Glowing Coast Technology
Audio Suite v2.03

Audio Suite is a fast, tightly integrated group of three Audio waveform processing programs which will turn your PC in to a powerful Digital Audio Workstation. Consisting of a Waveform Editor, a Digital Mixing Desk and a Format Converter.

Features include

Audio Wave Editor

Editing: Undo, Cut, Copy, Insert, Mix, Trim, Delete, Replace, Append.
Tools: Equalisation, Invert, Mute, Normalise, Gain Adjust, Remove DC offset, Pop Removal, Noise Gate, Compressor, Temporal Averaging, Filtering, Graphic Equaliser, Fade In, Fade Out, Ramp, DTMF generation.
Effects: File Multiplier, AM, FM, Enhance, Reverse, Time Compress\Expand, Delay, Echo, Flange, Phaser, Reverb, Vibro.
Special Tools : FFT, Spectral Averaging, Power Spectrum Analysis, Noise Reduction, DTMF Identification.
plus: Experimental Beat and Note Extraction routines.
Additionally for stereo tracks only: Cross fade, Bouncing cross fade, channel splitting.

Digital Mixing Desk (Real Time)

Works like a standard analogue mixer.
Mix up to 30 channels, (15 channels if all Mono, 30 channels if all Stereo)
Independent Left\Right volume control.
Delay the tracks to align correctly or create ambience effects.
Loop individual tracks, continuously or fixed number of cycles.
Save both the resultant mix and the mixer settings for a later date.
Apply Harmonic Enhancement and Equalisation to each channel.
32 bit processing throughout.
Re-Dithering available for increased Dynamic Range.

Audio File Converter

Convert from one set of frequency, resolution and channels, to another.
Anti-aliasing and noise reduction filtering available to minimise conversion artefacts.
Quick comparison with original and other formats allows fast determination of the required setting.
File formats supported - WAV, AU, AFF, AIF, IFF, VOC

Index

Audio Wave Editor

The Audio Wave Editor is a simple waveform editing tool which lets you visually manipulate any uncompressed *.wav file you may have.

Over View
The Audio Wave Editor is a hard disk based waveform editing tool, this means you can edit any size file you want. It is also an MDI program, that is many files can be open at once.

The first thing to do is Open a *.wav file from menu File|Open, unregistered users are limited to a 10 second section.

The Top panel contains the media player controls, a file length indicator, and a free Disk space indicator (this indicates space available on the Disk where the TEMP directory is located).

A second panel should have appeared, containing the Start and End section indicators, a button to play the selected section, a button to continuously Loop the selected section, a position indicator, and a playing position indicator.

The graphical representation of the selected file should also have appeared in it's own window, with it's name, sampling rate, resolution and channels contained in the title bar.

Selecting a Section of the track.
With the cursor over the selected file, click the left mouse button and drag the mouse to select a section of the file. Notice how the Start and End section indicators track your position. The Start and End section indicators define the position for the selected section that the Editing, Tools, and Effects commands use.

An alternative is to select the Section by manually rolling the up/down buttons on the Start and End selection indicators. You can adjust the step size, by altering Options|Fine Time Increments.

To quickly Select All the track press Q, or from menu Edit|Select All.
To Select and view All the track press Ctrl+Q, or from menu View|Full View.

Playing a Section of the track.
Clicking the Section Play button, or hitting the Space Bar, will start playing the selected section of track. If the Start and End points are in the same place then the track will be played from this point to the end of the track. With this feature you can dynamically jump about the track by clicking the mouse (left button) at the position you want to move to, and hitting the space bar.
Depressing the Loop button will cause the selected section of track to loop.

To stop the track playing you must click the stop button of the media player in the top panel.

Zooming In and Out.
To Zoom In, select the section to examine then press Ctrl+Z, or from menu View|Zoom In.
To Zoom Out, press Ctrl+A, or from menu View|Zoom Out.

Notice how the, Zoom In (Ctrl+Z), Zoom Out (Ctrl+A), Select All (Q), View All (Ctrl+Q) are all grouped together for easy access.

Keyboard Shortcuts.
Extensive use of Shortcuts has been made for all the most commonly used functions.
Editing functions are typically require two fingers (Ctrl + ?), and use the Microsoft convention for cutting and pasting, the Tools and Effects typically require one finger.

Getting started

The first thing to do is Open a *.wav file.

Click the Section Play button, or hit the Space Bar, to start playing the Track.

Editing.
Editing functions use the Microsoft convention for cutting and pasting. Cut and Paste features working similarly to those of a word processor.

The Audio Wave Editor has a powerful function which will Trim the file to the selected length, remove the DC Offset then Normalise it, all in one go. The function can be accessed through Edit|Trim-Offset-Normalise, or Ctrl+N. Notice how Tools|Normalise uses (N), related functions on same key.

