

ai

Nicholas Blachford

COLLABORATORS

	<i>TITLE :</i> ai		
<i>ACTION</i>	<i>NAME</i>	<i>DATE</i>	<i>SIGNATURE</i>
WRITTEN BY	Nicholas Blachford	June 24, 2025	

REVISION HISTORY

NUMBER	DATE	DESCRIPTION	NAME

Contents

1	ai	1
1.1	Aural Illusion v1.1 Manual	1
1.2	Aural illusion Version 1.1 © Blachford Technology 1993-1994	1
1.3	Loading a sample	2
1.4	Stereo Loading/Saving	3
1.5	File formats supported	3
1.6	IFF file type	4
1.7	AIFF file type	4
1.8	AVR file type	4
1.9	SAFF file type	5
1.10	RAW 8 bit file type	5
1.11	RAW 16 bit file type	6
1.12	RAW 16 LSB file type	6
1.13	Other file types	7
1.14	SAFF Specification	7
1.15	The Edit Window	10
1.16	Sound playback	11
1.17	Pull down menus	11
1.18	Sample looping	12
1.19	The Sample Effects	13
1.20	Manipulations	17
1.21	The Synthesizer	19
1.22	Hints and Tips	22
1.23	Synth hints	22
1.24	Filter Sweep	23
1.25	Mixing Samples	26
1.26	Resonant Filter	26
1.27	Phase Shifting	27
1.28	Miscellaneous	28
1.29	Problems	29
1.30	Aural illusion Version v2	30

Chapter 1

ai

1.1 Aural Illusion v1.1 Manual

This is the instruction manual for Aural illusion v1.17

by Nicholas Blachford Updated 12/9/95

Introduction to v1.1

Loading a sample

Stereo loading & saving

File Formats

The Edit window

The Pull-down menus

Sample Looping

Sample Effects

Sample Manipulations

The Synthesizer

Hints & Tips

Problems

Version 2

Miscellaneous

1.2 Aural illusion Version 1.1 © Blachford Technology 1993-1994

Welcome to Aural illusion v1.1 (here in known as Ai)

This program requires Kickstart/Workbench 2.04 or greater and 1 Meg RAM.

More memory is recommended, a hard disc would come in handy as well.

If you are intent on processing large samples an accelerator would be most useful.

To run the program first boot from Workbench then click the program's icon twice.

What is Aural illusion?

Aural illusion is a tool for producing and enhancing sounds, the program accepts sounds as samples which are loaded in from disc, the program cannot record from a sampler, you should use the supplied software for this purpose if you have a sampler.

Ai stores all data internally in 16 bit format and processes it in 32 bit format, this allows for much better quality results even if the samples are then stored in 8 bit format.

Ai can be used for creating a single sample from many samples or many samples from a single sample - very complex sounds can be the result in either case.

This program allocates as much memory as it can find. If you wish to get maximum memory run the program straight after booting Workbench. Booting Workbench from a floppy disc will increase the size of blocks allocated.

The memory found is allocated in three equal blocks, two of these are shown on screen and the third is used for sample processing. All sample data is stored in 16 bit format so this reduces the number of samples available but increases the quality of the effects, a basic Amiga A1200 will get about 250,000 samples per buffer. If you have fast RAM this will be allocated if possible, on some systems both Chip and Fast RAM will be allocated, this will depend on the amount of fast RAM available.

1.3 Loading a sample

There are 2 methods for loading a sample, The quickest is to press one of the load buttons, these are located beside the message window (the large black rectangle) Clicking on 1 will load the sample into the first buffer, 2 loads into buffer 2, Using the buttons samples loaded are mono IFFs.

The second method allows you to load a sample into a range, to select a range click the mouse pointer in one of the sample windows and drag the pointer across the screen with the left hand button held down, when the range you wish to select is displayed in white release the mouse button.

Once this is done select Load from the Project (pull down) menu, a standard file requester will appear allowing you to select your file, press "Volumes" before loading anything first time.

By default all samples are loaded or saved in mono, even if the sample is stereo.

When a sample is loaded the play rate will not change, so some samples may sound different, Play rate is saved however (not in the RAW formats).

For more info see [Stereo loading & saving](#).

A number of percussion sounds have been provided in the file Hits.saff, this file is a mono SAFF file which contains numerous samples which have all been created and processed using Aural illusion, To save disc space the samples are close together so you will need to split them up to judge each sound individually. Due to the Sample display routines some of the shortest samples will not show up unless you have closed in on them so you will need to be careful when splitting them up, There are 31 individual sounds in total in the sample.

1.4 Stereo Loading/Saving

To load in stereo you must select the stereo button which is below the Sample Rate slider, this will load the sample in stereo even if it is mono, One side will be loaded into Sample 1 with the other side put into Sample 2, The same applies for saving samples, the range will be saved with the equivalent area from the other Sample, If you load a Mono sample as stereo it will play at twice the speed, if you load a stereo sample as mono it will play at half the speed and sound slightly different. This can however be used to effect and it is for this reason loading and saving works this way.

The same applies for saving, Select Stereo if you want to save in stereo and the current range will be saved along with the equivalent area from the other side.

Select Mono for saving in mono.

1.5 File formats supported

7 file formats are currently supported:

[IFF 8 Bit](#)

[AIFF 16 Bit](#)

[AVR 16 Bit](#)

[SAFF 16 Bit](#)

[RAW 8 Bit](#)

[RAW 16 Bit](#)

[RAW 16 LSB](#)

[Other formats](#)

1.6 IFF file type

IFF are the standard 8 bit samples used on the Amiga compatible with almost all music software, Ai handles these files with restrictions.

Saves

8 bit

Mono & Stereo

Play rate

Loads

8 bit

Mono & Stereo

1.7 AIFF file type

AIFF is the Apple Macintosh sound format and is also used by several Amiga products, Ai supports these in 16 bit format with restrictions.

