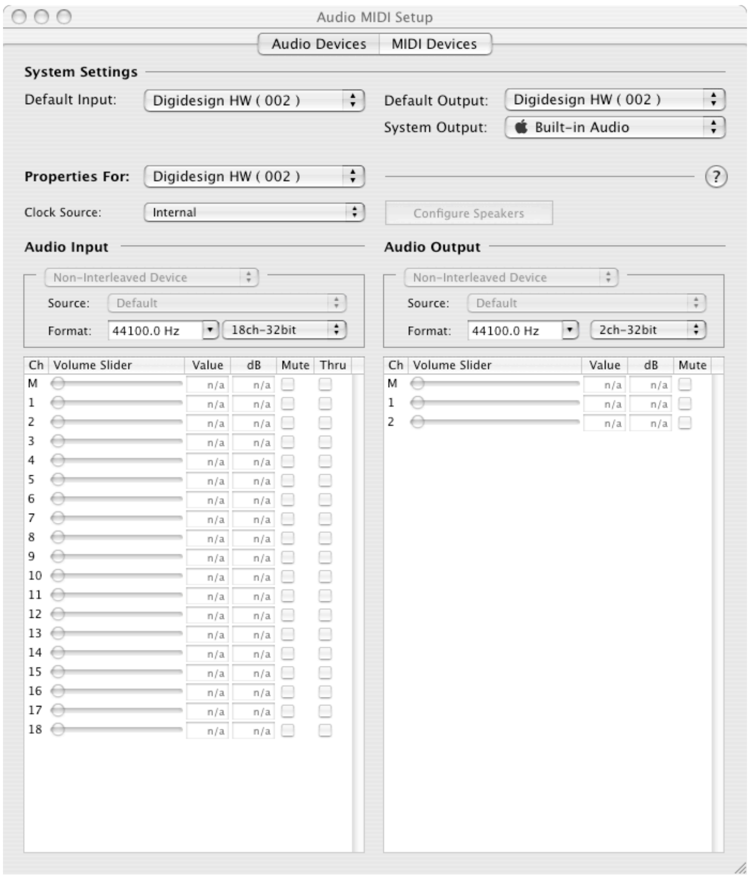

It may indeed be undesirable to set the Default output in “Audio MIDI Setup” to anything other than “Built-in Audio”.

File Paths Panel

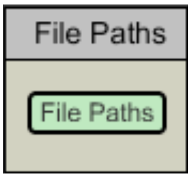

This opens a window where two folders can be chosen. Files in these folders, and sub-folders within them, then become accessible to 3DAudioScape. This is particularly useful when 3DAudioScape's audio file players are used. Audio files may be placed in folders (or even hard discs) other than that containing 3DAudioScape. Max/MSP needs to have these locations within its search path to be able to find and play them. See Section 4.1.6.

Listener Panel

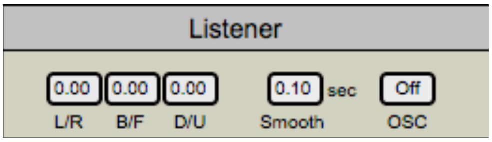

The listener position can be moved from the centre position to another position, even one outside the room defined as "V-Space". This enables the listener to move relative to all sources in the sound field simultaneously. The L/R, B/F and U/D distances are relative to the average distance of the loudspeakers. Movement is smoothed by the "Smooth" time. The "OSC On/Off" button displays whether OSC messages to control this movement are received. This is set in the "External Controllers" window.

Note that each Panner must be set to react to this movement in the "Inputs" window.

Further, note that the listener cannot move in a pre-recorded B-Format space, but only within the space where each source is independently encoded. The L/R, B/F and U/D ambisonic distances are subtracted from the L/R, B/F and U/D ones for each source, and these new coordinates are used to encode the sources. The Listener Position is not stored, and is automatically set to 0,0,0, the centre point, when the program is launched.

Memory/File Panel

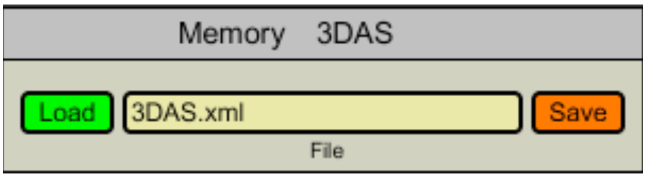

The Memory panel at the top-right of the main window stores all the main settings on all except one of the sub-windows (“Panners”) within the program. The settings in the audio “Set-up”/DSP Status window, are saved as a preference when the program is Quit.

Parameters stored in this memory are:

Midi Controller and Program Change Ports and Program Change On/Off
Midi Controller numbers, scaling and offset for each of the 16 Panners.
OSC message settings.
The physical audio input for each of the 16 Panners, its colour and its name
The physical audio inputs used as a four-channel B-format input
The settings of the “Output” window, where the positions of the loudspeakers in your array can be set up
The settings in the “3D Display Control” window. These alter the appearance of the 3D Display. This displays the position of the loudspeakers and audio inputs.

When 3DAudioScape is started a file called "3DAS.xml" is read and a previously stored state of all the above parameters is recalled. This file should reside in the same folder as the application.

If any changes are made they are stored automatically, and the user will be prompted to save when quitting the program even if they have just opened any relevant window, or scrolled a menu on the front page.

“Save” brings up a file selector box, and the stored content of the Memory can be named and saved to disc as an .xml file, preferably in the same folder as 3DAudioScape. By default this will have the name “3DAS.xml", and saving with this name will mean that any stored settings will be recalled when the program is next run. Thus your most commonly used set up will become the default.

Another name can be used, but this will have to be loaded manually after the program has initialized. “Load” brings up a file selector box, and a previously saved file can be loaded to the Memory. All settings are recalled.

NB *The “Panners” window has its own memory system, which is unaffected by the above Memory. (See Sec. 4.1.3.)*

Inputs and Outputs Panels

The “Inputs” button opens a window where the physical audio input to each panner can be selected, its colour chosen and its name entered.. (See Sec. 4.1.2)

Similarly the “Outputs” button opens a window where the position of up to 16 speakers can be defined. This sets up an ambisonic “Decoder”, or relevant calculations in other versions. (See Sec. 4.1.1))

The Inputs “Meters” button opens a window containing 16 meters, one for each input.

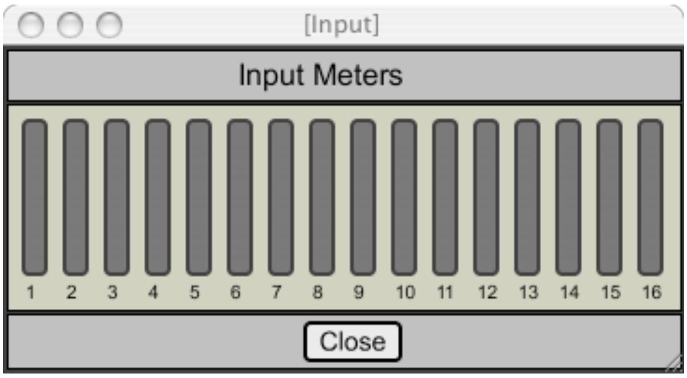

The Output “Meters” button opens a window containing 16 meters, one for each output.

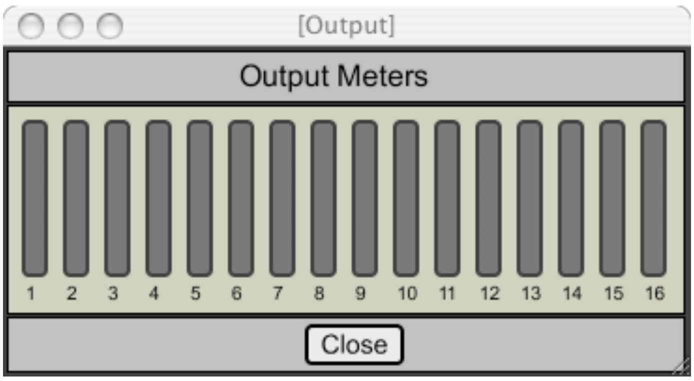

External Control Panel

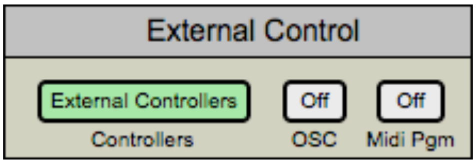

The "External Controllers" button opens a new "External Controllers" window, where the Midi controllers and OSC messages for each of the 16 panners can be set up. .. (See Sec 4.1.4)

The "OSC" and "Midi PGM" On/Off buttons display whether these are on/off. These are set in the "External Controllers" window.

Midi Program Change messages recall Memory states in the "Panners" window. (See Sec. 4.1.3)

The settings are stored in the "3DAS" Memory.

Audio MIDI Setup

Attached Midi devices should be set up first in "Audio Midi Setup". This can be found in Applications/Utilities on your system hard disc. If it is used often it may be advantageous to drag it to the Dock. This creates an alias of the application on the Dock, where it can be more easily located.

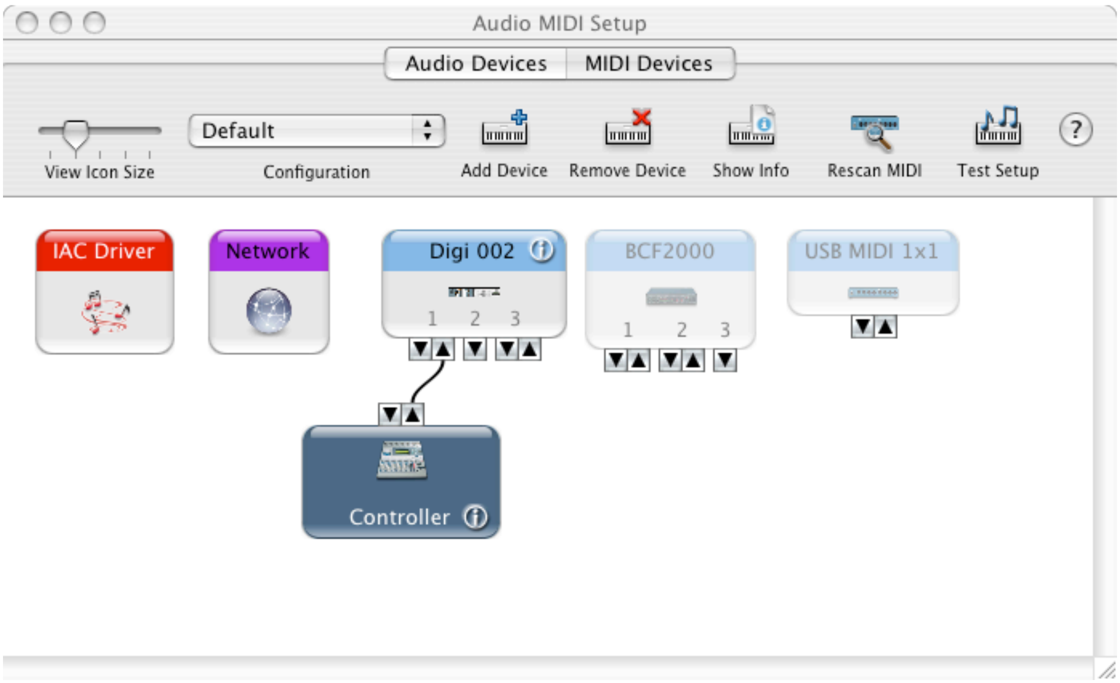

This shows a fairly simple system. The main audio/Midi interface in use here is a Digi 002. Two other previously used but currently not connected Midi interfaces are shown greyed out.

"Controller" is an external hardware Midi fader controller, and the output is connected to the Digi 002"s input port 1.

The "IAC" box is the "Inter Application Communication" system, which is installed with the Mac operating system. Clicking on it opens a new window.

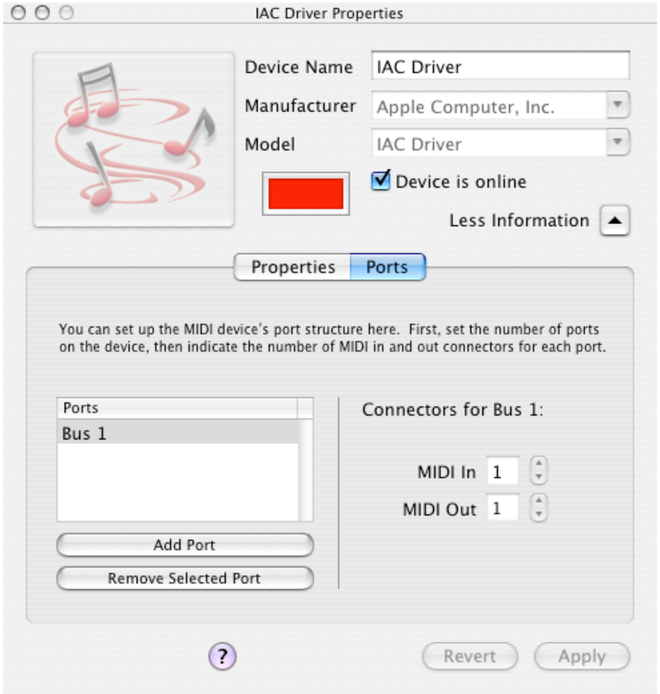

Only one bus is set up by default, though others can be created.

Further Midi set up options are to be found in the “File” menu in 3DAudioScape; “File/MIDI Setup”.

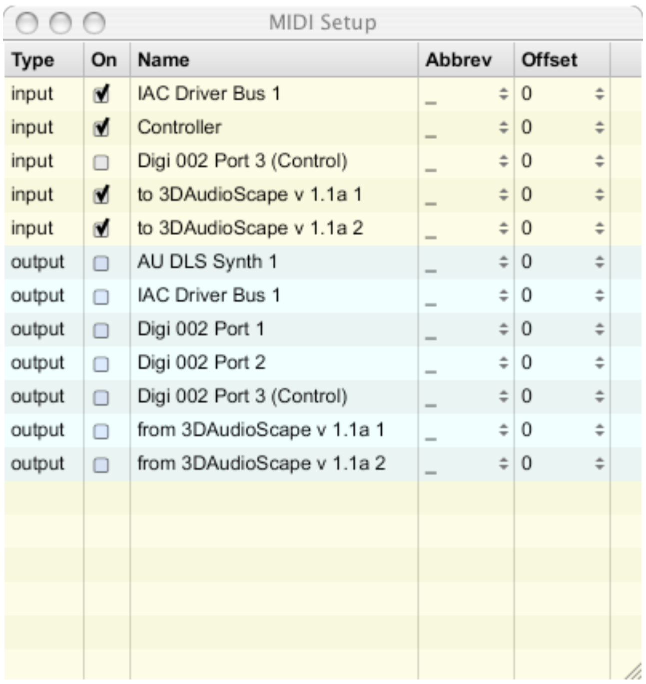

There is no Midi output from 3DAudioScape, so all output options can be disabled (unticked).

“Controller” is the external hardware Midi fader controller previously set up in “Audio Midi Setup”.

“IAC Driver Buss 1” is the “Inter Application Communication” bus shown above. This is possibly the most useful one to use when sending Midi Controller data from a DAW to 3DAudioScape, though it is also possible to use the two “to 3DAudioScape” busses.

B-Format File Player and Recorder Panels

NB *These only appear in the ambisonic version of 3DAudioScape.The Binaural version replaces the B-Format Recorder with a binaurally encoded stereo recorder.*

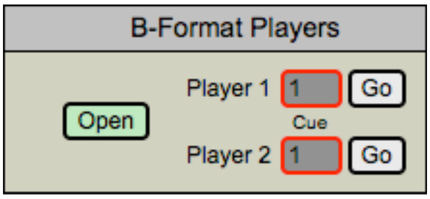

The “B-Format Players” panel contains basic controls for two B-Format File Players, in which 128 “Cues” can be stored.

A Cue number for each player can be set in the “Cue” box then played with the “Go” button.

“Open” opens a separate window, which contains more detailed control of the Players. (See Sec. 4.1.6)

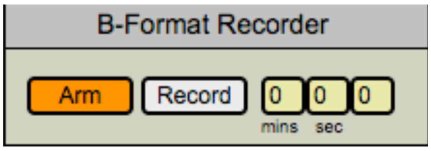

The “B-Format Recorder” panel controls the recording of an interleaved 4-channel B-Format file, which can be played with the above “B-Format File Player.” The signals recorded, and metered by the “B-Format Meters”, include everything that is audible: inputs, played audio files and the reverberation. It can be used to record a performance for archive purposes, or to record audio for use in a performance. Recorded files can be imported into the B-Format Player.

When “Arm” is pressed a window opens, allowing the naming of the file to be recorded, and its disc location. Pressing “Record” commences recording, the timer runs, and the button is labelled “Stop”. Pressing “Stop” stops the recording.

Files produced can be edited and re-saved in Logic, Cubase, Audacity (freeware) or similar audio editing package that can handle interleaved four-channel audio files.

B-Format Meters Panel

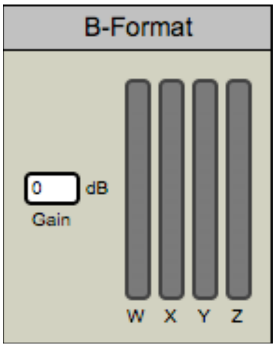

The meters show the levels of the W, X, Y and Z components of the internal B-Format signal, which is wired to the “B-Format Recorder”, and to the “Decoder”, which is set up in the “Outputs” window. The “Gain” box controls the output level.