Multiple files can be open simultaneously.
Flip between the files with (Ctrl + Tab), this will ensure Start/Stop markers aren't moved, this is particularly useful when working with big files.

Using the Tools and Effects.
The Tools and Effects are only applied to the Selected section.

Try this:
Click Time Compress/Expand from the Effects menu (T), reduce the selected length to about 90%.
Click the section play button again, the pitch of the selected section of track will have increased.

To flip back to the previous version, select Edit|Undo (U). Clicking U again will Redo the effect, only one level of Undo is available.

With the original waveform, select Effects|Echo (E). A panel will be presented requesting Delay time and Fading parameters, for now use the default settings, click OK. An Echo will be quickly added to the selected section of track.
Play the track to hear the effect.

All the Tools and Effects have been made as intuitive as possible.

More details of each function are available elsewhere in Manual.

Watching a Playing file.
Select All of the Track (Q), depress the Loop button and press the space bar to start playing.
Selecting Special|Frequency Analysis (Ctrl+F). This will display the instantaneous FFT of the current position, but it will probably be changing too fast to make any sense out of. Click Stop and Close to return to the Waveform.

Try using Special|Spectral Averaging instead, select 8 spectrums. Instead of displaying the instantaneous waveform the FFT is averaged. Now you can watch the tonal variations in the track.

As with all tools like this the way to learn, is to play.

Working with Big files.
Since an undo facility is a must at least two copies of the file are required, this slows down normal working, and gets significantly worse with big files. The bottleneck is writing the file to the disk, even the smallest modifications require the whole file to be rewritten.

This problem can be reduced by use of the File|Save Section and Open... command.
Select the section to be modified, use the above command to Save the Section and Open it for Editing. Do the modifications on this smaller file. You can flip between the two files until the modification is as desired.

Provided you don't move the start/end markers in the original file the Edit|Replace command can be used to perfectly align the modified file.
Once you're happy with the modified file; Select All (Q), and Copy (Ctrl+C) it all.
Go back to the original and Replace (Ctrl+R) the section, this deletes the original section and inserts this new section.

File Menu Commands

Record.
Initialise a new file for Recording too, the recorder is activated and ready to start.
The sound source must be specified by the software supplied with your sound card, the sound can then be recorded to disk. The file will be recorded in the format as specified in the Options| Recording Setup menu, the default being 16 bit, mono at 22kHz.

Whilst recording **do not** do anything else with the computer.

Playback File.
Open a file to be played whilst recording (only available for Full Duplex Sound Cards).

Open.
Open a new file.

Save.
Replaces the existing file.

Save As...
Saves a new or existing file.

Revert.
Restore the original file.

Save Section As...
Saves the selected section to a new or existing file.

Save Section and Open...
Saves the selected section to a new or existing file and Opens it ready for Editing.

Close.
Closes the active window.

Exit.
Quit the Audio Wave Editor.

Edit Menu Commands

Undo.
Undo previous editing.

Cut.
Copy and Cut section defined by Start and End positions.

Copy.
Copy section defined by Start and End positions.

Insert.
Insert copied section at Start position.

Mix Paste.
Mix copied section with selected file at Start position.

Trim-Offset-Normalise.
Trim the file, Remove the DC Offset and Normalise the amplitude.
Apply to the section defined by Start and End positions.
Fast way of implementing the most common functions used after Recording a new file.

Trim.
Trim the file to the section defined by Start and End positions.

Delete.
Delete section defined by Start and End positions.

Replace.
Replace section in one file with the copied section; particularly useful with big files.

Append.
Append copied section to End of file.

Insert Silence.
Insert up to 10 seconds of silence.

Select All.
Reset Start and End positions.

View Menu Commands

Full View.
View the entire waveform.

Zoom In.
Zoom In to position defined by Start and End positions.

Zoom Out.
Zoom Out by 50%.

Tools Menu Commands

Wave Adder.
Add one of the selected waveforms to the selected section.
Select the frequency of the waveform, and the amount of the original section required.
If 0% of original selected then the new waveform will have amplitude 100%.

Equalisation.
Produces a more uniform amplitude for the selected section.

Invert.
Invert the section.

Mute.
Zero the section.

Normalise.
Amplify section to maximum level without clipping.

Gain Adjust.
Amplify section by a defined amount.
<100% Reduce amplitude.
>100% Increase amplitude, be careful of clipping.

Remove DC Offset.
Remove any DC Offset which may have been introduced.
The offset amount is calculated by the computer, and removed from the entire file.

Pop Removal.
Attempts to remove any short discontinuities within the selected section, by interpolating across the discontinuity. This can, on occasions go wrong, particularly if the signal is transient in nature, or sampled at a low frequency. Also, if the signal is clipped whilst echoing or reverberating, the popping produced at a reduced gain probably won't be detected.