Saves

16 bit

Mono & Stereo

Play rate saved as 31.677 Hz

Loads

16 bit

Mono & Stereo

1.8 AVR file type

AVR is a format used on the Atari ST and in the "Clarity" 16 bit sampler on the Amiga, Ai supports these with no other restrictions.

Saves

16 bit

Mono & Stereo

Sample Speed

Loop points

Loads

16 bit

Loop points

Mono & Stereo

1.9 SAFF file type

SAFF is a new sample format designed by the author, reason being that the other formats have restrictions, or are simply unsuitable for transferring 16 bit sample data. SAFFs are not supported by any other software as they have just been invented but they are fast and easy to support.

Ai supports these files with 1 minor restriction.

Saves

Loop points

Sample speed

Mono & Stereo

MSB order

16 bit data

Loads

Loop points

MSB & LSB order

8 or 16 bit

Mono & Stereo

Sample loaded as RAW 16 bit if incorrect type, if a stereo file is loaded as mono the sound will not be changed in pitch or distorted.

The restriction is that Block interleaved SAFFs are not yet supported, some Public domain utilities will be released at a later date to handle SAFFs.

For more information I have included the [SAFF Specification](#).

1.10 RAW 8 bit file type

RAW 8 is sample data stored in 8 bit format with no header file, most sampler software will support loading samples in Raw 8.

Saves

8 bit

Mono & Stereo

No header

Loads

8 bit

mono & Stereo

Ignores any header

1.11 RAW 16 bit file type

Raw 16 is sample data also stored with no header file but this time in 16 bit format, some software will support this but not much, supported by some CD cutters.

Saves

16 bit

Mono & Stereo

No header

Loads

16 bit

Mono & Stereo

Ignores any header

1.12 RAW 16 LSB file type

Raw 16 LSB is the same as RAW 16 bit except that when files are loaded or saved the data bytes are reversed, this is so that PCs can read the files as their data is normally in reverse byte order (Least Significant Byte first) Audio CDs also store their data LSB first but include other coding, supported by some CD cutters.

Saves

16 bit

LSB order

Mono & Stereo

No header

Loads

16 bit

LSB order

Mono & Stereo

Ignores any header

1.13 Other file types

Most sample formats consist of a file header followed by raw data, Using RAW 8, RAW 16 or RAW 16 LSB will load the data but will include the header as part of the sample at the beginning, You should remove this by selecting it then Zeroing it. Then select the rest of the sample and select Rescale which will scale the sample to it's maximum volume (the header will normally confuse the Rescaler).

You may need to try a number of formats before the sample loads properly, try RAW 16 LSB if the sample is completely unrecognizable - PCs and CDs store their data in this format. Also try Stereo and mono loading.

Ai will always load something if you are using one of the RAW formats.

1.14 SAFF Specification

Simple Audio File Format Specification

Revision 1.1 10/3/94

---- SAFF Specification ----

An SAFF (Simple Audio File Format) file consists of a very small 32 Byte header file which holds the basic sample info followed by the samples itself, the samples can be any size you want and can contain up to 255 channels.

The samples can be as many bits as you want but the data is stored in minimum units of 8 (1 Byte) You can store 8, 16, 24, 32 or more bits per sample right the way up to 2040 bits (255 bytes)! Aural illusion v1.1 however does not support anything over 16 bit as yet.

You can have up to 255 channels worth of data, these are not interleaved like the other formats but are stored one after the other, the reason for this is speed, it is much faster to load in the sample data in single lumps without having to de-interleave them, it also means you don't have to write a different stereo loading routine - you just use the mono one twice.

The data can be interleaved in blocks of 25000 samples if the Block Interleave Byte is set to 3, in this case 25000 samples of channel 1 are saved followed by 25000 samples of channel 2 then back to channel 1 etc... This will allow data to be played from a hard disc with multiple tracks all in one file but without having to search or de-interleave data. Aural illusion v1.1 ignores the Block interleave byte at the moment.

SAFF also allows you to store your sample data in Least Significant Bit first format, this means sample data could be transferred to and from the PC without difficulty, by default data is stored in Most Significant Bit

first order. (Amiga / Mac / ST standard).

The header as well as being small is also easy to decode. All Longs are on Long offsets and the same applies for the sample Words/Bytes.

---- The SAFF Header ----

The SAFF header is at the beginning of a SAFF file and is 32 Bytes long.

Offset Bytes Description Type

0 4 Ascii code - "SAFF" Ascii SAFF

4 1 Channels Byte 1 or more

5 1 Bytes Byte 1 or more

6 1 Loop Byte 1=on 0=off

7 1 Byte Order Byte 0=MSB 1=LSB

8 4 Sample Length ULONG Samples per channel

12 4 Playback Rate ULONG Samples per second

16 4 Loop Start ULONG Offset from 0

20 4 Loop End ULONG Offset from 0

25 1 MIDI Byte 0=off 1=on

26 1 MIDI NOTE Byte Note number (0-127)

27 1 Compression type Byte 0=none See below

28 1 Block interleave Byte 0=none See below

29 4 Set to zero ULONG 0=sample data next

ULONG = 32 bit unsigned

Byte = 8 bit signed

The sample length is in samples NOT words, to calculate the size of the sample data multiply the number of samples by the number of bytes.

For the entire samples multiply the above resulting number by the number of channels.

There are no standard compression types as yet but I may add some in the future, so the compression byte should be set to zero for now.

Block interleave sizes (in samples - NOT Bytes or Words)

0=none

1=10000

2=15000

3=25000

4=50000

5=75000

6=100000

---- The sample Data ----

Data is saved as signed words, be they 8, 16, 24, 32 or however many bits, if your data is held as unsigned words convert them into signed words as

you save them. If you are storing 12 or 14 bit data store it as 2 bytes.

A 24 bit sample will be stored as 3 bytes.

Data is stored in sequential blocks starting with the first channel.

The sample data is not interleaved so loading and saving can be very fast achieved with memory dumps.

Block interleave allows the interleaving of channels in blocks, 25000 samples at a time, this will allow you to read or write more than one channel at a time direct from Hard disc, The blocks are stored :

Channel 1,2,3,4,1,2,3,4 etc...