NB *In the binaural version this is replaced with a stereo output meter, and in the AEP and AD versions is not present at all.*

3D Display Panel

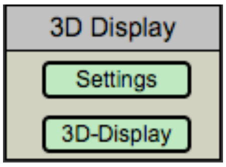

The “3D-Display” button turns the “3D-Display” window on/off.

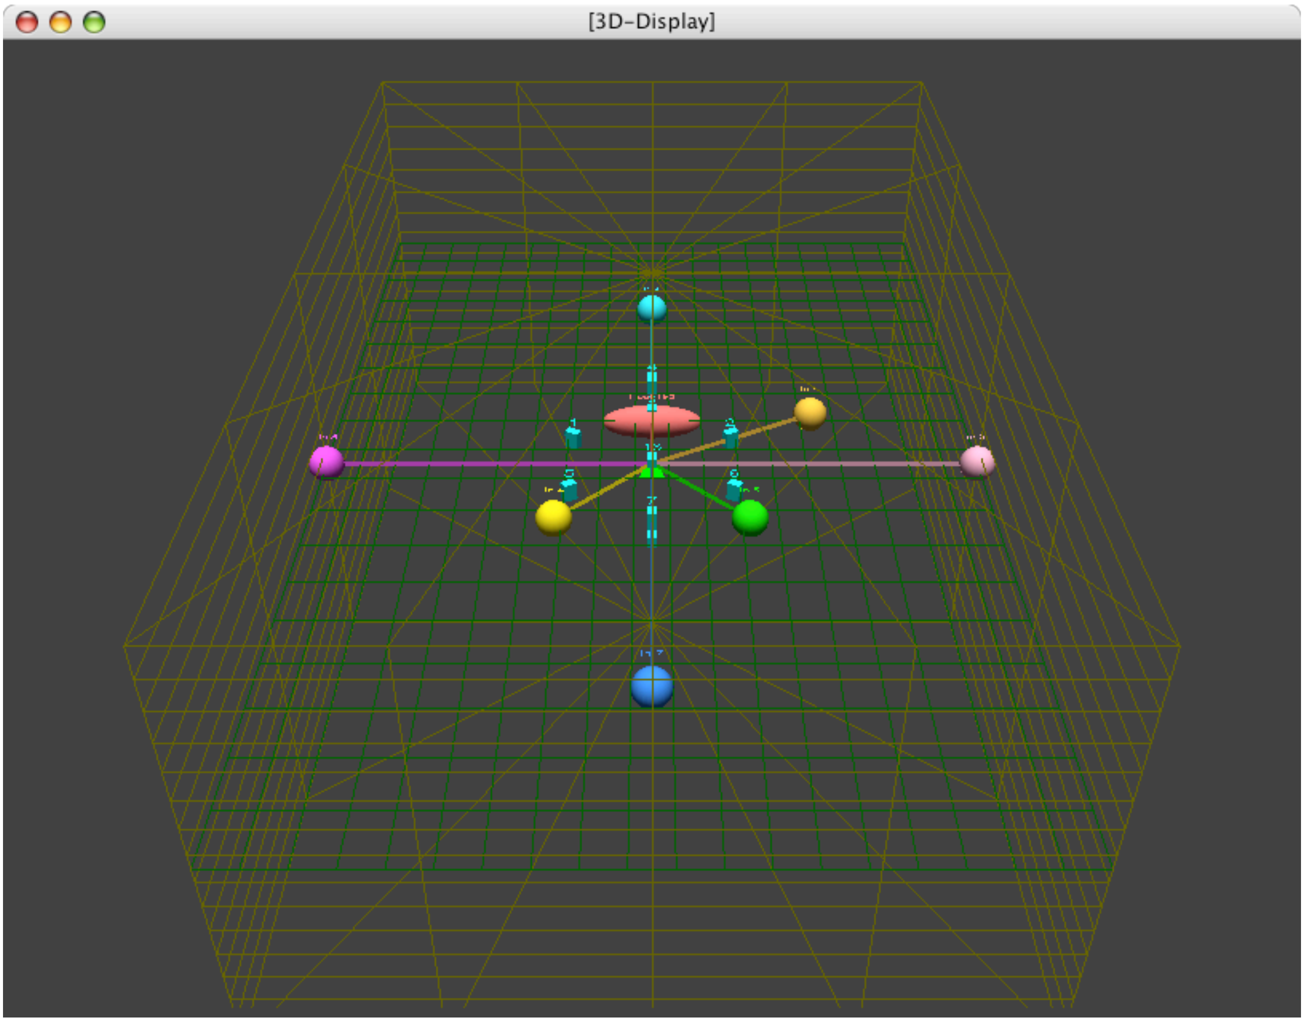

This contains a composite display of each sound source, each loudspeaker, and the virtual room set up in the “V-Space panel” (See Sec. 4.1.3).

This window can be dragged to any position on any monitor.
It shows the “Virtual Room”, displayed as a yellow grid box, “Ground” as a green grid plane.

The input sources are coloured spheres if mono and ellipsoids if stereo. The loudspeakers are cuboids. A central green triangle shows the “listening position” and the “forward” direction, for the purpose of orientation.

The image above shows an array of 8 speakers and 8 input sources.

Note that the position of the green arrow representing the listening position will move if the “Listener” position is moved. The loudspeaker array will also move correspondingly. In truth the loudspeakers themselves do not move, but the sources and the reverberation move relative to the loudspeakers. The sound field moves in the opposite direction to the listener.

The view in this display can be altered in the “3D Display Control” window, opened by clicking on the “Settings” button in the 3D Display panel. (See Sec. 4.1.5)

Windows Button

The “Windows” button opens a window where a collection of the various sub windows can be opened, dragged to desired positions on the screen, and this layout can be saved

4.1 Windows

These are presented in rough order of operational importance.

4.1.1 Outputs Window

Clicking on “Outputs” opens the “Outputs” window.

The settings here should be made to match the array of loudspeakers in use. This should be done prior to running any audio, or using the B-Format File Player, otherwise results will be incorrect and confusing.

The settings of the Decoder are stored in the “3DAS” Memory when this window is closed. All speaker settings and the output state of the Decoder, apart from the “Normal/Binaural” selection (which defaults to “Normal”) are stored.

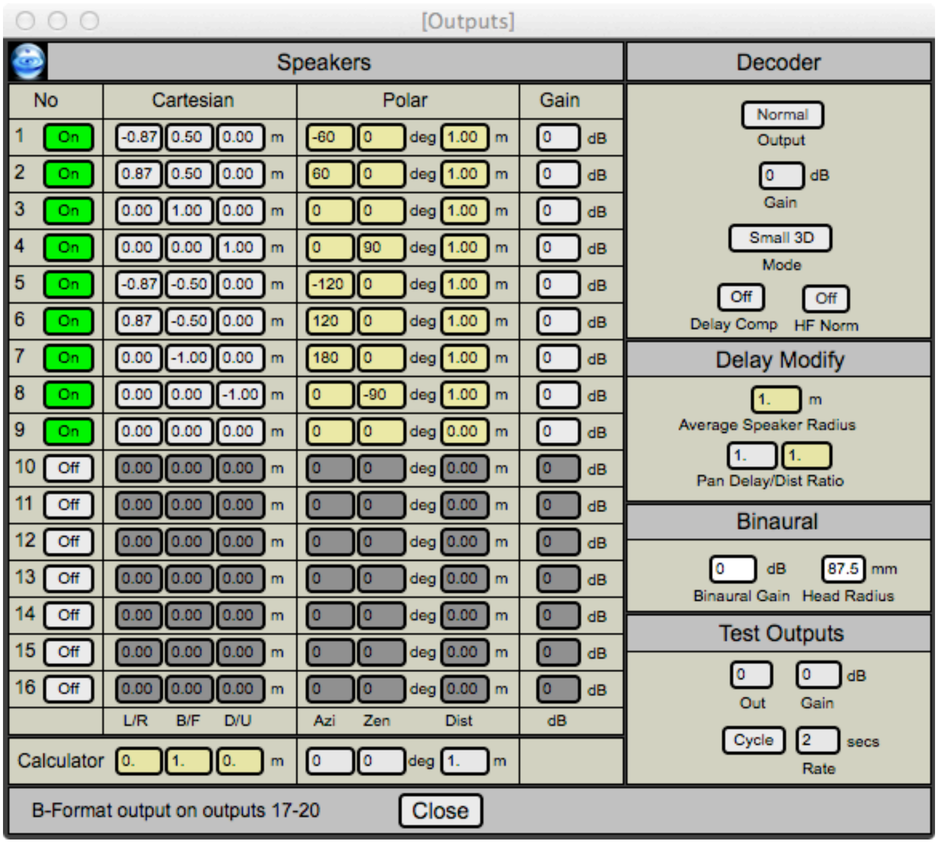

NB *This window will appear different in non-ambisonic versions of 3DAudioScape. Inappropriate and non-functional controls have been removed in those versions.*

The loudspeaker coordinates are used differently in calculations for the Ambisonic, Ambisonic Equivalent Panning, and Delta versions. The positions of the loudspeakers defined here set up any ambisonic decoder in use. The “V-Space” reverberation is ambisonic in all versions of 3D AudioScape.

Speakers can be turned On/Off individually. A switch turns green when a speaker is on. When a speaker is off all the other boxes turn grey to indicate that the controls have no effect, and that this speaker is mute. Unused loudspeakers should be muted, as this disables the relevant processing.

The position of the loudspeakers can be set in Cartesian co-ordinates and are constrained to lie within the “V-Space” virtual room. To change the speaker positions simply adjust the numbers in the “X”, “Y” and “Z” boxes. The position of the loudspeaker in polar coordinates is calculated and displayed. The average distance of the loudspeakers is calculated and sent to the “Panners”, so that their positions are related to that of the loudspeakers.

To change the speaker positions simply adjust the numbers in the “L/R”, “B/F” and “D/U” boxes. You should be able to view any changes you make in the 3D Display.

At the bottom of this window is a calculator, which allows entry of values in Polar coordinates and displays the corresponding Cartesian coordinates.

The gain of each loudspeaker can be modified to compensate for differences in perceived or measured loudness, if this is not possible externally. The maximum gain is 0 dB, so louder speakers should be attenuated. If necessary, external equalization could be applied to the loudspeaker signals to make their sound similar to each other.

Output

The “Normal/Binaural” menu determines the output mode of the decoder:

“Normal” is for output to a number of loudspeakers. This mode is selected automatically when the program is run.

“Binaural” is for headphone listening. When in “Binaural” mode, a binaurally coded stereo signal is sent out of the first two outputs, normally connected to the headphone output of a multi-channel audio interface, and the other outputs are muted.

NB *This switch only appears in the ambisonic version of 3DAudioScape.*

Gain

The overall gain of the Decoder (or output in other versions) can be altered with the “Gain” control.

Mode

NB *This switch only appears in the ambisonic version of 3DAudioScape.*

The “Mode” menu controls shelf filters at low and high frequencies in the decoder. There are ten options “Basic 2D”, Basic 3D”, “Energy 2D”, “Energy 3D”, Small” 2D, “Small 3D, “Medium 2D”, “Medium 3D”,”Large 2D” and “Large 3D”.

The effect of the filters is subtle. The theory is that the filters are required due to different mechanisms of directional perception at low frequencies than at high frequencies..

The 2D settings are nominally for 2D, horizontal, arrays and the 3D settings for full periphonic (with height) arrays. Some 3D arrays could be considered as a combination of two or more 2D arrays.

The “Basic” settings implement a “Velocity” decoder at all frequencies, which suits a single central listener.
The “Energy” settings implements” what is termed an “Energy” decoder, which is generally applicable.

“Small”, “Medium” and “Large” are not precisely defined.

The “Small” settings implement a velocity decoder at low frequencies (< 700 Hz) and an energy decoder at high frequencies, optimised for small spaces with a small number of listeners, such as studios and domestic rooms.

The “Medium” settings implement an “energy” decoder at low frequencies and an “in-phase” decoder at high frequencies.

The “Large” settings implement an “in-phase” decoder at all frequencies.

The “Furse-Malham “In-phase” Decoder, is useful in larger spaces when members of the audience are closer to some loudspeakers than to others, and may hear sources as coming from the nearest speaker rather than from the true direction. This setting leads to a larger “sweet spot” at the expense of some reduced directional accuracy.

NB *The modes of this decoder all assume “regular” loudspeaker arrays: speakers on the faces or vertices of regular polygons or polyhedra. Efforts should be made to match these criteria with all speakers the same distance from the listener, though some deviation can produce an acceptable result. It is not possible to produce a general solution for irregular arrays.*

Delay Comp (Compensation)

NB *This switch only appears in the ambisonic and AEP versions of 3DAudioScape.*

Ideally all speakers should be at the same distance from the centre of the array. This is often physically impossible due to constraints imposed by the shape of the listening space, layout of the audience, health and safety issues, and availability of suitable loudspeaker rigging positions. When loudspeakers are at differing distances from the centre of the array, the signal to nearer loudspeakers can be delayed, so that the signal arrival from them is co-incident with that from the furthest loudspeaker.

Suitable delays are automatically calculated from the entered speaker positions, and put into effect by turning “Delay Comp” “On”. Obviously, the accuracy of the calculated delay times can only be as good as that of the numbers supplied but as only one person can be at the ideal centre point, the delay times cannot be correct everywhere, so distances do not have to be excessively accurate and plus or minus 10% is likely to be accurate enough.

HF Normalisation

NB *This switch only appears in the ambisonic version of 3DAudioScape.*

Daniel also suggests that a high-frequency energy normalisation is advantageous, and this is provided as an option in this decoder with this “HF Norm” On/Off button.

Delay Modify

NB *This only appears in the ambisonic and AEP versions of 3DAudioScape.*

“Average Speaker Radius” and “Pan/Delay/Dist Ratio”

A yellow box shows the average distance of the loudspeakers in use. This distance is sent to the panners. All distances in ambisonics are related to a “unit sphere”, and the average distance of the speakers is considered to be “unity”. Sound sources can be anywhere inside, or outside this imaginary sphere, which has the same radius as the average distance of the speakers. The distance of a source is then relative to this distance, being “1.0” when on the surface of the unit sphere, between “0” and “1.0” when inside the sphere, and greater than “1.0” when outside the sphere.

In ambisonic terms the physical distance of the loudspeakers is not an important factor, as it is always considered to be unity. It is dimensionless: it is not expressed in units of feet, inches, meters or millimetres.

Sound travels at a finite speed of around 343 metres/second. Therefore it takes time for sound emitted from a source to arrive at a listener. If the source is 10 feet away, the sound from it takes around 10 milliseconds to arrive. If a source is moving towards or away from the listener, the delay time varies and this is perceived as a pitch shift, the well-known Doppler effect. The panners each contain a control, which allows this effect to be simulated to a variable degree.

The time taken for sound to travel from a source to the listener is calculated from its distance, and this distance does have a dimension, as does time: seconds or milliseconds. As all delay calculations are based on the distance of the speakers, this must be known in physical units (in this case metres). The delay time, and thus the amount of Doppler shift will then vary with the average distance of the speakers set up in the decoder. Thus problems could arise when replaying a multi-input sound field, programmed on a small loudspeaker array, on a much larger array, or vice versa.

The “Pan Delay/Distance Ratio” control can then be used to correct Doppler shifts which may become more or less extreme than desired and previously programmed.

The value for this should usually be set to a value of “1”.

Delay times increase as the number in the “Pan Delay/Distance Ratio” box increases. If moving from a small array to a large array, the “Pan Delay/Distance Ratio” should be reduced below “1.0”. If the average speaker radius used for programming was 1 metre, and the playback average speaker radius is 10 metres, this should be set at 1/10= 0.1

Conversely, when programming for a large array while monitoring on a small array, the “Pan Delay/Distance Ratio” could be increased above “1.0”. If the large array has an average speaker radius of 10 metres, and the small monitoring array a radius of 1 metre, the “Pan Delay/Distance Ratio” could be set at “10”.

The yellow box shows the result of the inverse of the “Pan Delay/Distance Ratio”

There should not be such delay problems with “rendered” B-Format recordings, which should be reproduced correctly. The angular dimensions remain the same. B-Format recordings are not subject to time modulation, so any Doppler effects in them will not be modified

Binaural

NB *This only appears in the ambisonic and binaural versions of 3DAudioScape.*

The radius of the (modelled) head can be modified. 87.5mm is the radius of the average human head. “Binaural Gain” controls the listening level. Binaural simulations are never entirely accurate, results being different for different listeners, and different headphones. It is provided here to allow work to be done and monitored when access to a loudspeaker array is not possible. A more accurate binaural rendering is supplied by the 3DAudioScape Bin version, which directly binaurally encodes each source rather than binaurally encoding the signals for six virtual ambisonically decoded speakers.

Test Outputs

The “Test Outputs” panel allows pink noise to be sent to each loudspeaker to aid identification, and can cycle the noise sequentially to the loudspeakers to aid setting them to produce a similar loudness at the central listening position. If required, external equalization could be applied to the loudspeaker signals to make their sound more similar to each other.

4.1.2 Inputs Window

Clicking on the "Inputs" button opens the "Inputs" window, where the connection of the physical inputs on the audio interface to the 16 panners is made. They can be named and colours chosen.

When this window is closed the settings are stored in the "3DAS" memory.