Noise Gate.
Removes any signal below the specified threshold, once the Hold condition is satisfied.
Threshold Amplitude at which the Noise Gate is activated.
Hold length of time over which signal must remain below the Threshold.
Release length of time before Noise Gate is switched off.

Compressor.
Compress the Signal above the Threshold by the amount determined by Ratio, once the Attack condition is satisfied.
Threshold Amplitude at which Compression is activated.
Ratio Amount of compression to apply.
Knee: **Hard** Activate compressor at Threshold.
Soft Activate compressor 6dB below Threshold, but with less compression, to produce a softer knee.
Any Signal >Threshold will be compressed by the amount Ratio.
Attack length of time after hitting the Threshold before Compressor is activated.
Release length of time before Compressor is switched off.

Filters.
Low Pass, High Pass, Band Pass and Notch Filtering are available.
If passband ripple is 0dB then a Butterworth filter is used, otherwise a Chebyshev is implemented.
The filters are designed with a maximum of 12 poles.

Graphic Equaliser.
The 10 band Graphic Equaliser lets you modify the characteristics of the sound, it works the same as a standard audio equaliser. This tool is suited to creating smooth tonal variations. If you need sharper cut-off then use the **Filters** tool.

Temporal Averaging.
A fast waveform smoothing technique, the selected section being smoothed over the number of samples selected. Can be useful for smoothing out noise and glitches.

Fade In.
Fade In the section of Track.

Fade Out.
Fade Out the section of Track.

Ramp.
Fade the Track between the Left and Right amplitudes.
Align Left: Left = 100% Right: Input amplitude.
Align Right: Right = 100% Left: Input amplitude.

Split Channels.
Split the Left and Right channels of a stereo Track in to separate mono files. A single filename is required for the new files, _l and _r are appended to the respective filenames.

Effects Menu Commands

Multiplier.
Multiply the selected section by another file. The file used as the multiplier must be 8 bits, mono, and of the same sampling rate as the current waveform.

Amplitude Modulate.
Multiply the selected section by one of the waveforms selected.
Select the frequency of the waveform, and the amount of the original section required.
If 0% of original selected then the new multiplied waveform will have amplitude 100%, a selection of >0% but <100% will mix the original with the multiplied waveform.

Frequency Modulate.
Multiply the selected section by an FM signal.
Select the centre frequency(Fc) of the waveform, the amount of the original section required, and the percentage Deviation of the centre frequency.
If % Deviation set to 100%, the FM signal will swing from; 0.5Fc to 1.5Fc
or Fc(1 - Deviation/200) to Fc(1 + Deviation/200).
If 0% of the Original section is selected then the new modulated waveform will have an amplitude of 100%, a selection of >0% but <100% will mix the original with the modulated waveform.

Enhance.
Increase the harmonic content of the section.
This sonically tends to give most emphasis to the lower frequencies, where the additional harmonics occur in the 300-4000Hz region; ie. where our hearing is most sensitive.

Reverse.
Reverse the section.

Time Compress/Expand.
Compress or Expand the selected section.
Specify the desired length of the selected section, the section will be expanded or compressed to fit.

Delay.
Creates a Delay effect on the selected section.
Define the amount of delay required, and the Gain of the delayed section.
Similar to an echo, but without feedback.

Echo.
Creates a simple Echo effect on the selected section.
Define the amount of delay required, and how the Delayed section will Fade out.
Setting both Fade parameters the same will fade the Echo at a fixed rate.

Settings the Fade parameters differently allows the front and back of the delay to be Fading at different rates.

Flange.
Flanging is created by modulating the pitch of the selected section.
Select the offset delay, the amount of the original section included, the amount of flanging (% deviation), and the frequency of the flanging.
The % Deviation is the % of the Delay. If Delay = 100mS, % Deviation = 20%, then the Delay will be modulated +/-10%, at the specified frequency rate. The flanging frequency is typically low <10Hz.

If 0% of original selected then the flanged waveform will have amplitude 100%, a selection of >0% but <100% will mix the original with the flanged waveform.

Phaser.
Phasing is an effect created by shifting the phase of the signal then adding it back to the original, causing some frequencies to be attenuated and others enhanced.
Depth is the amount of the Phased signal to be added to the original.
Feedback is the amount of Feedback to apply.
Phase Shift is the amount to shift the signal by.

Reverb.
Reverberation is a complicated, computationally intensive means of adding ambience to the sound.
A later section is dedicated to this function.

Vibro.
Multiply the selected section by one of the waveforms selected.
Select the frequency of the waveform, and the amount of the original section required.

If 0% of original selected then the new multiplied waveform will have amplitude 100%, a selection of >0% but <100% will mix the original with the multiplied waveform.
The Vibro effect is created by Amplitude Modulating the original, at a low frequency <30Hz with a large amount of the original signal included >60%.