If you were using a block size of 25000 (Block interleave=3) and 2 tracks of 16 bit data you would read 50000 bytes of channel 1 then 50000 bytes of channel 2 then 50000 bytes of channel 1 and so on...

By reading 100000 bytes at a time and setting the hardware audio pointers to the beginning of each block then using double buffering you could easily play multitrack audio from a hard disc without using almost any CPU power.

It is unlikely that the sample you are saving will have a size that is an exact multiple of the Block interleave size, so the final blocks should be shortened to allow for this, padding out with zeros would be a waste of space so this method is not used, file readers should take account of this.

If you are using a PC and wish to store data in SAFF format set the LSB byte to 1 so other platforms know the data bytes will need to be reversed.

If when reading the file check the first 4 bytes ASCII code, they should read as "SAFF" but if they read ASFF the entire file will need to be byte swapped.

---- Copyright & future changes ----

Write zero into the final four bytes as these are reserved for future use, perhaps to indicate another header before the sample data. If more is added they will be in the form of more 32 byte headers, a zero value in the last ULONG will indicate sample data is next.

I will retain the copyright of the SAFF format but only to stop it getting messed up and confused. Anyone can use the SAFF format if they wish, just don't change anything. It's designed as a simple format for storing and transferring sample data, if you wish to add to the specification contact me FIRST and I can update the format specification. There are no charges for using SAFF - it's making transfer of samples easier, not profit.

Some public domain programs will be release at a later date to deal with SAFFs, these will probably be Amiga only but I may do other platforms as well.

SAFF is © Blachford Technology 6th March 1994

1.15 The Edit Window

The Edit window controls:-

Cut - Removes selected section of a sample.

Copy over - First select the range in one sample then a point in the other, the range will be copied over the area from the point.

Copy into - Similar to copy over but this time the rest of the sample is moved along, not destroyed.

Copy all - Copies the entire selected sample into the other window.

Play view - Plays the current view - mono only.

Play range - Plays the current range - mono or stereo. If stereo is selected the area equivalent to the selected range will be played from the other sample.

Play all - plays all of the sample - mono or stereo.

Play Loop - Plays the current Loop, in mono or stereo.

See [Sound Playback](#)

Stopping play

To stop a sample playing press one of the play buttons, the sound will then cut off after a short delay - a long delay if the sample rate is low.

Stereo playback will stutter if you are using a 68000 CPU, On other CPUs stuttering will occur if any other program is using a lot of CPU time,

Stuttering may also occur with mono and stereo sample looping.

Show Rng - Shows the current selected range.

Show all - Shows the entire sample.

Zoom in - Shows less of the sample.

Zoom out - Shows more of the sample.

Range all - selects entire sample as a range.

Zero - Clears selected range.

Scale % - Increases / Decreases selected area by amount set in slider immediately below.

1 - Loads Mono IFF sample into buffer 1.

2 - loads Mono IFF sample into buffer 2.

Sample rate - Controls playback speed of sample, defaults to 31677Hz.

Mono/Stereo - Selects Loading/Saving and Playback mode .

Effects - Opens effects selection window.

Manips - Opens the Manipulations window.

Start - Sets 1st Loop point.

End - Sets 2nd Loop point.

- If Start and End are pressed on the same spot looping is switched off.

Displays

Message window - indicates current operation, When clear this means it has finished operation and is waiting for your next command.

Cusr 1: - Displays position of cursor in sample 1.

Cusr 2: - Displays position of cursor in sample 2.

Range: - Displays size of last range selected.

1.16 Sound playback

Any rates above 28185Hz will be of lesser quality since the Amiga's hardware is not designed to play any faster and are included just to give an indication of how the result will sound if played on 8/16 bit audio hardware at that rate. Any play rate below 44190 will be the same as OctaMED, This is a very popular Amiga Sequencer/tracker and the same rates are used to prevent samples going out of tune when played using it. The Amigas audio hardware is only 8 bit so sound is only played at that resolution, if you wish to hear your samples in 16 bit quality you will need additional hardware.

Playback Rates Compatibility Notes

4181 - 28185Hz OctaMED etc..

29829 - 31677Hz OctaMED etc.. Will crackle slightly

44190Hz CD/DAT Muffled

48046Hz DAT Muffled

50062Hz Atari Falcon Muffled

57734Hz Other Amiga products Muffled

1.17 Pull down menus

Project

Zero all - Clears all samples

Open - Loads selected file type into the selected range/s.

If nothing is loaded you may be trying the wrong file type, try one of the RAW formats if unsure.

Save - Saves the current selected range/s as the selected file type.

Almost everything supports Iff samples but these are only 8 bit,

Most Amiga 16 bit packages support AIFF.

If you wish to store data not interchange it SAFF is recommended
as it is the fastest.

File Formats in Ai v1.1

About - Displays info about Ai

Quit Aural Illusion - Quits program - CAREFUL! It doesn't ask if you're sure!

Edit

Cut - Same as on window

Copy over - Same as on window

Copy into - Same as on window

Copy all - Same as on window

Rescale - This will rescale the selected range to it's maximum possible
volume without distortion.

Swap Byte Order - Switches the order of the bytes around, Will produce white
noise if tried on a normal sample.

Toggle Filter - Switches filter on or off, if your power light is dim the
filter is off and the sound will be brighter.

1.18 Sample looping

Sample looping allows you to make a short sample last forever.

You set a Start and End point and press Play Loop, what is in the loop is
then repeated.

The loops can be Mono or Stereo and can last as long as the buffer, It is
possible to have a loop as drums or a backing rhythm.

To set the loop points you use the Start and End buttons on the Edit window,
you click the left mouse button wherever you wish the point to be and then
press Start or End, you then set the other point by choosing where it is to
be and pressing the other button. Two red lines should appear, these are the
loop points.

To hear the loop press Play Loop, the loop will then play until you press
one of the play buttons to stop play. If no loop is set nothing will play.