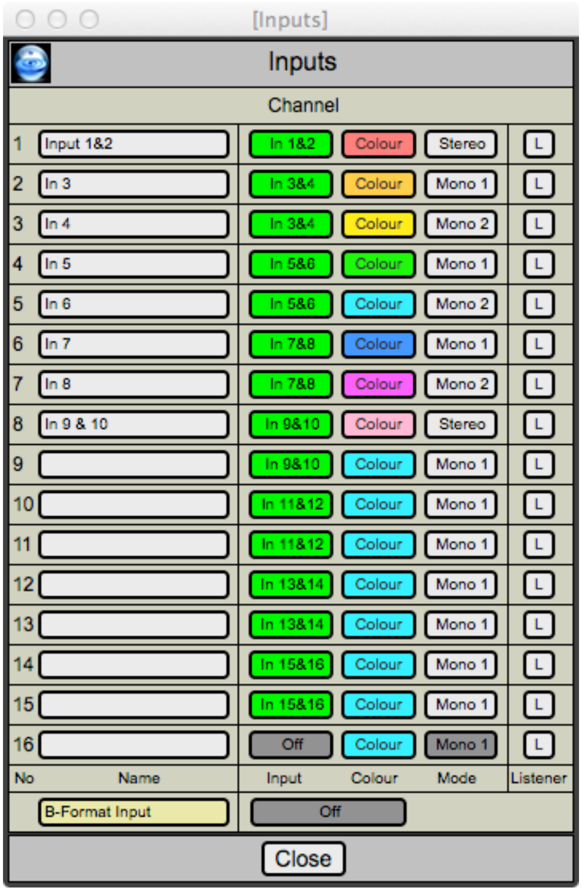

The "Input" box is a menu where the input to the each panner can be selected. These are selected as stereo inputs. Options are: Off, In 1&2 to In 15&16. These are the analogue or digital inputs of the attached audio interface.

Individual panners can be used as mono or stereo.
The "Mode" menu selects the way the input is treated.
"Mono 1" is a mono panner the input of which is the first (left) channel of the stereo input selected for that panner.
"Mono 2" selects the second (right) channel of that input.
"Stereo" configures the panner as a stereo panner, the two channels being encoded as a linked pair.
"MS" is for use with "Mid/Side" stereo inputs, which are encoded as a linked pair.

Each input can be named. Click in the "Name" box and type the name. This name then appears here and in other windows in the program.

The colour of the spheroid representing the input source in the 3D Display is selected by clicking on the "Colour" button. This opens a window where the colour of the object can be chosen.

The Listener “L” button selects whether the panner responds to changes in the “Listener” position When selected this button turns green, and the “Listener” position affects the position of the source controlled by that panner.

B-Format Input

NB. *This is only available in the ambisonic version of 3DAS.*

The “Input” selector box for the four-channel B-Format input has the options: Off, In 1-4, In 5-8, In 9-12, In 13-16. The inputs are the analogue or digital inputs of the attached audio interface. If an option that uses an input selected for one of the panners is chosen, this will be deselected for that panner and the input box for that panner will be greyed out.

NB. *The B-Format input does not appear in the 3D Display.*

4.1.3 Panners Window

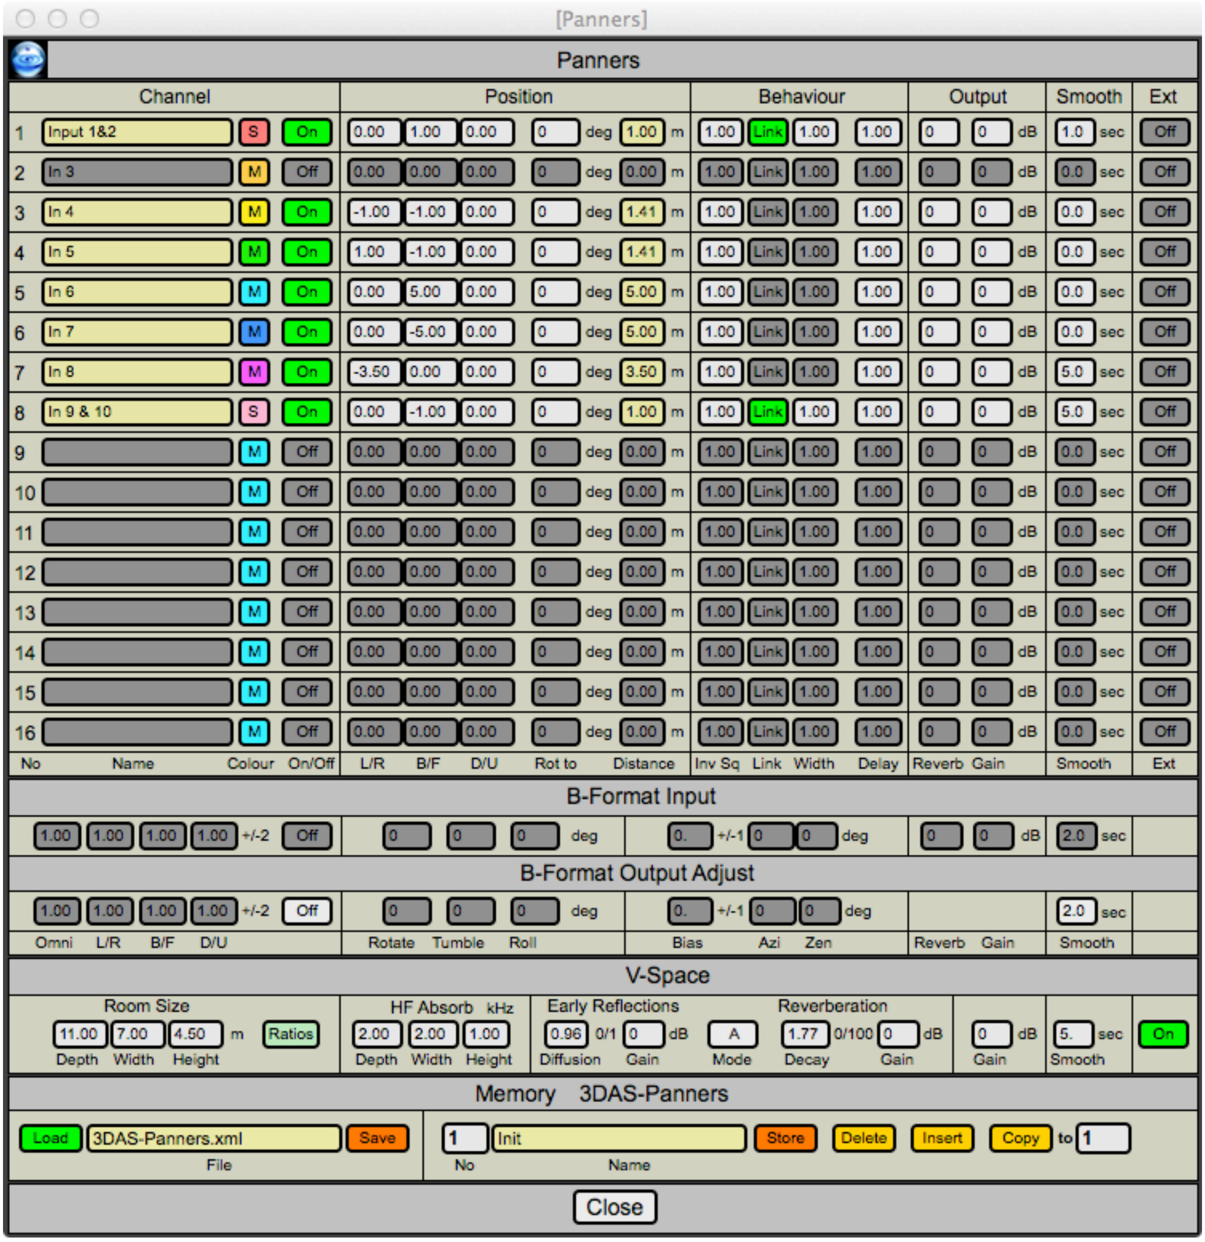

This window controls the sixteen “panners” for mono and stereo signals and the “V-Space” reverberation engine. In the ambisonic version it also controls the four-channel B-Format input and the total B-Format sound field.

Unused channels should be turned off, to disable the processing.

All settings can be stored in Memory slots in the “3DAS-Panners” memory

3DAS-Panners Memory Panel

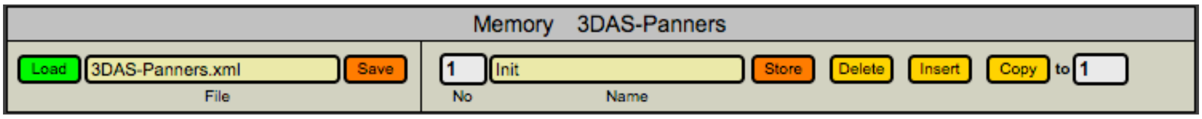

The “Panners” window has its own memory system: “3DAS-Panners. This stores the state of all the controls in this window as a “cue” in numbered Memory Locations. There are 512 memory locations. These can be given names.

A state is recalled by scrolling the “Select” box to a new number and then releasing the mouse button, or by clicking in that box, typing a number then “Return” or “Enter”.

Midi Program Change messages recall Memory states if this is enabled with the On/Off button in the “External Controllers” window (See Sec 4)

Memory states 1-128 are recalled with Midi Program Numbers 0-127 on Midi channel 1, states 129-256 are recalled with Midi Program Numbers 0-127 on Midi channel 2, etc.

.

The memory system is set only to register any parameter that has changed. So, if a location that has no stored information for a given parameter is recalled, there will be no change of that parameter.

A cue is programmed by selecting a Memory/Cue number, then setting the boxes in the “Panners” window. It is stored by clicking on the “Store” button. This brings up a dialog box where the cue can be named.

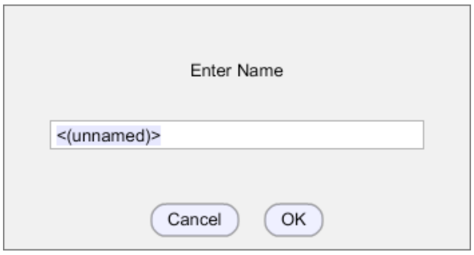

If the cue already has a name, this will appear instead of “<unnamed>“.

Cues can be deleted, inserted, or copied to other locations.

“Delete” removes the selected cue, and decrements the number of following cues by one.
“Insert” inserts a new cue prior to the selected one, and increments the number of the selected and all following cues by one.
“Copy” works by recalling a cue, selecting a cue number that the cue is to be copied to, then clicking on “Copy”.

“Save” brings up a file selector box, and the stored contents of all spatial cues within the Memory can be named and saved to disc as an .xml file. You can modify the contents of the Memory Locations and save the complete Memory content with the default name, or with another name at any disk location. If the name is the same as that of the default the file will be loaded automatically next time 3DAS is run. Files with other names can be re-loaded manually using the “Load” button.

NB. *When 3DAS is quit any changes in the stored memories for “3DAS-Panners” will cause a save dialogue box to open.*

“Load” brings up a file selector box, and a previously saved .xml file can be loaded into the Memory. Location “1” is automatically recalled.

An .xml file with the name “3DAS-Panners” is loaded automatically, and Memory Location number “1” is recalled when the program is started. This is to ensure that the program is properly initialized.

This file is placed in the folder containing 3DAudioScape. You can save your settings with the same name, in the same place, so that it is auto-loaded when the program is run.

The Difference between Store and Save

Store is like a taking a snapshot. The snapshot is placed in a memory slot. If the program is quit and the changed contents of the memory slots have not been saved to disc, you will be prompted to save. If you quit without saving the snapshots will be lost. You can, of course, save at any time.

Channel Panels

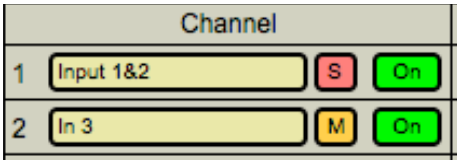

The channel number, name and colour as set in the “Inputs” window are shown here. The “On” button turns the channel on and off, turning green when on. When the channel is off this button, the name box and all the other number boxes to the right of it are greyed out.

The “Colour” button is just an indicator showing the colour, with a superimposed “M” or “S” to display whether the panner is mono or stereo.

Position Panels

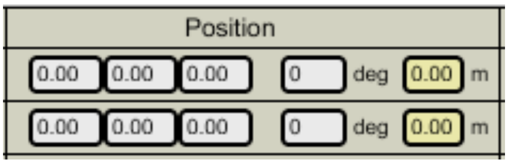

The “Position” of each input is adjusted by varying the “L/R” (Left/Right), “B/F” (Front/Back), “D/U” (Up/Down) and “Rot to” (Rotate to) parameters. All numbers in these boxes are constrained so that sources cannot go outside the “V-Space” virtual room. Sources can be positioned both inside and outside the loudspeaker array. When the distance of the source is zero, the source is centrally placed, and all loudspeakers produce the same signal. Each “panner” controls distance as well as direction.

The sound field is produced by the loudspeaker array, and the average distance of these from the listener depends on the size of this array. Thus, when moving a mono or stereo sound source with a “distance panner”, the position is not represented by distances measured in metres or feet and inches, but by ratios to the average distance of the loudspeakers. This is to ensure that using decoders for different sizes of loudspeaker array does not change the relative spatial quality of the reproduced encoded sound field.

A “B/F” value of 1, together with “L/R”, “U/D” and “Rot to” values of zero would then result in a source positioned at a loudspeaker centre-front. Decreasing the value of “B/F” would result in the signal moving towards the listener, being at the listener when zero, with all loudspeakers receiving the same signal. Decreasing this value further, going negative, moves its position to the rear and when it is –1, the position would match that of a centre-rear loudspeaker.

The distance of the source, in metres, is calculated and displayed in a yellow box, the colour indicating that this cannot be changed directly.

In the ambisonic version the output of each panner is a B-Format ambisonic signal, which is mixed internally with all other B-Format signals before passing to the “B-Format Adjust” section and then to the “Decoder”. In the other versions loudspeaker signals for each source are produced, and these are mixed to produce the final signal for each loudspeaker.

The “Rot to” control maintains the same source distance, but alters the azimuth of the source, to allow sources to travel in a horizontal arc around the centre. The range of this is unlimited, so that a source can be made to rotate endlessly around the centre. –360 and +360 degrees are identical to 0 degrees. Positive is in a clockwise direction. Combinations of “L/R”, “B/F”, “D/U” and rotation angle produce a mixture of these two motions.

Note that the displayed “Distance” value will be modified when the “Listener” position is modified

NB *Note that if the “Rot to” control is not at zero or +/- 360 degrees, the “L/R”, “B/F”, “D/U” controls will act in a correct but confusing manner. Effectively, the coordinate system has been rotated.*

Behaviour Panels

The “Behaviour” section allows adjustment of the way the sound output varies with position.

“Inv Sq” controls the amount of the inverse square law applied as distance varies. Zero is off, meaning that there is no variation in level with distance, and 1 is the theoretically correct variation. When a source is outside the speaker array, it is only the inverse square law that has any effect. If this control is set to zero, sources outside the speaker array will not be adequately simulated. Inside the speaker array (in the ambisonic or AEP versions), as a source approaches zero distance from the centre, the X, Y and Z (and higher order) components all go to zero. There is a compensating increase in the level of the W signal, so that the apparent loudness of the source does not decrease as the source nears the listener, as this would be contrary to normal perception. The effect of the inverse square law is superimposed on this W compensation.

The “Delay” box has a range of between 0 and 1. Sound sources that move radially towards or away from the centre listening position should display Doppler effect pitch variation. Examples are a police siren or train passing at speed. This has been implemented using a variably tapped delay line. The delay time is related to the simulated distance of the source, reflecting the finite speed of sound. When the distance of the loudspeakers is large, this delay will be added to the acoustic delay due to the distance of the speaker. This may adversely affect the timing of events. The Doppler pitch shift may be disturbing for musical sounds, leading to “out of tune” effects. Thus the delay is made continuously variable between 0 and 1 for each panner, “1” being the theoretically correct value.

Note that many existing recordings of moving sound sources already contain Doppler shifts.

In the panel shown above, the top channel is in a stereo mode. The “Width” control is then active. It can be varied between -2 and +2. A value of 1.0 is the normal setting. Values greater or less than 1.0 increase or decrease the width. A value of zero is equivalent to combining left and right to mono. Negative values reverse the stereo image.

The “Link” button links the Width to the inverse square law. As stereo sources become further away the width would be expected to diminish. As a stereo source moves past the listener it would be expected to increase in width to a maximum when it is either side of the listener and then decrease as the source recedes. This effect will occur when the “Link” button is green. When it is off the width will remain constant with distance.

For a “Mono” panner the “Width” and “Link” controls are non-operational and are greyed out.

Output and Smooth Panels

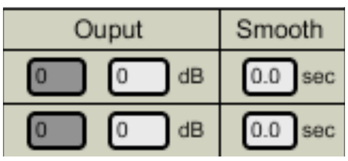

“Gain” operates as expected. It affects the output level of the source. A gain increase of 24dB is allowed, so that distant sources, affected by the inverse square law can be made louder. Care should be exercised if these sources are subsequently made to be close to the centre (zero distance), as the signal could become very loud and overload the output stages. This could cause hearing or loudspeaker damage.

“Reverb” sends a second B-Format signal to an ambisonic reverb simulated in the “V-Space” virtual room. Here the number box is greyed out, indicating that “V-Space” is muted.

The “Smooth” time is the time over which all the parameters for that channel change from their current state to a new one, whether recalled by changing to a different Memory state, adjusted manually, or remotely via Midi. This ensures a glitch-free result. For control by an external Midi controller, or a sequencer, this time can be made short, or even zero.