To remove the loop press both Start and end immediately after on another,
if the Start and End point are the same there is no loop to play so looping
is switched off, If you attempt to set the Start after the End or End before
the Start looping will then also be switched off.

Where to set loop points

Knowing where to set loop points can take a lot of practice as it can be difficult getting the loop to be clean, you want the sound to be heard but not the loop points. Try to set the loop points on a zero point (the red line in the sample displays) also try to make sure the points are at the same point in their waves, if the Start point is on a downward slope the End point should also be on a downward slope. Try and make sure the waveforms are as similar as possible otherwise the loop point will stand out a mile. Listen to the loop also, if the beginning has a different frequency mix to the end this will show up badly, so try and make the beginning sound similar. One method of making a loop not "visible" is to make it long, remember the points above but a long loop won't sound like its being repeated.

1.19 The Sample Effects

When you press Effects on the Edit Window the Effects selection window will open, This window has 22 buttons, 21 of these select the effects and the Cancel button closes the window. Pressing one of the effects buttons will open a different window which allows you to control and execute the Effect, when the Effect window opens the Effects selection window will close. Executing an Effect

All Effects windows are laid out in a similar manner, Two gadgets are always present:

GO - This executes the effect, The message box on the Edit window will indicate what effect is being executed, you should not press anything while the effect is calculating. When the effect is complete the message box will clear.

Cancel - This button will close the effect window if you had set any levels these will be lost.

Add/Subtract - Most effects also have an Add/Subtract selector, Subtract does not attempt to remove the effect but changes the frequency mix present in the effect.

Effects only work on the selected range, it is usually wise to select a range bigger than the sample if you use an Echo or Reverberation type effect. Effects also only work in Mono, even if Stereo is selected, if you wish to effect both sides select the other side as a range and redo the effect once the first side is complete, by changing the settings slightly a better stereo separation will be heard.

---- The Effects ----

REVERB - Creates a reverberation effect with indistinct echos.

Delay - Sets reverb delay time.

Decay - Sets initial level, best kept low.

Passes - Number of times the reverb is repeated, a small number can sound unnatural.

Reverb/Subverb - Reverb is added, Subverb is subtracted, There is a subtle difference between the two.

ECHO - Repeats the original sound a number of times.

Delay - Sets delay time.

Decay - Sets initial level, best kept low.

Passes - Number of echoes.

Add/Subtract - selects between added or subtracted echoes.

DELAY - Delays the sound and adds it to original.

Depth - Level of sound.

Add/Subtract - selects between added or subtracted delay.

BOUNCE - Gives very mechanical sounding short echo.

Delay - sets delay time.

Decay - sets initial level.

Add/subtract - selects between added and subtracted bounce.

BACKVERB - Creates a backwards reverberation effect.

Delay - Sets reverb delay time.

Decay - Sets initial level, best kept low.

Passes - Number of times the reverb is repeated, a small number can sound unnatural.

Add/subtract - There is a subtle difference between the two.

PHASE SHIFT - Has a filtering sort of effect if kept low.

Shift - Sets phase shift length.

Passes - Number of phase shifts.

Add/subtract - Changes frequency response of result, very pronounced difference.

PHASER - Gives sweeping effect.

Value - sets speed of phasing, small numbers here will give a more pronounced effect ending like a delay.

Add/subtract - Selects added or subtracted effect, changes the frequency mix of the sweep.

FLANGE - This is the most powerful effect with a variety of uses, it has a very powerful sweeping and frequency changing effect.

Delay - Amount of shift.

Level - Number of shifts.

Depth - Number of repeat phases.

Phase - Speed of phases.

PhaseD - (Phase Delay) Stops the sharp beginning or end to the flange which sometimes occurs.

Add/Subtract - On more than one pass this can have quite a difference.

Straight/reverse- Straight is best for pure tones, whereas reverse best on recorded samples, if reverse is used set the Phase high.

CHORUS - Gives impression of many sounds.

Spread - Amount copies are spread out.

Delay - Phasing speed.

Passes - Number of times effect is done, at higher levels this can give a very unusual effect.

Add/subtract - Added or subtracted effect.

REPHARB - Creates a phased reverberation effect.

Delay - Sets reverb delay time.

Decay - Sets initial level, best kept low.

Passes - Number of times the reverb is repeated.

Phase - Phasing speed.

Add/subtract - Subtracted is usually louder.

DISTORTION - Distorts sound.

Level - Level of distortion shift.

Depth - Level of distortion added to original.

Upper - Upper half of sample effected only.

Lower - Lower half of sample effected only.

Both - Both halves of sample effected.

Add/subtract - Added or subtracted distortion.

CONTORT - Computerised distortion.

Level - level of Contortion / Distortion - best kept low.

F-Distort - (Fast distort) Faster version of normal distortion.

Contort - Replaces distorted values with other values which are related to the original sample gives a very crackly form of distortion, the suitability of this effect very much depends on the sample being effected.

RESONANT - Resonates sample around a specific frequency.

Length - Length of delay.

Passes - Number of times effect is repeated.

Add/subtract - Added or subtracted, gives different sound.

FEEDBACK - Similar to resonant but sharper and stronger.

Length - Length of delay.

Passes - Number of times effect is repeated.

Add/subtract - Added or subtracted, gives different sound.

FILTER - Removes high or low frequencies.

Passes - Number of time filter is applied.

Low/High - Low acts as a low pass filter removing most high and mid band frequencies, High does the opposite removing all but the high frequencies.

RESAMPLE - Shortens or lengthens sound.

Value - Size change

100 = double size when increasing.

100 = half size when decreasing.

Increase/decrease- Selects between increasing or decreasing size.

DETUNE - Shortens or lengthens sound and adds to original.

Value - Size change, values same as Resample

Depth - Amount added back.

Increase/decrease- selects between increasing or decreasing size.

Add/subtract - Added or subtracted.

MAKE STEREO - Makes a single sound seem like it's in stereo.

Depth - Amount effecting other side when Add is selected (otherwise ignored) If you use the overwrite option the operation is much faster.