Ext Panel

This contains an “On/Off” button to enable Midi/OSC controllers to control the position of the panned source. The controllers are set up in the “External Controllers” window.

B-Format Input and Output Adjust Panels

NB. *This is only available in the ambisonic version of 3DAS.*

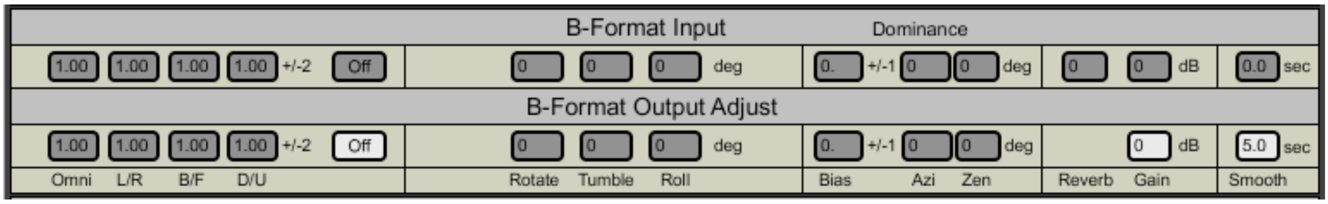

These panels affect the B-Format signals of a B-Format input and the entire mixed sound field, comprising all inputs, the reverberation and the “B-Format” Player.

The “B-Format Input” panel allows spatial manipulation of the B-Format input signal.

The “B-Format Output Adjust” panel allows global spatial manipulation of the output sound field for corrective or creative purposes, without necessitating detailed re-programming of all sources. The “On/Off” button switches the adjustments on/off, and turns the processing on/off. In the picture above it is shown as “Off”, and the other controls are greyed out.

In the “B-Format Input” panel the “On/Off” button turns the input on/off. It is shown above as greyed out, indicating that no input has been selected in the “Inputs” window.

“Omni”, “L/R”, “B/F” and “D/U” control the directionality of the sound field. A value of “1” for all is nominally correct, The range of adjustment is -2.0 to 2.0 , “L/R”, “B/F” and “D/U” allow emphasis or de-emphasis of the directional signals independently, allowing variable gain on each axis. Negative values reverse the sound field on the corresponding axis. “Omni” should normally be left at a value of “1.0”, but can be altered for unusual effects.

“Rotate”, Tumble” and “Roll”, allow rotation of the sound fields about the D/U, L/R and B/F axes by numbers of degrees.

“Bias”, “Azi” and “Zen” control a “Dominance” effect, which biases the sound field in a particular direction. “Azi” and “Zen” set the direction, and “Bias” the amount, normally zero. The range of the “Bias” control is -1 to +1, negative values producing the effect in the opposite direction.

The Gain controls affect the volume of the B-Format input and the overall volume of the B-Format output. The “B-Format Input” panel has a send level control to the “V-Space” reverberation engine.

The “Smooth” boxes smooth any changes made over the times in the boxes.

V-Space Panel

Reverberation is an essential cue for distance perception. As a sound source becomes more distant it gets quieter, the sound coming from it takes longer to arrive at the listener, and at distances of 50-100m it becomes noticeably “duller” due to high frequency absorption by the air. None of these effects gives a reliable indication of distance, particularly with sounds that do not change position, or are unfamiliar.

In a reverberant space (i.e. all spaces to a greater or lesser degree), the level of the direct sound from a source varies with distance to a much greater degree than the level of sound reflected from walls, floors, ceilings and other surfaces and objects, which remains more constant. The ratio of direct to reverberant sound levels aids our perception of the relative distance of sound sources.

Headphone listening always has a tendency to sound “in the head”. Reverberation aids in producing an “out of the head” illusion.

All sources and loudspeakers are contained within a virtual room, the “V-Space”. This is a “cuboid” room, the “Room Size” controls altering the “Depth”, “Width” and “Height” dimensions of this room. The positions of sources and loudspeakers are then constrained to be within this space.

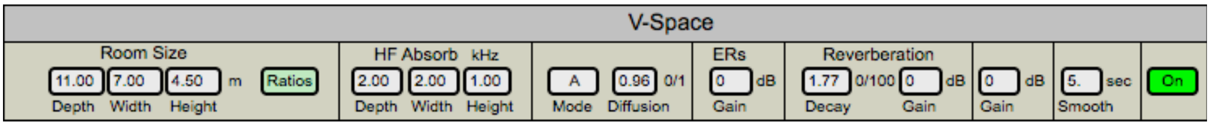

The parameters of this virtual space affect a “global” reverberation model. Each channel can be sent to this reverberation model via its own level control. The model is built along the lines of most “algorithm” reverberation units or plug-ins, with the aim of providing a useable “ambisonic reverberator”, where the spatial reverberation reflects the position of a sound source. It may not be “the best reverb” in the world, but efforts have been made to make it sound natural and usable, and provide useful effects and a common acoustic space for all sources. A better reverberator would require an early reflection generator for each source, and more processing power, and this version was felt efficient enough to be suitable at this stage. A “convolution” reverb is unlikely in the near future.

Many possible sound sources have their own built-in reverb or recorded “space”, and a global reverb can be used to help different sounds sit together, or to separate them spatially. The use of this reverb can also give a rough impression of what a soundscape composed with 3DAudioScape might sound like in another “playback” space, which would have its own reverberation.

The “HF Absorb” boxes set the frequency above which high frequencies are absorbed on each axis.

Clicking on the “Ratios” button opens a new window, which shows the ratios of the dimensions of the virtual room of “V-Space”.

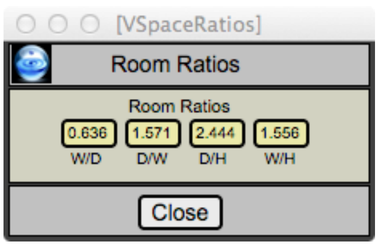

The “Room Ratio” boxes display the ratios of the three dimensions, to ease the selection of rooms where the ratios of the dimensions are not integer values (e.g. the golden ratio), which could lead to excessive colouration of the sound of the reverberation, although this can be achieved if desired. The number boxes are coloured yellow, indicating that they cannot be changed directly.

Controls

The reverberation has early refections (“ERs”) and “Reverberation” sections. Each has its own gain control.

"Mode" has three possible values, "A', "B" and "C". In "A", the levels, and decay times, of the spatial components of the reverberation are affected by the ratios of the "Room Sizes". In "B" they are affected by the inverse of these rations, and in "C" the ratios are all equal.

"Diffusion" controls the amount of feedback around a number of delays, which provide the early reflections, and affect the density of the late reverberation.

"Decay" controls the length of the reverb tail.
The "On/Off" switch turns the reverb on and off, showing green when on. When "Off" the signal processing is disabled to free up processing power. The reverb signal fades when turned off, over a time related to the size of V-Space, and fades in over the same time when turned "On".

The "Gain" control, in dB, controls the output level of the combination of early reflections and reverberation.

The "Smooth" box controls the time over which a change in any reverb parameter occurs, to avoid audible glitches, and smooth transitions between possible spaces.

4.1.4 External Controllers Window

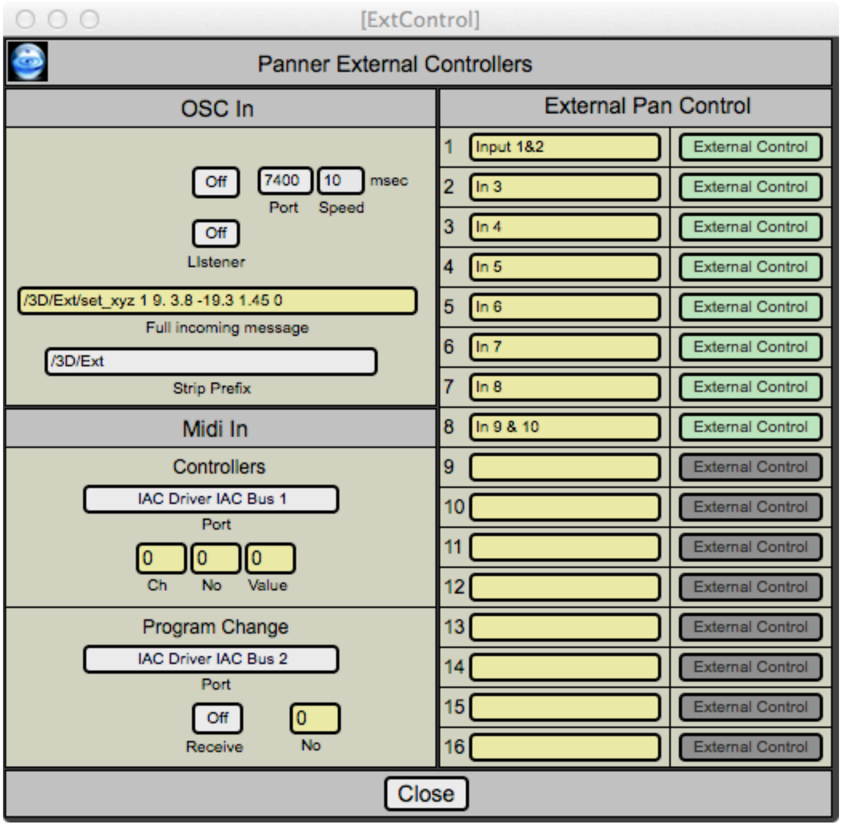

The left-hand side of this page sets the general properties of OSC and Midi messages, while the right-hand side has buttons which opens a more detailed set up window for each panner.

All settings here are stored in the “3DAS” Memory when this window is closed.

OSC In

The On/Off switch selects whether OSC messages are to be passed to 3DAS. The Listener On/Off switch selects whether OSC messages to control the Listener position are to be active.

The “Port” box defines the computer input port to receive messages. This should match the port number set in the program or device sending messages. Click in the box, and type the port number to replace the default 7400.

The “Speed” box sets a limit on the rate at which messages are read. OSC messages can contain a lot of data which change very rapidly, particularly if coming from a sensor. This could cause 3DAS to struggle to implement a message before other messages arrive, slowing its operation to a crawl as it cannot keep up. The lowest value here is 10 milliseconds, the largest 500 milliseconds, i.e. half a second.

The basic messages expected by the panners are of the form

Position
/set_xyz <int id> <float x> <float y> <float z>
where <int id> is the number of the panner,
<float x> <float y> <float z> are the L/R, B/F and D/U coordinates, typically between -1 and +1

Rotate
/rotate <int id> <int degrees>
where <int id> is the number of the panner,
<int degrees> is a clockwise rotation angle in degrees

The yellow “Full incoming message” box shows any incoming OSC message so that it can be confirmed that messages are being received, and what their form is.

OSC messages allow one or more prefixes before the data messages so that data can be sent to a unique receiver. In a complex set up not all receivers need all data so it can be advantageous to send a receiver only data that it can respond to. Data may be filtered or altered by a receiver and then sent to different receiver.

The “Strip Prefix” box allows prefixes to be stripped or added, so that the message is just of the basic form above. The default prefix is
/3D/Ext.

Again, click in the box, and type the prefixes (each starting with a slash /) to replace the default.

Midi In

For Midi Controllers, the “Port” menu is pre-populated with any Midi ports on the system. Clicking on it will reveal these and allow a choice of port.

The yellow boxes below show the Midi channel, Controller Number and Value of any incoming Midi Controller.

Similarly, the Program Change “Port” menu allows choice of the port accepting Midi Program Change messages. The “On/Off” switch sets whether such messages are to be accepted, and the yellow “No” box shows the value of any incoming Midi Program Change message.

Individual Panner External Controls

Clicking on a green “External Control” box for a panner opens a window where more detailed settings of how the data for that panner is handled.

OSC

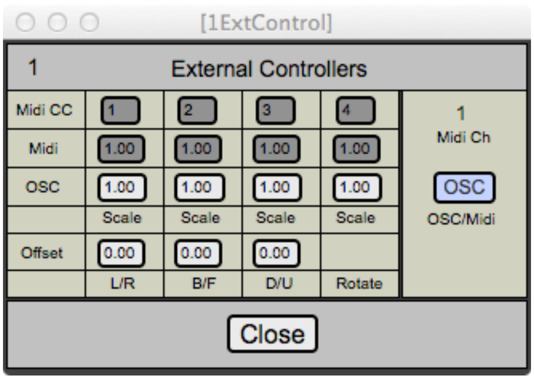

The “OSC/Midi” switch in the right-hand panel selects whether the panner is to be controlled by OSC or Midi Controller. The settings for the one not chosen are greyed out.

The x, y, z, values can be scaled, that is multiplied by numbers. This can be used to define the maximum distances of the movements of the sources; if these are not greater than 0 your source will not move anywhere. Each Controller for L/R, B/F and D/U can be scaled by a factor between + and – the maximum scaling. A value of “1” allows the source to move to the distance of the loudspeaker array (assuming that the incoming value varies between -1 and +1). The maximum scaling on each axis is half the dimension of the virtual room defined in the “V-Space” window (see below). Negative values invert the incoming Midi controller value.

Similarly, the x, y, z, values can be offset from zero to confine movement to a given volume of the entire 3D space.

Midi In

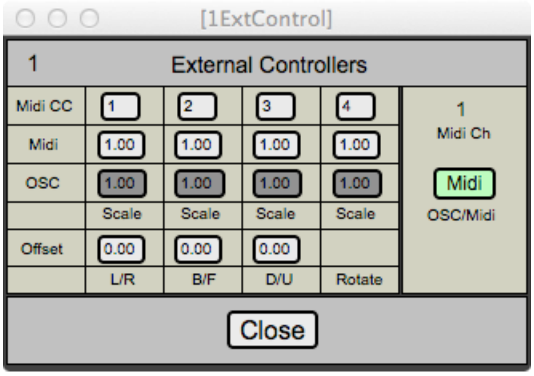

For the sake of simplicity Panner 1 is controlled by controllers on Midi channel 1, Panner 2 by controllers on Midi channel 2 etc. up to Midi channel 16. Each channel will need to have its Midi controller information set up separately.

The “Midi CC” boxes set the continuous Midi controller which will control that parameter..

There is a controller for each of the three axes (L/R, B/F and D/U) and for “Rotate” for each channel. To avoid confusion it is suggested that you use the same controller number (0 to 127) for the same axis for each Channel. In the above example Controllers 1, 2, 3 & 4 have been used.

Note that some programs and hardware controllers start their Midi Controller numbers at zero (0-127), and others at one (1-128).

Midi Controller values vary between 0 and 127. A value of 63.5 would then position the input source at the centre point of the relevant axis.

The “Scale” parameters can be used to define the maximum distances of the movements of the sources; if these are not greater than 0 your source will not move anywhere. Each Controller for L/R, B/F and D/U can be scaled by a factor between + and – the maximum scaling. A value of “1” allows the source to move to the distance of the loudspeaker array. The maximum scaling on each axis is half the dimension of the virtual room defined in the “V-Space” window (see below). Negative values invert the incoming Midi controller value.

The “Rotate” control can be scaled at values between -10 and +10.

“Offset” allows each Controller to be offset by a factor between + and – the maximum scaling value. A value of “1” on a given axis offsets the centre point of the source to the positive edge of the loudspeaker array on that axis. Negative values offset the centre point in the opposite direction. Thus individual channels movement can be constrained to lie within a defined area.

There is no offset for the “Rotate” control.

NB *When setting “Scale” and “Offset” parameters it may help to view the 3D Display.*

The “Ext” button in the “Panners” window must be set to “On” to enable external OSC/Midi control.

4.1.5 3D Display Control

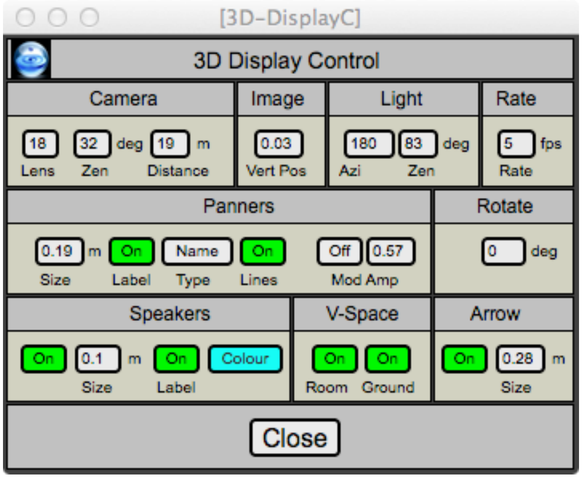

The user can change the vertical position, lens angle and distance of the “Camera”. You can also move the vertical position of the whole image and rotate it, change the position of the lighting, and the “Frame Rate”.

The size of the panners (input sources), speakers and green arrow, can each be adjusted globally.

The Speakers, the “V-Space” “Room” and “Ground”, and the “Arrow” can be shown or hidden.

The position of the “Arrow” represents the Listener, and it and the loudspeakers will move to match that position.

The Panners” “Label” (its number or name can be chosen) can be shown or hidden. “Lines” from each input source to the central listening position can be shown or hidden.

The “Mod Amp” controls affect the brightness of the panner images with the level of the sound source. The “On/Off” button enables/disables this. Number box next to this sets the gain globally. One is maximum gain and zero is off (the panners will remain dark).

The size and colour of the speakers are adjustable, their label numbers shown or hidden.

All settings here are stored in the “3DAS” Memory when this window is closed.