Overwrite - Overwrite other sample.

Add - Add to other sample.

ENHANCE - Enhances treble response of sample.

Depth - Level of enhancement, best kept low.

Add/subtract - Subtract gives a more distorted effect, Add is cleaner.

UNQUALITY - Reduces the quality of the sample, from 16 bit at 31KHz down to 1 bit 0.7KHz with both sliders set at maximum

If you want your sample to sound like a 1980s home computer or a 70s video game then this is the effect.

Bits - Number of bits to be removed.

HFloss - Amount of bandwidth to be removed.

PRESETS - These are preset phase shift effects which have different filtering type actions.

Bass - Enhances low bass response.

Tin - Makes sound tinny.

Tube - Makes sample sound as if it were in a tube.

Add/subtract - Changing to subtract gives the opposite of the effect

- Bass and subtract will remove the bass.

For more info see [Hints & Tips](#).

1.20 Manipulations

Manipulations

The Manipulations window contains 30 different manipulations which can be applied to one or more samples. Some of these require two inputs.

The way 2 input Manipulations work is the opposite way to copying.

If you wish to Add a drum with a string sound with the result overwriting the string heres what to do:-

- 1) First click the pointer at the beginning of the drum sound.
- 2) Select the String sound as a range - this range must be in the other sample window.
- 3) Click Add.
- 4) The result will be placed in the current range, you can hear this by pressing Play Rng.

Manipulations operate in the same way as effects, the result of the operation is placed in the selected range.

Most Manipulations are better suited to pure waveforms rather than recorded samples, However it very much depends on what you are manipulating as to the quality of the results, some manipulations change the waveforms shape which can sound from strange to horrible on some samples but pleasant or interesting on others.

---- The Manipulations ----

** indicates 2 inputs required

INVERT - Turns sample inside-out.

Gives very distorted sound on recorded samples.

HALF WAVE - Zeros all data below the center line.

FULL WAVE - Flips all data below the center line upward.

gives strange results on samples.

FLIP - Flips sample upside down.

Virtually no Audible effect on anything.

CENTRE - Centres sample around zero line.

Can remove or introduce some loud clicks.

REVERSE - Reverses sample.

Result is backwards.

ADD - Adds two samples.

** Useful for mixing samples.

SUBTRACT - Subtracts one sample from the other so $\text{Sam1} = \text{Sam1} - \text{Sam2}$.

** Has very similar result to add.

BSUBTRACT - Subtracts in reverse order so $\text{Sam1} = \text{Sam2} - \text{Sam1}$.

** Similar result in different place.

MULTIPLY - Multiplies two samples.

** Best on pure tones.

DIVIDE - Divides one sample by the other so $\text{Sam1} = \text{Sam1} / \text{Sam2}$.

** Also best on pure tones.

BDIVIDE - Divides in reverse order so $\text{Sam1} = \text{Sam2} / \text{Sam1}$.

** Yet again best on pure tones.

AND - ANDs two samples.

** Useful for producing a noisy radio sound.

OR - ORs two samples.

** Another Noisy radio.

XOR - Exclusively ORs two samples.

** Yet another noisy radio.

GREAT - Finds Greatest of two samples.

** Strange mixing effect.

LESS - Finds least of two samples.

** Another strange mixing effect.

OVER - Deletes all sample values below 8000.

Removes quiet noise but adds some as well.

THICKEN - Makes sound thicker.

Useful on anything but use sparingly unless everything
will sound the same.

THICKEN + - Makes sound much thicker.

Can give a pronounced phasing effect on some samples.

STUTTER - This uses one sample to modulate another sample in one of 2
ways,

** on or off, where the modulating sample is positive no change
occurs but when it is negative the result is zero, if you stutter
a sample with a low frequency wave the sample will switch on and
off sounding like it's stuttering.

MODULATE - In this case the modulating sample determines the level or

** amplitude of the other sample using a low frequency wave will
cause the sample to go up and down in volume.

OCTAVE + - This will double the frequency of any sample.

OCTAVE - - This will half the frequency of any sample, the length of the
sample will also double so ensure the range selected is big
enough.

CLEAN - This removes the highest frequencies in a sample, it was added
to remove noise generated when Octave - was used.

UNCLEAN - This is the opposite of Clean, the noise being removed is kept while the rest is removed, the result is a very high pass filter.

M CLEAN - Another filter but this removes both the highest and lowest frequencies leaving the high mid band.

M FILTER - This is also a mid band filter but leaves the lower mid band.

+ DISTORT - Version 1.0 of Ai had a different Enhance algorithm which produced a distinctive distortion of the sound when turned up to maximum, this is that same effect.

- DISTORT - The subtract option of the above produces another distinctive form of distortion.

For more info see [Hints & Tips](#).

1.21 The Synthesizer

The Synthesizer

Aural illusion Includes a Synthesizer with which you can generate sounds.

The synth can produce an almost infinite variety of waveforms and if thats not enough you can grab waveforms in from samples. The synth is designed to act as a starting platform from which to create your sounds, You generate a sample with the synth which will sound pretty basic and then using the Manipulations and effects shape the sound into whatever you choose.

Basic operation of the synth

When you press the Synth button a large window will open which is covered in controls most of which are sliders.

The controls are in four columns. At the top are four blank windows which will display the waveforms when they're generated, below the windows are eight wave sliders which are used to make the waveforms you will then use to make your sound. In a normal synth you chose a wave from what is available and use it, Ai's synth can make waves using the sliders and then "morph" the wave into the other waves over time so you can make sounds which change through time.

On the left hand side you will see an Add/Sub selector, this determines the way in which the waves are combined to make the wave used for synthesis, When Add is selected the waves are simply added in the amounts you have chosen, when Sub is selected the waves are subtracted in turn from the total.

Making and Grabbing waves

Below the wave sliders are two buttons labelled Make A and Grab, pressing

Make A will construct and display the wave from the settings in the wave sliders in column A, You don't need to press this if your generating or morphing a sample as this is done automatically, the Make buttons however allow you to see exactly what wave you're making.

The Grab buttons allow you to take and use a section of a sample and use it as a wave, All samples when grabbed are squeezed into a buffer 1000 samples long so it is not possible to take large sections of samples and use these as waves, you can grab large samples but they will be squeezed to fit into the buffer. Once a wave has been grabbed the wave sliders in that column will become useless - you can't mix a grabbed wave to make another wave, A grabbed wave will remain in its buffer until either you quit the program or you press Make, this will clear the buffer and over write it with whatever you've set in the wave sliders.

Envelope controls

Below the Grab and Make buttons are the envelope shaping sliders. These are used both in the synth and outside it, The envelope you've set is automatically applied to the sound you are creating but can also be applied to normal samples, this allows things like creating a drum sound from white noise like an analogue synth. The envelope shaper is a simple 4 stage ending at zero. The length settings are set with the length sliders, The lengths you set are ratio based since the length of the sample is determined by the selected range, so you can set four values of 1 and the stages in the envelope will be of equal length. The length settings also determine the length between which waves morph.

The Volume sliders are used to set the levels for the envelope, using these together with the length sliders you can generate any envelope shape you wish.

Generating a sound

To make a sample you first set a frequency you want it to be by moving the Frequency slider, if you change the sample play rate the frequency sliders will be wrong, ALWAYS select the frequency after setting the play rate.

Frequency values for 5 Octaves

Note Freq | Freq | Freq | Freq | Freq |

A 55 | 110 | 220 | 440 | 880 |

A# 58 | 117 | 233 | 466 | 932 |

B 62 | 123 | 247 | 494 | 987 |

C 65 | 131 | 262 | 523 | 1047 |
C# 69 | 139 | 277 | 554 | 1109 |
D 73 | 147 | 294 | 587 | 1175 |
D# 78 | 156 | 311 | 622 | 1245 |
E 82 | 165 | 329 | 659 | 1319 |
F 87 | 175 | 349 | 699 | 1397 |
F# 92 | 185 | 370 | 740 | 1480 |
G 98 | 196 | 392 | 784 | 1568 |
G# 104 | 207 | 415 | 830 | 1661 |

if these values can't be generated get the closest fit.

To actually make a sample you will first set the wave sliders and the envelope settings, you also need to select a range which should be over 10,000 samples long, if it is not the synth will not operate.

To create a sample using a single wave press Generate, this will make the sample using only the wave created or grabbed in column 1, no other columns are used for waves but the full envelope is still applied.

To create a sound using the four waves press Morph, this will create a sample and will use all 4 waves morphing between them in turn.

The synth will only ever generate a single wave although its volume and shape will change, it is like a single oscillator synthesizer, if you wish to add more oscillators you can use the manipulations Thicken or Thicken +, these have been added for this purpose, Using low values of the Effect Detune will have similar results but using Thicken all the time will make all your sounds very similar.

Applying an Envelope to a sample

To do this all you need to do is to select the range you wish to be effected and after setting the Envelope sliders press Apply Env, the envelope settings will then be applied to the range.

After you've made your sound

Once you've made a sound it may not exactly jump out at you, this is where the Manipulations and Effects come in, by using these you can change your sound into something completely different or simply shape it into a better form.

To now make a sound see [Hints and Tips](#).

1.22 Hints and Tips

This section contains some information on how to get more out of Aural illusion, It is possible to do things with Ai which you may not think possible - There are no proper mixing facilities but you can mix sounds and filter sweeps are possible even though the filters can't be swept.

Synth sounds

Filter sweep

Mixing samples

Phase shifting

Resonant Filter

1.23 Synth hints

OK, You may have read the notes but do you know how to work the synth?

Here I will take you through the steps of how to make a couple of sounds.

Number 1:

- Select a range of 24,000 samples, (doesn't have to be exact).
- Press Synth.
- Press Sub.
- Set Ramp in all columns to 256.
- Set the D-Spike values in the columns :
A=256 B=209 C=85 D=3.
- Set the Length values all to 1.
- Set the Volume values all to 256.
- Set the frequency to 327 (E).

Press Morph

The Synth window will now disappear behind the other windows and the word Morphing will appear in the message box.

When the message box clears the sample will have been generated and you can hear it by pressing Play Rng.

- When this is done press Cancel on the Synth window to close it.

Now to try a few effects:

- Copy the range to the beginning of the other sample (click at the beginning and press Copy over).
 - Click after the end of the copied sample and press copy over again.
 - Do this a few times until you have 5 copies of the sample.
 - Now press copy all, the 5 copies will be copied over the first sample.
-

- Select the first copy as a range and press Manips for the manipulations window (make sure the Synth window is closed first).

- Press Thicken.

When this is done press Play Rng to hear the result of the effect.

It may sound slightly crackly.

- Select the next copy and press Invert.

- When this is done change the play rate to 28185 and then press Play Rng.

Every rate above 28185 will distort the sound slightly so 28185 should be used for most accuracy, playing this sound at 31677 will show this difference in quality.

There will be a click at the beginning and end of the sample, this is due to what Invert does.

- Select the third copy

- Press Full wave

- Press play Rng when complete for the result.

- Select the fourth copy

- Press M Filter

- The sound will be much thinner as the filter has removed the low and highest frequencies.

- Select the fifth copy

- Press + Distort

- When complete play the sound and it will also sound as if it had been filtered in some way but it hasn't.

To make the above results even better try Thicken on them all, you will notice this does exactly what it says, thickens.

On the 5th copy try distort + again, this time the result will not sound like it's been simply filtered.

You will now have 5 different sounds which all started as the same one, and it is here where Ai's power is, By using a variety of effects on each sound you can produce a variety of different results.

- Press Copy all, this will copy your 5 sounds into the other sample buffer, Try other Manipulations or Effects on these, the Flange Effect can be a multitude of effects on its own.

1.24 Filter Sweep

Making a filter sweep

The filters in Ai use set values and can't be changed over time, this

however does not mean that you can't do filter sweeping effects.