4.1.6 File Paths Window

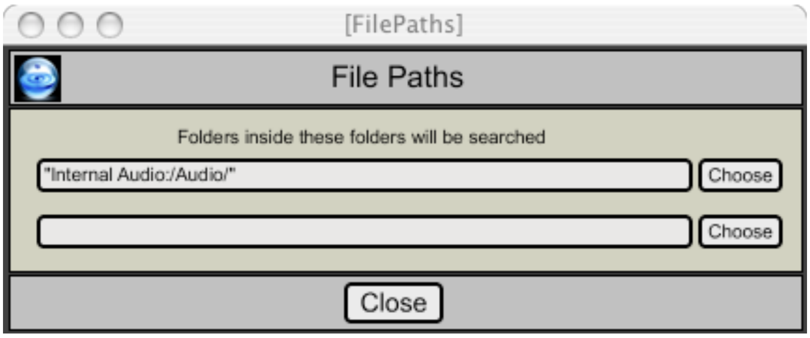

3DAudioScape has been built using Max, from Cycling 74. Max-built applications usually have restrictions on the locations from which files can be loaded. By default, files, particularly audio files, can only be loaded from a folder that contains the application, or from sub-folders within that folder. This window allows the user to select two further folders from which files can be loaded.

Clicking on “Choose” opens a dialogue box where a folder can be chosen for each slot. When chosen the name of the folder appears in the name box.

When this window is closed the settings here are stored in the “3DAS” memory.

4.1.7 B-Format Players Window

NB. *This is only available in the ambisonic version of 3DAS.*

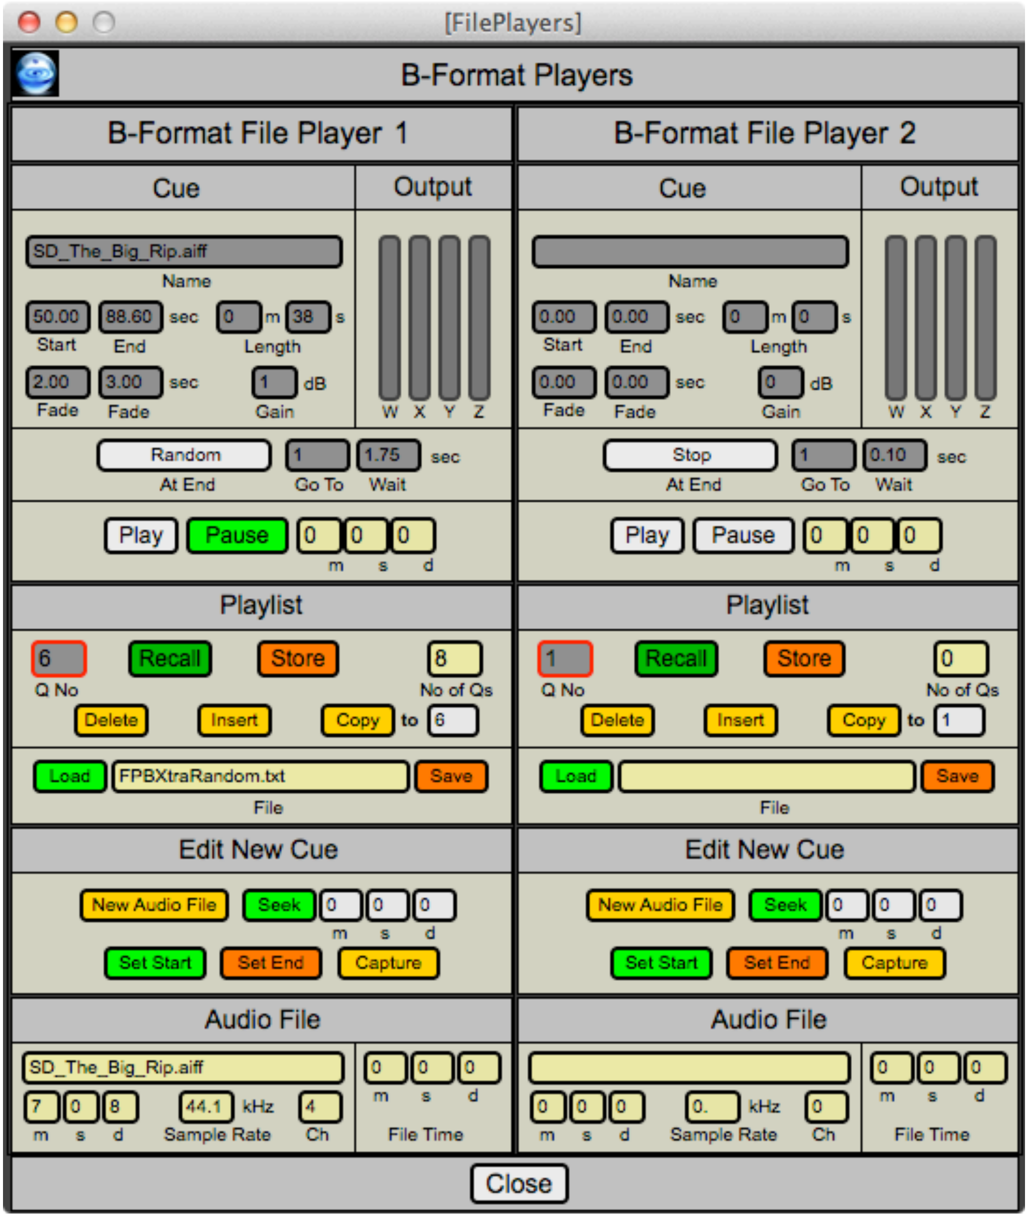

Two programmable B-Format File Players are provided.

Each Player can play any interleaved sound file of up to four channels, but the spatial results will not make much sense unless it is a ambisonic B-Format encoded file. A stereo file will be treated as a UHJ encoded file, and be encoded to horizontal B-Format.

Audio Files can be placed in, or copied to the same folder as 3DAudioScape, so that the Players can find and play them. Otherwise the “File Paths” window (See Sec 4.1.6) can be used to specify up to two folders containing audio files, and these files can then be loaded and played.

The B-Format Players can playback files recorded with the B-Format Recorder.

A large number of “Cues” (128) can be stored in each player. A “Cue” stores the audio file to be played, the start and end times within that file, a fade in and fade out time, the playback level, and what happens at the end of the “Cue”.

A “Playlist” file can be loaded into each player. These are text files. A programmed “Playlist” can be saved to disc as a .txt file with the “Save” button, and loaded with the “Load” button in the “Playlist File” panel. The name of the loaded file is displayed in the yellow box.

By default a “Playlist” with the name “Empty.txt” is loaded into each player. It contains no programmed cues, and is just intended to initialise the player.

In an “Empty” Playlist

i) The grey “Q No” number box with a red outline in the “Playlist” panel will be set at “1”. The “No of Qs” will be set at “0”.

ii) Click on “New Audio File”. This opens a dialogue window where the file to be loaded can be located and selected. A “Cue Name” dialogue box opens, containing the file name and suffix (.aif, or .wav) of the file e.g.

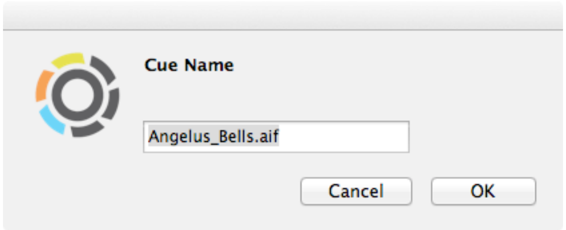

This name can be edited as desired and clicking on “OK”, or pressing “Return”, will store the “Cue”, and it will start playing.

In the “Audio File” panel the file name, its duration, its sample rate and number of channels will be shown.

iii) The file is also loaded into a temporary buffer. In the “Edit New Cue” panel clicking on “Seek” will play the file from the start. The “File Time” boxes in the “Audio File” panel will show the playing time. By setting the numbers in the “m” (minutes), “s” (seconds), and “d” (deciseconds) boxes next to “Seek” and pressing “Seek” again, the time from which you wish the cue to play can be found. Pressing “Set Start” will copy this time into the “Start” box in the “Cue” panel. Similarly an end point for the cue can be found using the “m”, “s”, and “d” boxes and “Seek” button, and copied into the “End” box in the “Cue” panel by pressing “Set End”. Note that the file may continue to play past the “End” time but not be heard.

“Store” the modified “Cue” again with these Start and End times. It can also be re-named in the dialog box that appears.

iv) The “Fade” times at start and end and the “Gain” can be set now if desired, or at a later time. Alter them and “Store”.

v) The Cue can be played and paused using the “Play” and “Pause/Resume” buttons. The Length of the cue is shown in the Length boxes, and the yellow boxes next to “Play” and “Pause” show the running time of the cue. Note that this may differ from the time shown in the “Audio File” panel, which shows the time in the file itself.

vi) Other cues can be programmed at other locations. These players are designed so that there are no empty locations. Having stored “Cue 1”, the Playlist “Q No” can be advanced to 2. Go to step ii) and prepare “Cue 2”. Repeat as many times as necessary.

vii) To recall a cue the “Q No” box in the “Playlist” panel is scrolled to a number and the mouse released, or by clicking in the number box typing a number and pressing Return or Enter. Then press “Recall”. The Cue can then be edited by altering the start and end times, fade in and out times and gain, and then “Stored” again.

viii) The “Playlist” can be automated by programming what happens at the end of each Cue. Options are: Stop, Next, Next Wait Play, Go To, Go To Wait Play, Random, Random Wait Play, Repeat, Wait Repeat.

‘Stop’ is the default.
‘Next’ advances the “Q No” to the next Cue awaiting “Recall”.
‘Next Wait Play’ advances the “Q No” to the next Cue, and “Recalls” it after the “Wait” time.
‘Go To’ goes to the programmed “Q No” awaiting “Recall”.
‘Go To Wait Play’ goes to the programmed “Q No”, and “Recalls” it after the “Wait” time.
‘Random’ and ‘Random Wait Play’ behave similarly, choosing a programmed Cue at random.
‘Repeat’ and ‘Wait Repeat’ repeat the Cue endlessly, the latter with a timed gap between repeats.

“Recall” a Cue, choose the desired behaviour and then “Store”.

ix) Cues can be deleted, inserted, cleared and copied to another location.

NB *"Delete", "Insert" and "Copy" require the audio to be turned off to allow the File Players to be re-programmed. This happens automatically and it is turned on again when the process is complete.*

4.18 Windows Window

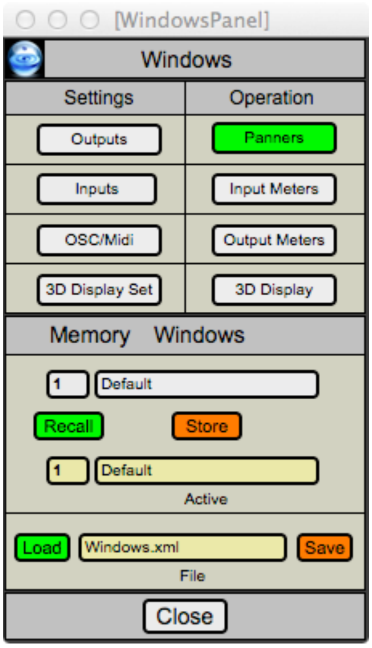

The most useful windows can be opened and closed by clicking on the named buttons. Any opened windows can then be dragged around the screen to desired positions. The result can be stored to one of ten memory locations. The white number box above "Recall" is used to set the number of the memory location. "Store" will then store it, with an option to name, or rename, the memory location.

To recall a stored memory, scroll the number box above "Recall", or type a number into it and press return. The name of the memory location will be displayed in the text box next to the white memory number box. Clicking on "Recall' will then recall that memory, and its number and name will be displayed in the yellow boxes below.

The entire contents of the memory can be saved to disc by clicking on the "Save" button. The file can be renamed, and its storage location chosen before it is saved as an .xml file. By default a file named "Windows.xml" (the same name as the Memory) will be loaded when 3DAudioscape is activated, and memory location "1" recalled. This should be present in the same folder as the 3DAudioscape application. It can be overwritten with your own preferred version with the same name, or stored as a version with a different name.

4.19 Other Versions of 3DAudioScape

There are currently three other versions of 3DAudioscape: “3DAudioscape Bin”, “3DAudioscape AEP”, and “3DAudioscape Delta”. They implement different spatial audio algorithms than the first-order ambisonics used in 3DAudioscape.

All versions are designed to use the same control pages, so all appear nearly identical, though they control different spatial audio engines. Furthermore, they are designed so that .xml files made with one version can be loaded into another version and produce an effective and similar result, which may require only minor modification. Thus, any spatial programming done in one version can be passed into another version to reproduce it using a different spatial audio algorithm.

Various control sections have been hidden, where they are not appropriate. For instance the binaural “3DAudioscape Bin” version uses only the two loudspeakers of a set of headphones, so the ”Outputs” page where multiple loudspeakers are set up, and the multichannel output meters have no effect, and are not displayed and. Any such hidden controls maintain their settings, so that any states imported from another version are loaded and can be re-saved in a “3DAS.xml” file.

Similarly, any per channel ambisonic order controls set in the “3DAudioscape AEP” version will be loaded into any of the other versions (where they have no effect) and will be resaved in any modified “3DAS-Panners.xml” file.

3DAudioscape Bin

“3DAudioscape Bin” implements direct binaural coding of mono and stereo audio sources, with a 3D first-order ambisonic reverb. This produces a better and more precise binaural result than that provided in “3DAudioscape”, which binaurally encodes the outputs of virtual loudspeakers.

No audio file players are presently provided, neither is a direct binaural input. This is largely because this was originally envisaged as a binaural encoding mixer for production. Audio files that are already binaurally encoded can be mixed in any standard DAW. Unlike ambisonic B-Format audio streams, which can be spatially manipulated, binaural encoding cannot be undone and then redone without serious effects on the quality of the sound, so the directional qualities of the stream cannot be altered.

In the “Panners” window, the input and output B-Format control panels are hidden, as they have no meaning for the binaural processing.

On the Main page, the level meters show the main output levels, and there is a master level control. The Head Radius control is on the Main page, and there is also a Decorrelation control. Decorrelation in headphone listening reduces the similarity of the signals to each ear. This reduces the effect of sound appearing to be “inside the head”. It has some similarity to short reverberation. There is an “On/Off” button, and a control for the amount, values between zero and one. One is definitely excessive, and the most useful values are less than 0.5.

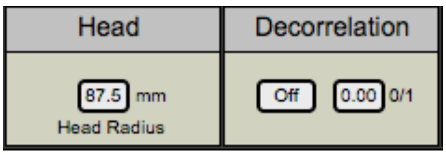

3DAudioscape AEP

“3DAudioscape AEP” implements direct Ambisonic Equivalent Panning of mono and stereo audio sources, with ambisonic order up to fifth-order, and a 3D first-order ambisonic reverb. The ambisonic order can be varied on each panner. Ambisonic Equivalent Panning combines encoding and decoding to produce loudspeaker feeds. There is no intermediate B-Format signal. By this means higher order ambisonics, and more precise spatial location, is achieved with less signal processing.

It appears similar to the ambisonic version “3DAudioScape” but, since there is no B-Format signal path, all sections controlling that have been removed.

On the “Panners” page there is an extra control to determine the ambisonic order of each channel.

This order is continuously variable between 1.0 and 5.0. Strictly speaking, higher orders require a larger minimum number of speakers, sixteen for third-order 3D, thirty-six for fifth-order 3D, for example.

Generally the minimum number of speakers required is
$N = (m + 1)^2$ for 3D
$N = 2m + 1$ for 2D
Where N is the number of speakers, and m is the ambisonic order.

However, the extra spatial precision may be useful with smaller numbers than optimal, though the panning will not be as smooth, with holes appearing between loudspeakers, and the loudspeakers becoming more dominant.

3DAudioscape Delta

“3DAudioscape Delta” uses ampltude/delay panning of mono and stereo audio sources, with a 3D first-order ambisonic reverb. The name “Delta” is derived from Delta Stereophony, proposed by Steinke and Wolfgang Ahrnert and others in East Germany in the1980’s for use in theatres and large multi-purpose halls. It is broadly similar to the algorithms used in “TiMax”, and to those involved in Wave Field Synthesis, though it uses far fewer loudspeakers than required for WFS. It would be more suitable for more irregular arrays of loudspeakers than those required for ambisonics.

It appears nearly identical to “3DAudioscape AEP” though, of course, there is no control for ambisonic order.

For each panner, sixteen signals loudspeaker signals are produced, the amplitude and time delay of each signal determined by the distance of the sound source from each speaker. The outputs of the panners are then mixed together for each speaker.

The amount of the inverse square law (“Inv Sq” affecting amplitude) and “Delay” are variable between zero and one (mathematically correct) for each panner. Thus, a wide variety of spatial behaviour is available.

5 Using 3DAudioScape with a DAW

Soundflower can be used as a bridge to pass audio from one audio application to another. Once installed it is available in a two-channel or a sixteen-channel version. The latest versions for MAC OS 10.8 have replaced the sixteen-channel version with a 64 channel one.

Jack is an alternative and more flexible alternative, although it is more tricky to set up correctly and complex arrangements seem to become fairly processor intensive.

The general principles are the same for Logic, ProTools 9, Reaper, Digital Performer, Ableton Live and Cubase/Nuendo. It is also possible to run your sequencer on a separate PC as long each has a suitable multi-channel audio and Midi interfaces.

First set up the default Audio Driver using Audio MIDI Setup. This is done to ensure that different audio applications can be set up to use their own driver, rather than all attempting to use the same audio interface.

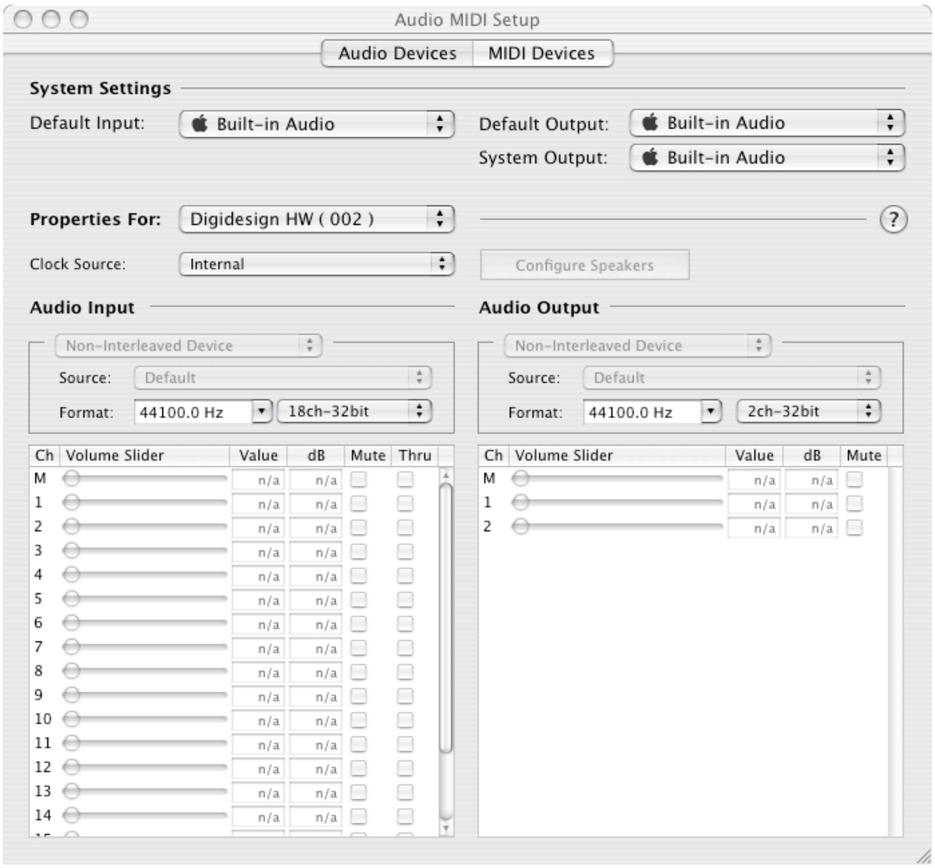

Choose Built-in Audio for “Default Input” and “Default Output”.

5.1 3DAudioScape Set-up

3DAudioScape Audio Interface Set-Up

Press the “Set-up” button (top left) to set up the audio driver.

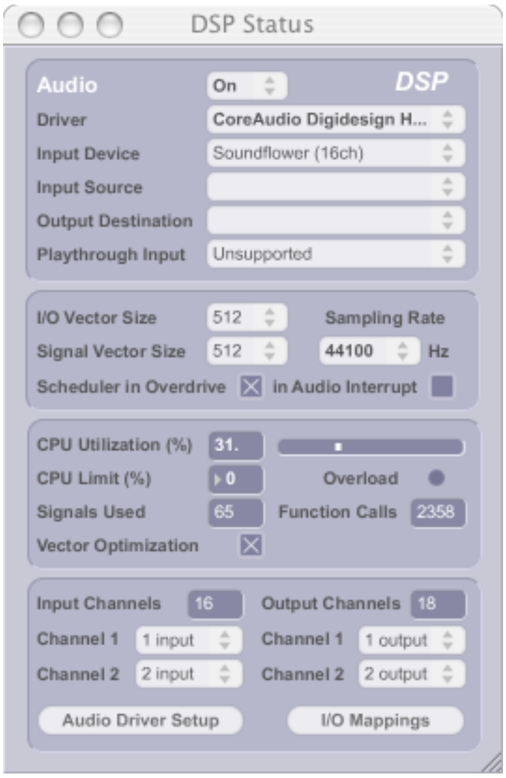

Select your audio interface as your “Driver”. Here a Digidesign audio interface is used.

Select Soundflower (16 channel) as the “Input Device”.

The “Sampling Rate” and “I/O Vector Size” (i.e. buffer size) should match the settings in your DAW and audio interface driver.

The I/O and signal vector size boxes should be set to avoid audio glitches or CPU overloads.

Press the “I/O Mappings” button on the bottom of the “Set-up” panel. Each channel output by your sequencing software needs to be input into 3DAudioScape. Switch on as many inputs as you have channels playing in your DAW. Close both windows.

Midi Set up

There are two main ways of controlling 3DAudioScape using Midi.

i) Sending Midi CC (Continuous Controller) messages. Each audio track in a DAW will have to have a corresponding Midi track. This track will contain Midi CC data for the four coordinates (L/R, B/F, D/U and Rotate).

ii) Sending Mid Program Change messages. These recall states of the “Panners” window.

The former leads to a lot of Midi data to be entered and maintained.

NB
Some DAWS (notably Logic) do not output intermediate values of Midi CC between automation points, leading to erratic and discontinuous movements.

The Midi Program Change method is much simpler to operate and leads to smoother movements. It requires only one Midi track.

Set up your Midi system using “Audio Midi Setup”.

Audio MIDI Setup

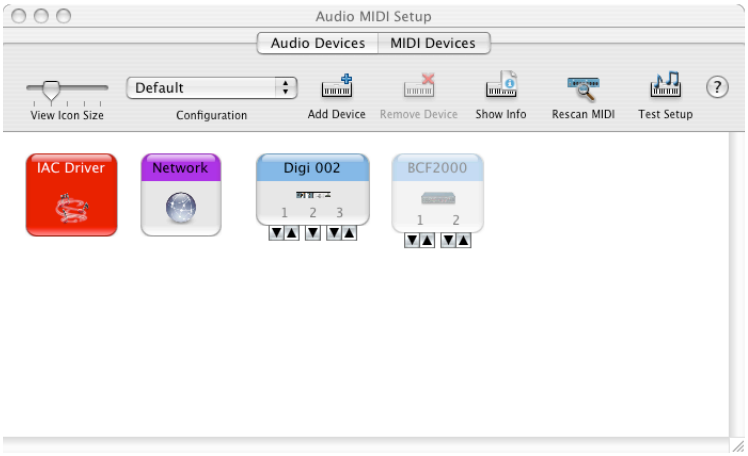

Here a very simple Midi set up is shown. We are going to use the IAC Driver to send Midi data from a DAW to 3DAS. Double click on the IAC Driver icon.

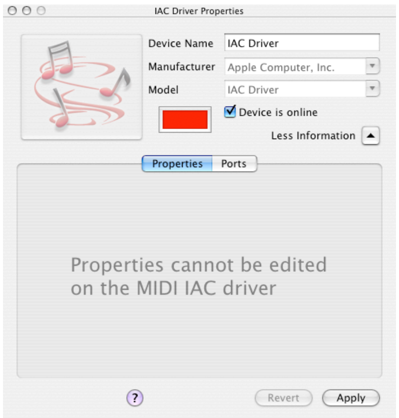

Select “Ports”.

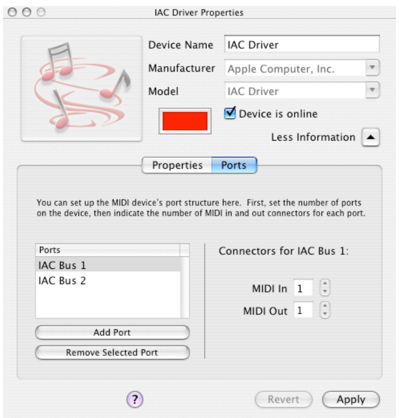

Here the first port has been re-named “IAC Bus 1” by double clicking on the name and re-typing. A second port has been added by clicking on “Add Port”, and it has been been re-named as “IAC Bus 2”. Each port can handle 16 Midi channels.

5.2 Using 3DAudioScape with Logic

Logic does not have a very good architecture for surround sound mixing other than the Dolby N.1 (N = 5, 7 etc.) formats. It can only host the most basic ambisonic plug-ins.

Hence it can only really be used for sophisticated ambisonic work by using it as a multi-track, multi-output playback device. Its audio outputs feed the audio inputs of 3DAudioScape. It can then also send Midi to 3DAudioScape to control the "spatialisation" of those outputs.

Logic Audio Set Up

Logic 8 is used as an example here. It is assumed that the user is familiar with the audio and Midi facilities available in Logic.

Launch the Logic program. Set its" audio driver to Soundflower (16 channel).

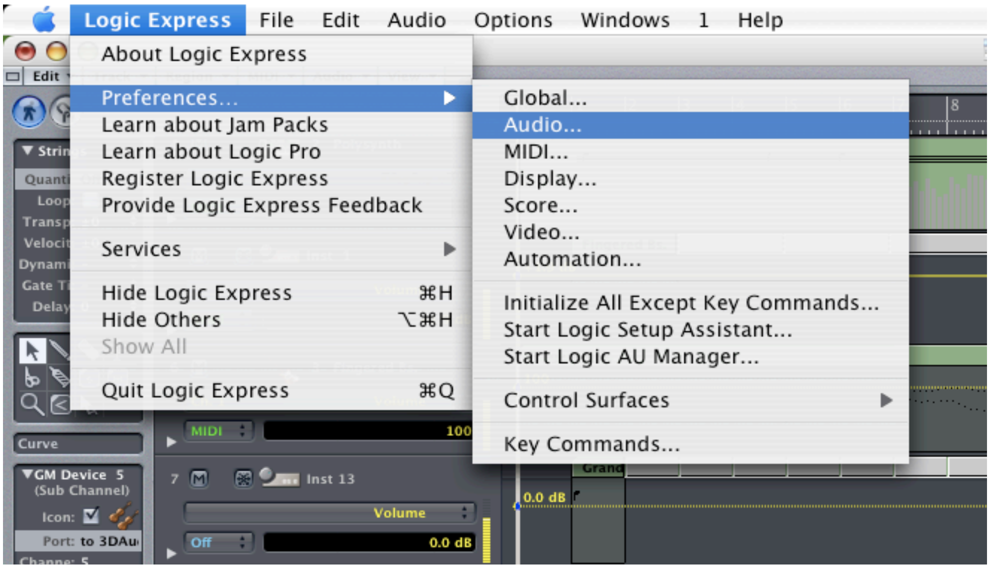

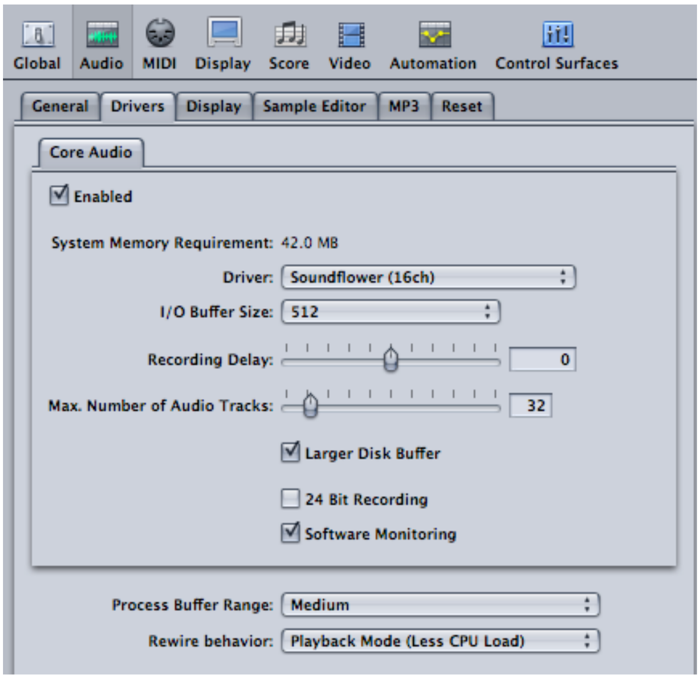

The "I/O Buffer Size" should be set to the same value as the "I/O Vector Size" in 3DAudioScape.

Logic Midi Set Up

Open the “Environment” window, Window/Environment or Cmd-8.

Environment

Select “MIDI Instr”” from the menu at the top left of the page.

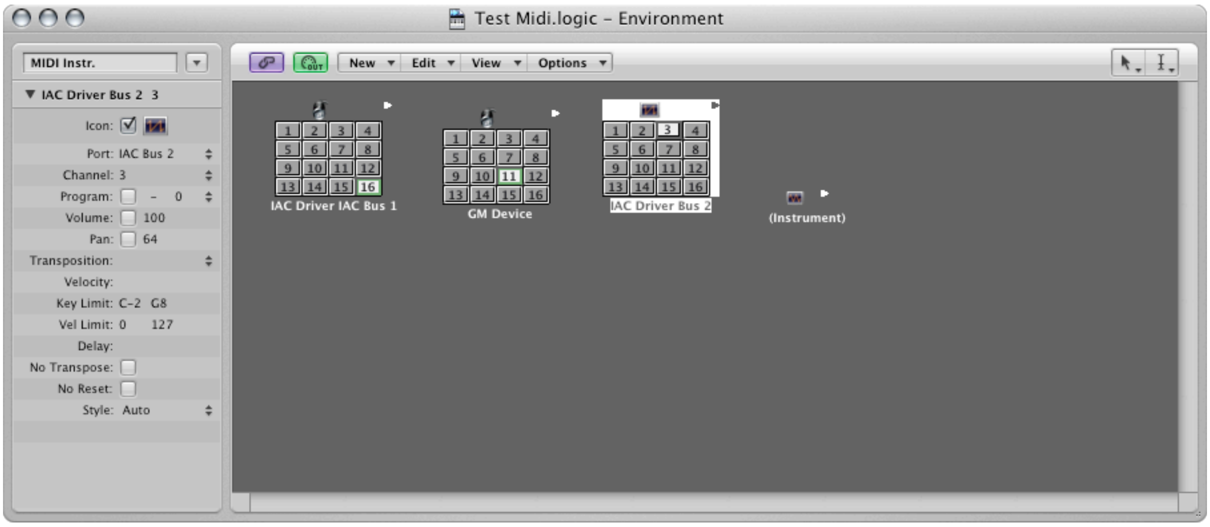

Here, two new Multi Instruments have been created from the “New” menu. When each is clicked on its settings are shown in the Inspector panel at the left hand side.

A new Multi Instrument is shown below.

Double clicking on the icon at the top of the Multi Instrument opens the following window, where the instrument can be named

By default a new Multi Instrument shows all 16 Midi channels with a diagonal line through them, showing that they are disabled. Click on each one. It will become uncrossed, and its settings are shown in the Inspector.

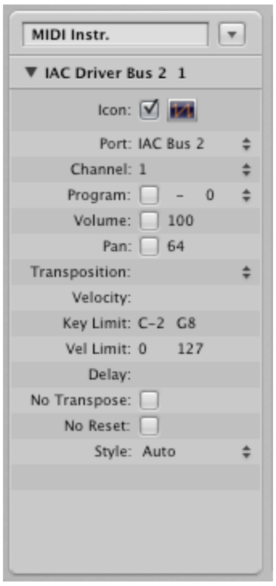

Clicking on the Icon next to the check box at the top of the Inspector allows selection of an alternative icon to the default one.

The Inspector shows that we have selected channel 1 on a Multi Instrument named IAC Driver Bus 2. We will use this to pass Midi Controller data from Logic to 3DAudioScape. We use a simple system to control the Panners in 3DAS: Panner 1 uses Midi channel 1, Panner 2 Midi channel 2, etc. We will use the same MIDI controllers for each L/R, B/F and D/U parameter: Control 1-Modulation, Control 2- Breath, and Control 3-Ctrl 3. There are of course others ways to use the Midi controllers, but this method means that the same control number controls the same parameter for each panner.

Set the port to IAC Bus 2 if this has not been done already. Set the Channel to 1, and deselect the Program box. Select each channel on the Multi Instrument icon and set in the same way: Channel 2 to Channel 2, Channel 3 to Channel 3 etc.

Now select “Click & Ports” from the menu at the top left of the page.

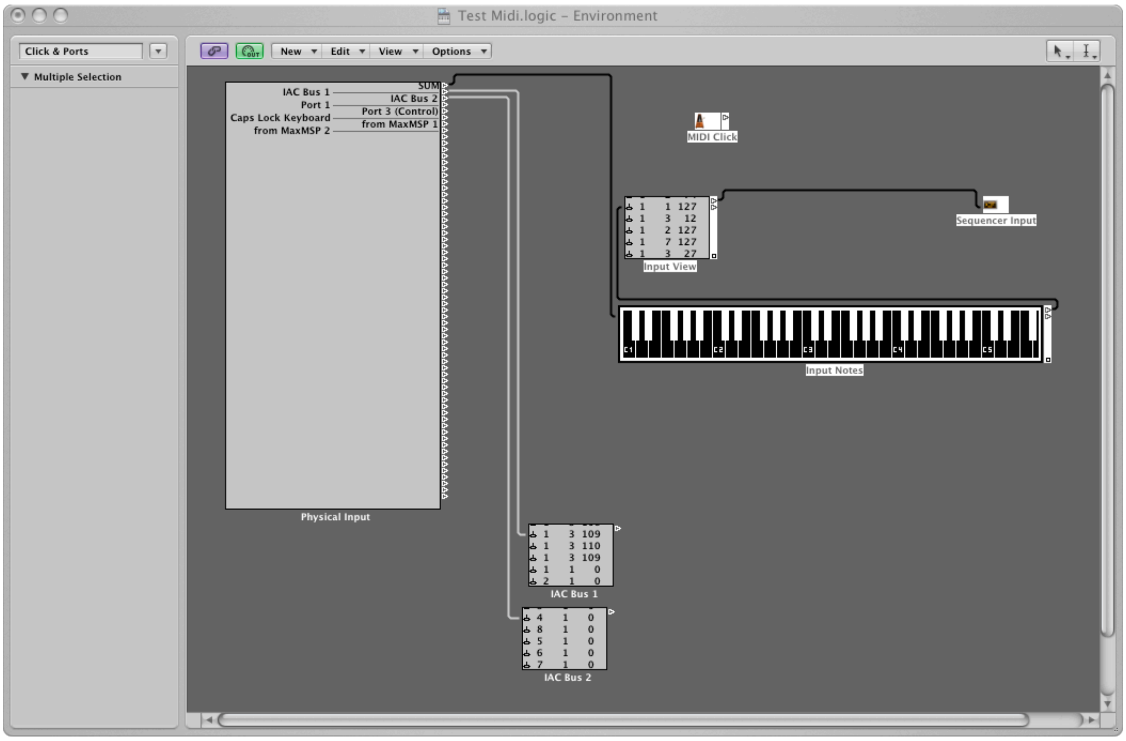

Curiously, by default Logic routes the outputs of the “IAC Driver Bus 1” and “IAC Driver Bus 2” Multi Instruments to the “Physical Input” and back into the “Sequencer Input”, creating a Midi loop, which causes all sorts of problems.