This is how to do a simple sweep with a sound where the sample will begin sounding dull and end up sounding bright and will sound neither dull nor bright in the middle. It is best to use a fairly long sample for this as you can hear the effect better.

- First select your sample and make 3 copies of it, make sure you have a spare of the original in case you need it.

- Apply 3 filters to the sounds

Filter (Low) to the first.

M Filter to the second.

UnClean to the third.

That's the three sounds you're going to use, but how do you get them to change over time?

To get the sounds to change you are going to modulate them using a very low frequency sound, the sound will have a frequency of less than 1Hz and you are only going to use part of a single waveform.

You want the bass to be loud at the beginning and quiet at the end, for treble you want the reverse, quiet at the beginning and loud at the end.

For the midband you want it loud in the middle and quiet at both ends.

To make the sounds you are going to use the synth, It cannot produce the low notes we want so you will use the Manipulation Octave -.

- Open the Synth and set up the envelope as lengths=1, and volumes=256.

The wave sliders should be set to Sine=256. Select the lowest frequency and select a range no smaller than 10,000 samples long.

- Press Generate to make the sample.

- Select a bigger range which should be close to the size of your sample, Make sure the sample just generated by the synth is at the beginning of the range.

- Press Octave - until the individual waveforms are visible in the sample.

You will only need the top half of the waveform so you should press Octave - until this is all that remains.

- The note will now be rather noisy as the bandwidth has been reduced so much, to clear this open the Effects window and bring up the Filter, Use a setting of Low and 10 then press go, this should remove all the noise, the sample may now sound hissy but this is the 8 bit sound hardware making this noise so don't worry about it.

- Do the same process as above but this time with a triangle wave, you also only want the top half of it, this time make the final waveform twice as long as the sine wave.

- Split the Triangle into two by selecting a section from the other buffer and using Copy into.

You will now have three waveforms and 3 filtered sounds, the waveforms will be used to modulate the sounds.

Modulate uses one wave to determine the volume of another, thus one is used to modulate the other.

- Make sure the sine and triangle parts are in a different buffer from the filtered sounds.

- Click at the beginning of the raising triangle half (the part which starts low and ends high). Then select the sound which used UnClean (the treble) and press the Manipulation Modulate.

The result should be a sound which starts quiet and ends loud.

- Now click the beginning of the sine part and the midband filtered sound which used M Filter and press Modulate.

The result will start and end quiet getting louder in the middle.

- Click at the part of the triangle which starts high and ends low, the point you click should be in the shape not before it.

Select the Bass sound which used Filter and press Modulate.

The result should start loud and end quiet.

You should have three filtered sounds which get loud or quiet at different points, all you now have to do now is to mix them.

Aural illusion has no proper mixing facilities but can add sounds.

If you are adding more than two sounds you will need to remember that each time a sound is added its volume halves. If you are not careful you can add sounds which will all end up at different volumes.

- Copy the bass sound to the other buffer, click the beginning of the treble sound and then select the bass sound as a range, then press the Manipulation Add.

The first two sounds will have been mixed OK but this is where the problems start, If the third sound was added now the bass and treble sounds volumes will be quarter what they were while the midband would be half. To correct this you half the midband sound by Scaling it to 50%.

- Select the Midband sound as a range and set the slider below Scale on the Edit window to 50, Press Scale and the ranges volume will be cut to 50%, select the beginning of the sample and then select the sample in the other buffer as a range, press Add.

You will now have mixed the three samples correctly and the result is a sound which starts with bass then changes to midband then to treble.

1.25 Mixing Samples

As you may have seen in the filter sweep section you can mix samples but problems can occur.

This is how to mix sounds correctly.

All sounds should be mixed in pairs :

4 sounds correct way:

$1 + 2 = \text{Result 1}$

$3 + 4 = \text{Result 2}$

$\text{Result 1} + \text{Result 2} = \text{Final mix.}$

The above result will contain equal amounts of all four sounds.

4 sounds wrong way:

$1 + 2 + 3 + 4 = \text{Wrong result.}$

The above method will result in sounds all at progressively higher volumes where 4 is the loudest and 1 the quietest.

3 sounds

$1 + 2 = \text{Result 1}$

3 Scaled to 50% = result 2

$\text{Result 1} + \text{Result 2} = \text{Correct volume for each sound.}$

Scaling sound 3 to 50% has the same effect of adding it to another sound and corrects the level.

1.26 Resonant Filter

How to simulate an Analogue Resonant Filter

Analogue sounds are currently in vogue and a particular part of an analogue synth is the filter, these will go into a resonant state if you turn them up and this can be used to produce some interesting sounds which have a "squelchy" sound.

Reproducing this on a digital system is difficult since software does not go into resonance and Aural illusion has not been programmed to simulate the effect - this however does not mean it can't be done.

The below method works on some samples I have tried but you may need to mess around with the settings to get it to work properly, it is not an exact method by any means.

First get the sample you wish to try it on, copy it to the second buffer and use Filter (High) to get the highest frequencies, Then use Over to remove the quieter parts of the sample.

There should now be very little left of your sample - if not you may be using the wrong sample, this method is a bit fussy.

Select Reverb with delay set to 8 and 2 passes, once this is done try 1 pass at a delay of 16. Once done use Enhance set to 18 to increase the treble energy, then Add the processed sample back to the original, You should now have a resonant filter sound which may or may not be any good. This method works best on filter sweep type sounds and sounds which last a very short time, if your sample is very bassy you may need to Enhance it before trying the method.

1.27 Phase Shifting

Phase shifting

Phase shift can be one of the most powerful effects, Uses include:

De-hancing Modulation

Enhancing Modulation

Adding unpredictable Modulation

De-hancing Modulation

Sometimes you may have a sample which has modulation of any sort but whose level is too high, Phase shift can be used to reduce it.

First you will need to know the wave length of the note your sample consists of. If you look closely at a sample it consists of repeating waves, the wave length is the length of one of these waves.