The loop can be broken by creating two new Monitors from the “New” menu, and connecting the outputs of IAC Bus 1 and IAC Bus 2 to them, as shown above.

Logic Arrange Page

Create new Midi track. Track/New

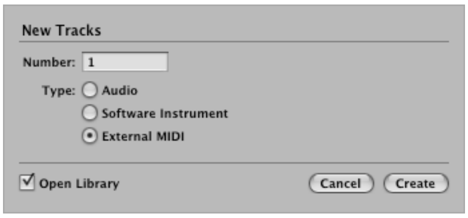

Turn on View/Track Automation.

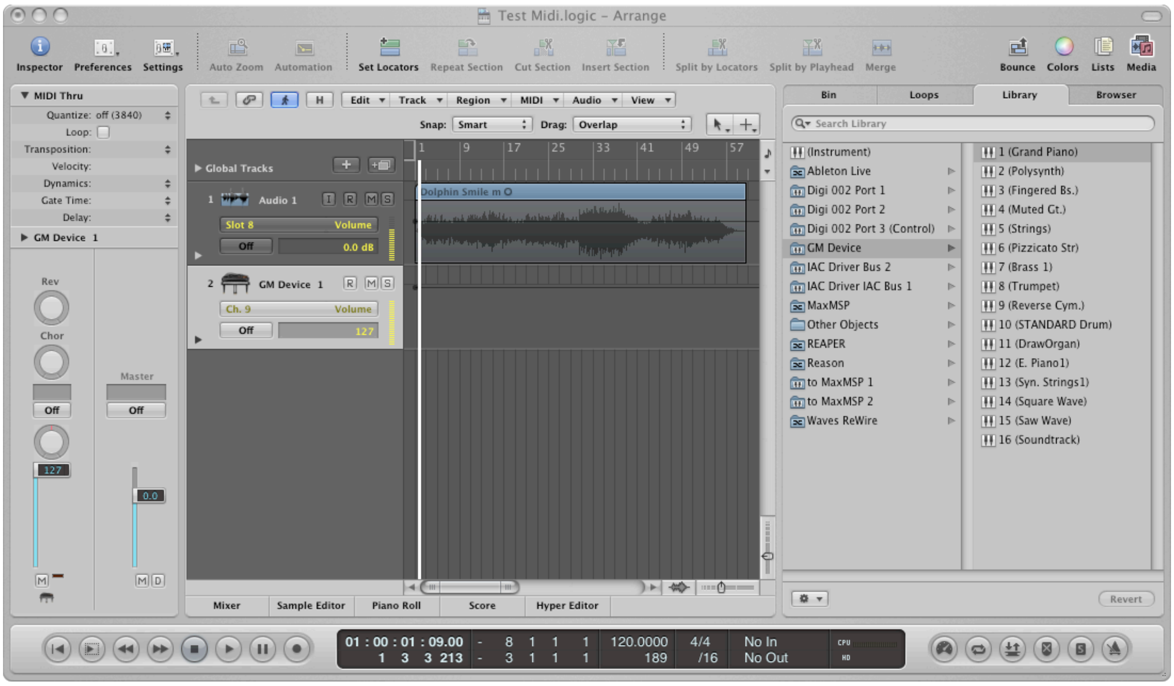

Ctrl/click track icon. Reassign Track Object. Midi Instr/IAC Driver Bus 2/Channel No 1, or select from the Library Panel on the right hand side.

Click on Ch.- display and choose Ch. 1, and Midi Control 0-63. Controller numbers should be chosen to match those set in 3DAudioScape. (See Sec. 4.1.4)

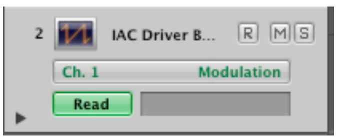

Track automation data can be recorded from incoming Midi, or written using the Pointer tool.

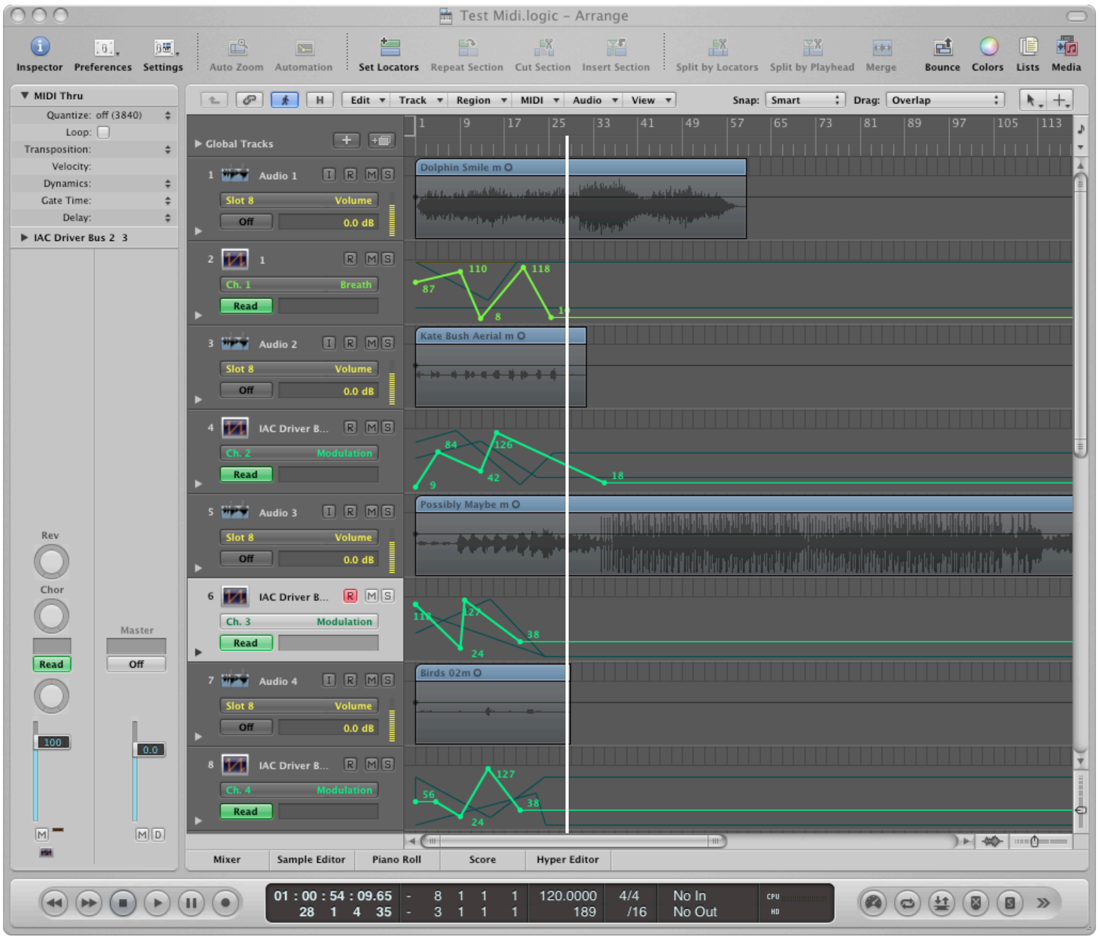

The picture above shows four mono audio tracks and four Midi tracks that control the movements of the associated panners in 3DAudioScape.

Audio Routing

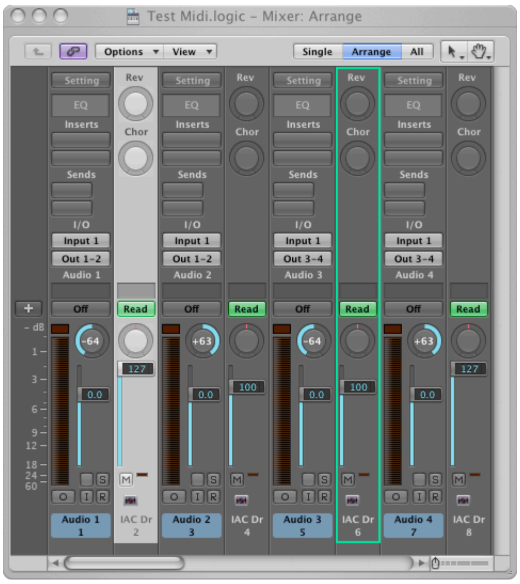

In Logic the outputs should be set to the desired Soundflower channel.

The audio routing is shown in the Mixer page, i.e. Audio 1 and Audio 2 to Output 1-2 and Audio 3 and Audio 4 to Output 3-4, panned so that each audio track has its own output through Soundflower.

Using Program Change Messages

Arrange Page

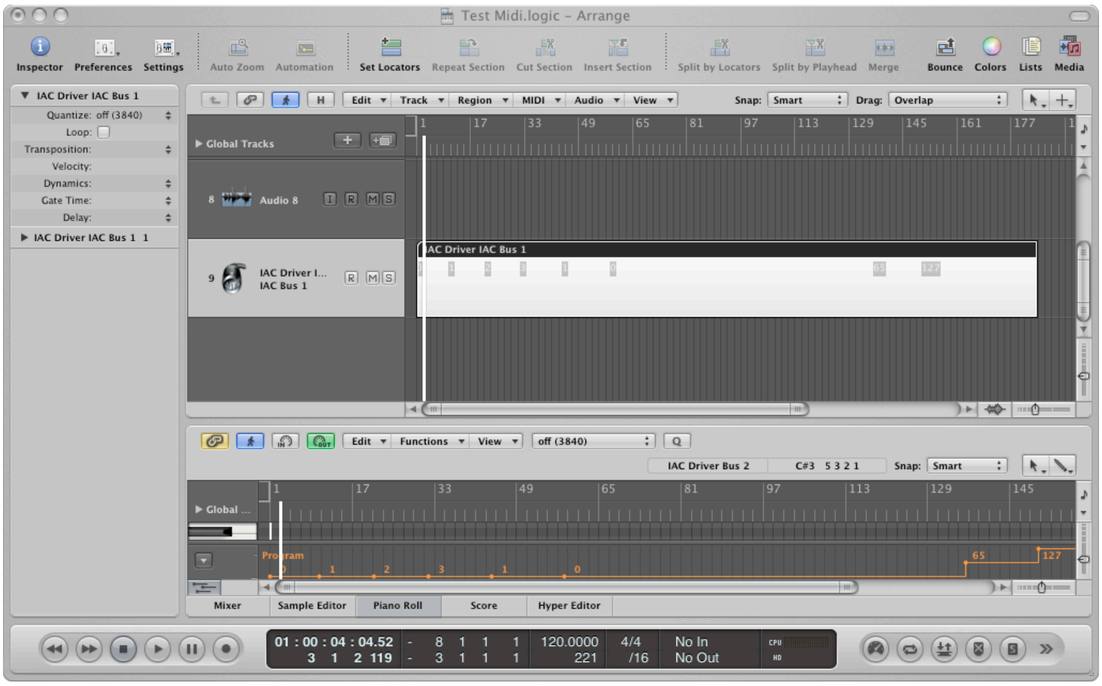

Event List Window

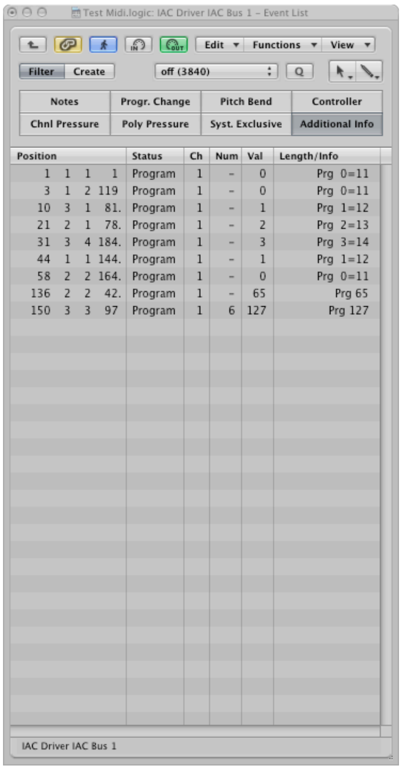

5.3 Using 3DAudioScape with Reaper

Reaper is a very affordable and flexible DAW. Its routing possibilities make it an ideal host for ambisonic sound production. Though not natively supporting ambisonics, it can use ambisonic plug-ins. These are available from a variety of sources

Bruce Wiggins" Wigware plug-ins
http://www.brucewiggins.co.uk/?p=95

Daniel Courville"s B2X plug-ins
http://www.radio.uqam.ca/ambisonic/b2x.html

The ICST Pluggo plug-ins
http://www.icst.net/research/former-projects/projectsarchive/

The Amb set of Pluggo plug-ins from Dave Hunt.

These can be used to produce a B-Format output that can be routed to the B-Format input in 3DAudioScape to use its decoder. Automation of moving sources is made considerably easier, and the output can be bounced to disc as a file playable in 3DAudioScape.

Audio Device Setting

Reaper/Preferences

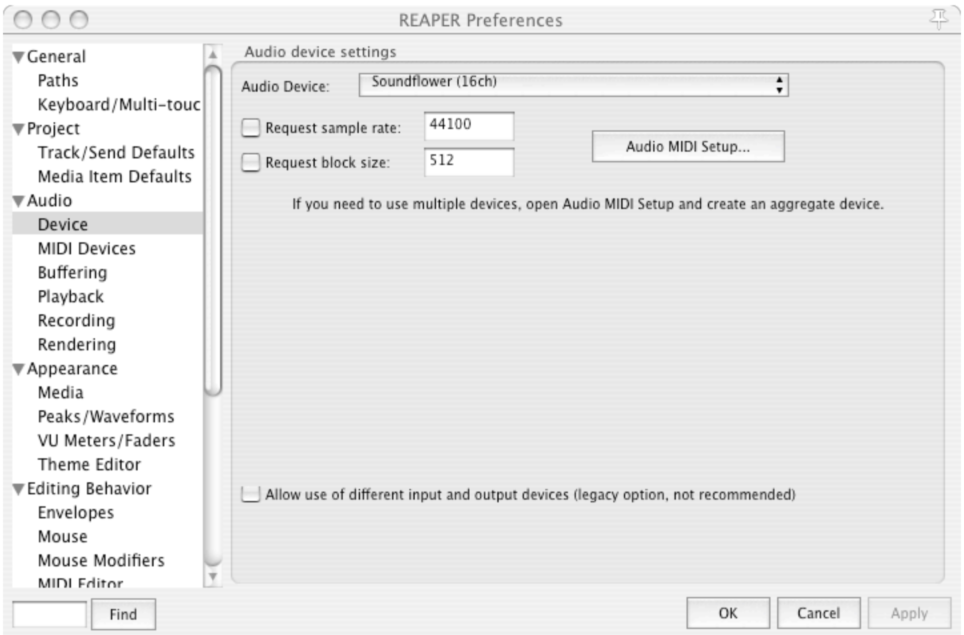

Select Soundflower as the "Audio Device".

Midi Devices Setting

Both methods of controlling 3DAudioScape with Midi are described. In practice the method using Midi Program Change messages is preferable.

As before we are going to use IAC Bus 1 for Midi Program Change messages and IAC Bus2 for Midi CC messages.

Reaper/Preferences

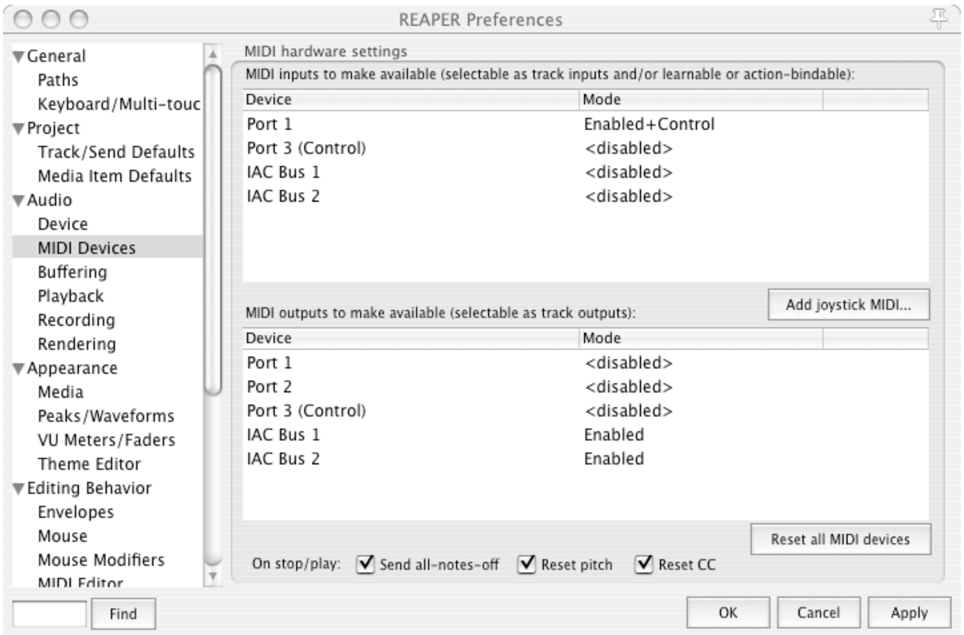

Double click on “IAC Bus 1, Mode”

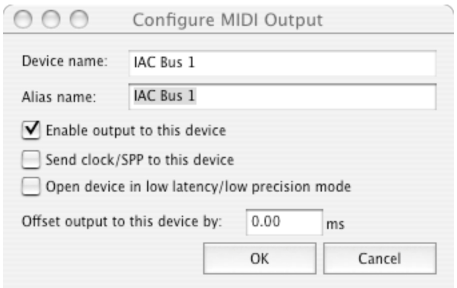

Similarly for on “IAC Bus 2, Mode”

If you wish to record incoming Midi Controller Data, double click on the Input Port Mode that you will use for this (in this case “Port 1”).

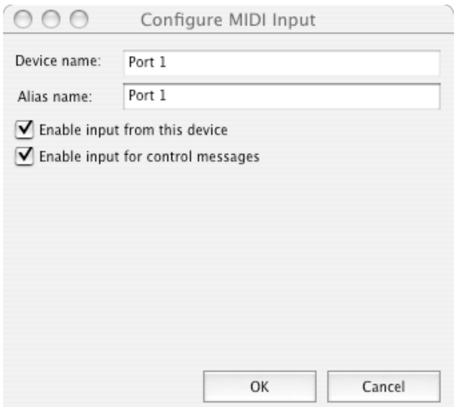

Reaper Arrange window

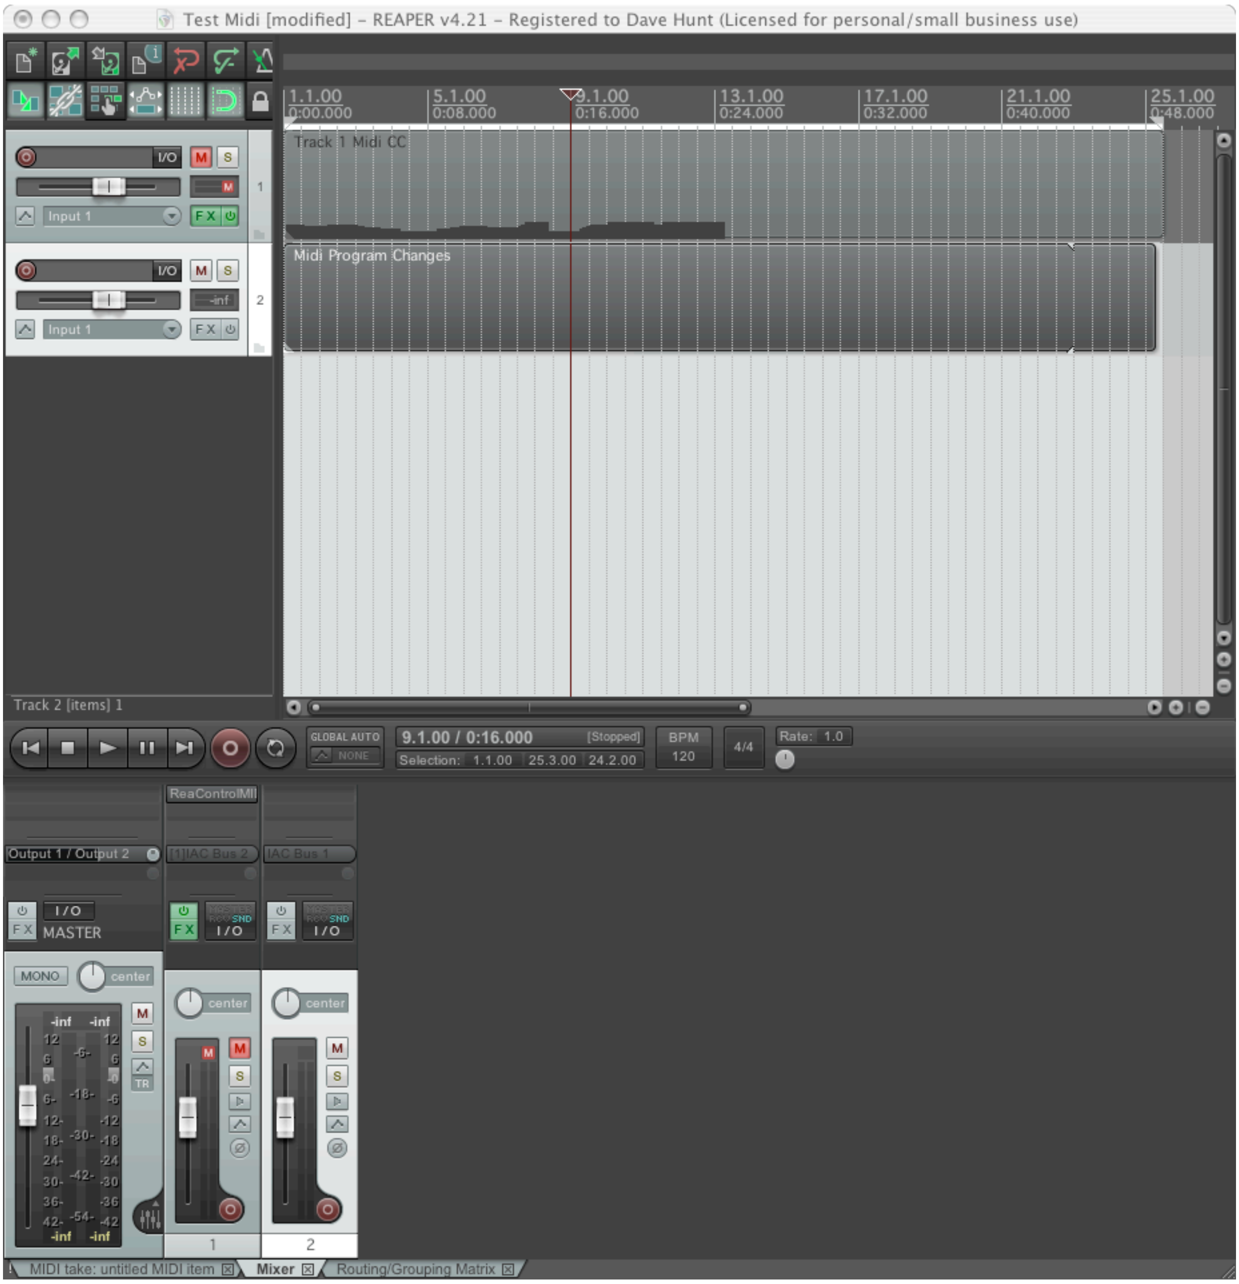

Add two new tracks. Create a new empty MIDI item on each track. To do this, select the required track and (optionally) make your time selection to define the length of the item. Then choose **Insert, New MIDI Item** from the main menu.

Select one of these items then select View/Media Item PropertiesTrack

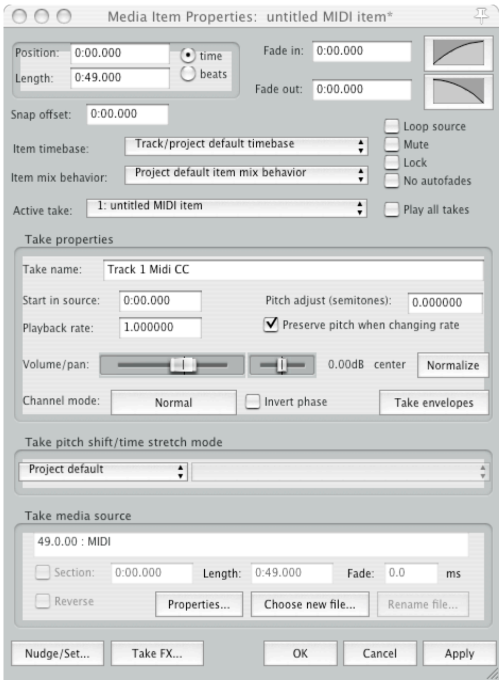

Double click on “Take Name” to label it, and click on “OK”.

Do the same with the Midi Item on the second track.

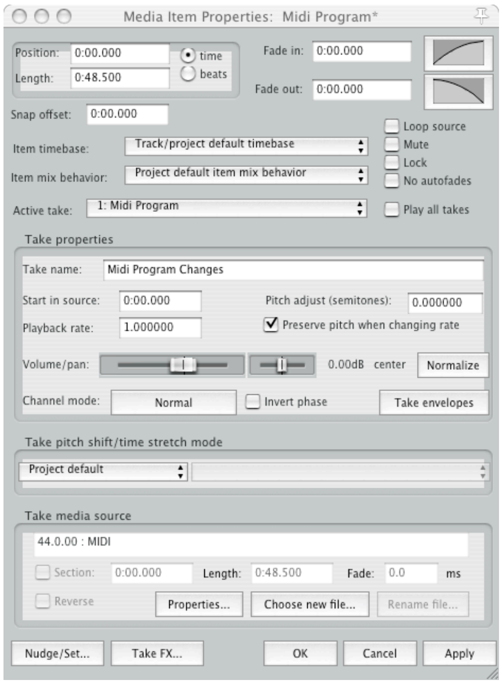

The routing for each track can be made in the I/O control on the Arrange or Mixer panels.

Track 1 is shown as muted, as the intention is to use only Midi Program Change messages.

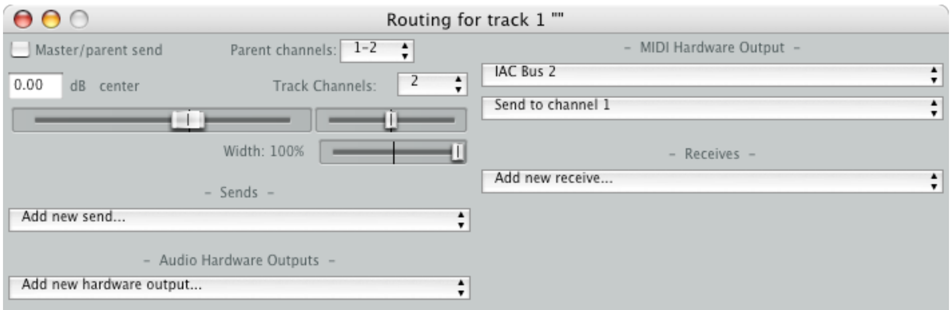

Track 1 is sending Midi CC data to IAC Bus 2 on Midi channel 1, and will control "Panner 1" in 3DAudioScape.

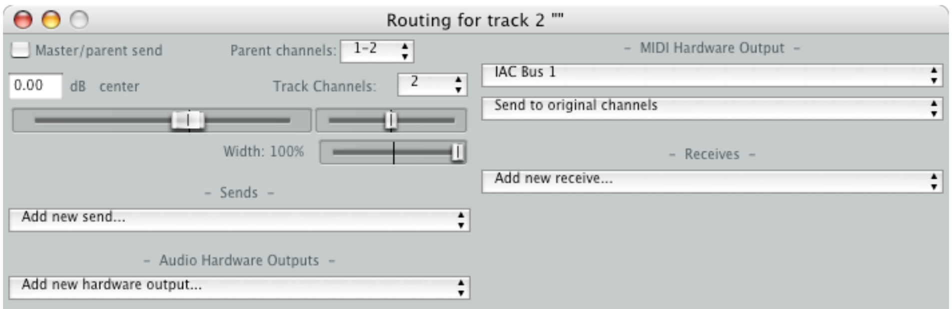

Track 2 is sending Midi Program Change numbers to IAC Bus 1 on the original channels.

Double clicking on a Midi Item opens an editing window.

Midi CC edit

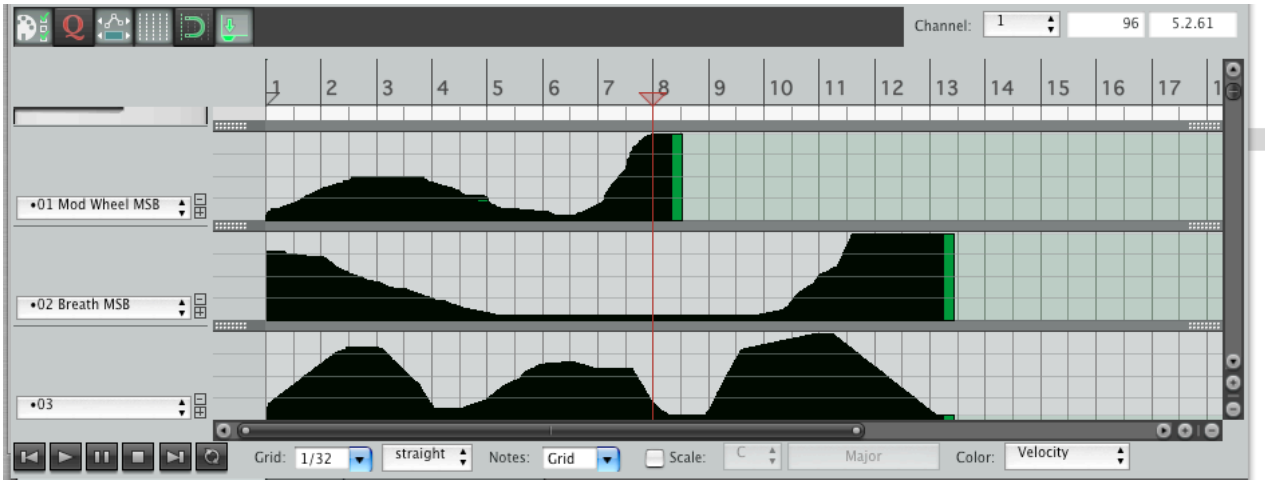

Clicking on the menu for each lane at the left-hand side enables the controllers to be selected. These should match the controllers chosen in 3DAudioScape.

Automation data can be drawn in by and, or recorded from a hardware Mid controller.

Midi Program Change edit

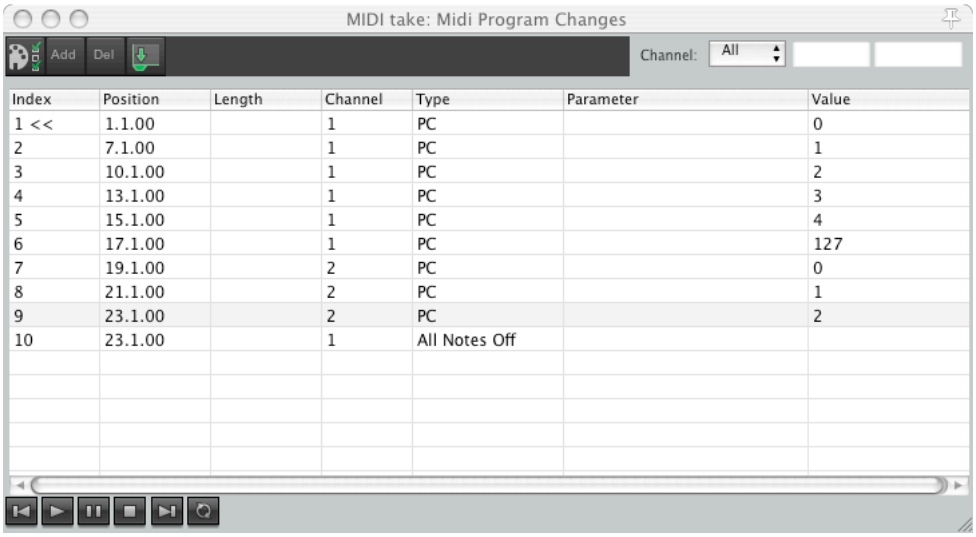

Index	Position	Length	Channel	Type	Parameter	Value
1 <<	1.1.00		1	PC		0
2	7.1.00		1	PC		1
3	10.1.00		1	PC		2
4	13.1.00		1	PC		3
5	15.1.00		1	PC		4
6	17.1.00		1	PC		127
7	19.1.00		2	PC		0
8	21.1.00		2	PC		1
9	23.1.00		2	PC		2
10	23.1.00		1	All Notes Off		

Items can be added to the list with the “Add” button, and deleted with the “Del” button. The position of the event can be edited by double clicking in the “Position” box, and typing. Similarly, the Midi channel, message type and value can be edited.

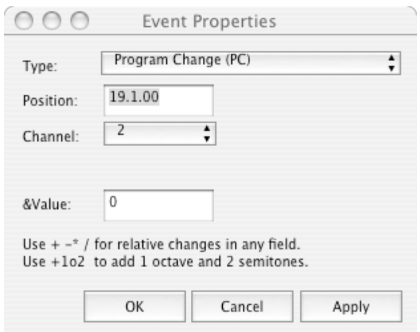

6.0 3DAudioScape Schematic (ambisonic version)