Open the Phase shift window and set the Length to equal that of your sample, make sure you have Add selected and at least 1 pass, Select Go to execute the effect. The result will be quite subtle with one pass so you may need to try more, or repeat the effect, you should notice the modulation reduced as the sample "moves" closer to the pure note.

This technique works so well in fact that you can tune noises which have no note as such - try it on white noise.

Enhancing Modulation

This is the opposite of the above, but make sure you select Subtract, the sound will then "move" away from the pure note if the correct wave length is set.

If you keep repeating the effect the modulation will take over completely almost completely removing the pure note.

Adding Unpredictable modulation

Above you have to select the correct wave length to work properly, if however you use the wrong wave length the sound will move to or from a different note altogether, the result of this is an "interference" of notes which at a low number of passes is rather unpredictable.

Subtracting the wrong wave length also has unpredictable effects.

Phase shift if kept low can have filter type effects with some frequencies made louder and some quieter.

1.28 Miscellaneous

Aural illusion v1.17 was written by Nicholas Blachford using :

Amiga A1200

SAS/C v6.2

GUI code generated by GadToolsBox V2.0

Many thanks to :

SAS Institute for SAS/C.

Jan van den Baard for GadToolsBox and the help with it.

Sharon McGuffie etc.. at Commodore UK developer support.

C.A.T.S.

Prince's Youth Business Trust for the money to get SAS/C.

John Kennedy for letting me "borrow" his A4000 and various suggestions.

And special thanks to :

Jimmy Hill for lending me his Computer, Printer, Monitor, just about every CD he possesses and for figuring out Amiga Guide documents.

Grant Mearns for reporting numerous bugs (which have been fixed), loads of Suggestions and lots more besides, - and of course for being my first customer.

AmigaGuide, AmigaGuide.info, amigaguide.library,

(c) Copyright 1991-93 Commodore-Amiga, Inc. All Rights Reserved.

Reproduced and distributed under license from Commodore.

Aural illusion, Aural illusion.info, Ai.Guide

and Aural illusion startup guide

(c) Copyright Blachford Technology 1994-95, All Rights Reserved.

AURAL ILLUSION AND AMIGAGUIDE SOFTWARE IS PROVIDED

"AS-IS" AND SUBJECT TO CHANGE;

NO WARRANTIES ARE MADE. ALL USE IS AT YOUR OWN RISK. NO LIABILITY OR RESPONSIBILITY IS ASSUMED.

1.29 Problems

Q & A

Q1 - When I press Rng all the Range: indicator says I don't have many samples even though the memory in my Amiga is quite big and I can get much bigger samples in other programs?

A - Ai allocates almost all the memory it can find in 3 equal size blocks, you can see two of these blocks in the sample windows but the third is hidden, Samples are stored in 16 bit format which take up two bytes each so you have half the number of samples. To find the amount of memory Ai has allocated press Rng all, then multiply the number displayed in Range: (number at the bottom on the right) by 6, this will give you the amount of memory taken by Ai.

Q2 - Why can't I have one big sample in memory instead of two?

A - Some of the manipulations require two inputs and thus two samples to apply these.

Q3 - Why does sound playback stutter when it plays in stereo or when playing a loop?

A - Playing back 16 bit sound with 8 bit hardware requires it to be converted in real time, without using "clever" playback routines this becomes quite a chore for the CPU and anyone with a 68000 may have stutter in stereo mode, A1200 users with the more powerful 68020 processor will not have stutter - providing they don't run anything in the background that uses a lot of CPU time.

Q4 - The program won't run, why?

A - If there is not enough memory or if the sound channels are in use the program will not run.

Q5 - I have found that when I do a certain operation the program or computer crashes, why?

A - I hope no one ever asks this but I can't test this program on all Amigas and bugs do tend to crop up when you're least expecting them - if this happens to you tell me about it as soon as possible.

Other Problems

If you find any problems with this program please let me know of them so I can get them sorted, if you do find a bug please tell me what Amiga / configuration you have and if possible how the bug occurred.

1.30 Aural illusion Version v2

Aural illusion v2.02

A much improved and updated version of Aural illusion with:

30 Effects including Time Stretch, Graphic Equalizer, Compress, Mixer, Noise Gate and Filter Sweep.

55 Manipulations including 8 Tunings, 8 Filters, 6 Modulations and various miscellaneous functions.

An improved easier to use Synthesizer with Phase controls for better sounds.

New file formats - WAV, VOC and RAW 8 bit PC

New miscellaneous functions including Loudness, Unloudness and Get Frequency.

Requires Workbench 2 or greater and 2 Meg RAM.

Aural illusion v2.02 costs £20 + P&P

Aural Synthetica v1.1

An all new sample generation program which uses Modular synthesis to provide the most versatile sound generation program on any platform.

This program does away with the limitations of normal synthesizers allowing the user to create their own synthesizer setups creating unheard-of sounds. It will create anything from simple blips to monster Analogue style lead sounds and everything in-between as well as sounds you can't even imagine.

A comprehensive waveform editor creates waveforms for the 18 powerful sound generating Oscillators which feed the processing modules. 29 Preset synth setups with 136 variations give the user scope to generate a huge number of sounds without any patch programming. There is also a Patch programmer for the brave with an almost infinite number of routing and modulation possibilities, these can be loaded, saved and edited at will.

Samples are created in CD resolution for the best possible quality and the program saves samples in IFF, AIFF, SAFF and WAV.

Aural Synthetica will require Workbench 2.0 or higher and 2 Meg RAM.

An accelerator and additional RAM are highly recommended.

Aural Synthetica v1.1 £21 inc P&P (UK pounds)

Aural illusion v2.0 £21 inc P&P (UK pounds)

Both packages can be bought from :

Nicholas Blachford

Glendale House

77 Southwell Road

Bangor

Co.Down

Northern Ireland

BT20 3AE

EMAIL minator@hearit.demon.co.uk

or

Seasoft Computing

Unit 3

Minster Court

Courtwick Lane

LittleHampton

West Sussex

BN17 7RN

Phone UK: 01903 850378