Cross Fade.
Cross Fade the Left/Right channels of a stereo Track. The length of the fade is determined by the Start and End positions.

Bouncing Cross Fade.
Bounce the Left/Right channels of a stereo Track. The length of the area is determined by the Start and End positions, with a selected number of Bounces performed within this area.

Special Menu Commands

Frequency Analysis.
Perform an FFT at the current position, the size of the transform being determined by the FFT Setup option. The windowing function applied is determined in the FFT window.

Spectral Averaging.
Average the selected number of FFT sequences.

Power Spectrum Analysis.
Compute the FFT and average the signal over the required bandwidth.

FFT Setup.
Select the size of the FFT to perform.

Noise Reduction.
Reduce the Noise within the selected section.
A later section is dedicated to this function.

DTMF Identification.
Try to identify any DTMF tones in the selected section.

Two experimental functions for the conversion of Beat and Note information in to a midi file.
Originally developed to detect drum beats (well, fingers on a table) for direct midi input. The routines work OK with simple, uncluttered sounds. You will need to clean the sound up as much as possible prior to using the routines.

Beat Extraction.
Try to identify the Beats within the selected section, the results are written to a Midi file.

Note Extraction.
Try to identify the Notes within the selected section, the results are written to a Midi file.

Options Menu Commands

Open CD Player.
Activates the CD player if one is available.

Recording Set-up.
Select the format of any recordings to be made.

Fine Time Increments.
Modify the Fine Increments setting of the Start and End Positions.

Zero Crossing Detection.
When selecting a section with the mouse, force the computer to ensure the section starts and ends at a zero crossing point if possible. If a section small enough that each sample is discernible is selected, then the ZCD is automatically disabled.

The Reverb Set-up Dialog.

Reverberation is a means of adding ambience to the sound.
This is created by adding multiple (up to 8) delayed copies of itself at varying amplitudes to create the Early Reflections, feedback and cycling of these Early Reflections is used to generate the Reverberation Tail. The resulting additive and subtractive nature of the delayed waveforms produces the added ambience.

This consists of a white Setup sheet, on which delay and gain **Taps** can be placed. The length of Setup sheet can be scaled from 320mS down to 10mS, with the arrow buttons. Above the red line indicates +ve gain, below -ve gain.
The panel below the Setup sheet contains two edit boxes for accurate Delay and Gain settings.
An additional two edit boxes Left and Right are to adjust the entire Gain of all non-delayed taps, so each channel can be adjusted individually; these Gains can only reduce or invert the gains of the original Taps.

A **Tap**, is delay and gain indicator, up to 9 taps can be on the sheet, 8 plus the original to which only a gain adjustment can be applied.

The Original Waveform is indicated by a Tap of zero delay (red), additional Taps are blue.

To add a Tap to the sheet, place the mouse close to the left of the sheet, left click the mouse and drag out a Tap to the desired location, the delay and gain being indicated in the edit boxes. Fine adjustments can be made of the currently selected Tap via the edit boxes.

Whenever you left click the nearest Tap will be picked up, if the nearest is at 0mS (which all cleared Taps are) then a new Tap is selected.

The act of adding a Tap, is placing the position of an Early Reflection.
The Tap with the greatest Delay controls the amount of feedback applied to generate the reverberation tail. Temporal Averaging (it's faster than filtering) of the Reverbed signal is performed, dependant on the position of the Reverb Tap.

Note: Sound travels at approximately 1 foot per millisecond.

The Buttons.

Open	the specified reverb settings.
Save	the current reverb settings.
Reset	the reverb settings.
OK	perform the reverberation as defined by the current settings.
Cancel	Exit and do nothing.

Reverberation Design

This is created by adding multiple (up to 8) delayed copies of itself at varying amplitudes to create the Early Reflections, feedback and cycling of these Early Reflections is used to generate the Reverberation Tail. The resulting additive and subtractive nature of the delayed waveforms produces the added ambience.

How Do I create Reverb effects ?

With the correct placement of Taps.
The act of adding a Tap, is effectively placing the position of an Early Reflection.
The Tap with the greatest Delay controls the amount of feedback applied to generate the reverberation tail.

The Reverberation Tail ?

This is where things can get difficult, the Early Reflections also control the creation of the Tail.

The Tap with the greatest Delay controls the amount of feedback applied, and the start position of the reverberation tail.
Placing only One Tap will create an effect similar to the simple echo.

Look at Reverb Setup: Hall04.rvb then Hall05.rvb, notice how only the last tap (with greatest Delay) has moved.

With **Hall04.rvb**, the first 3 taps are the Early Reflections. After the time determined by the last tap, the Early Reflections are repeated and continue to repeat with decreasing gain, specified by the last tap's height. With these settings the Reflections and Echo repeatedly overlap to produce increasing numbers of reflections, and a complex reverberation.

Hall05.rvb, the first 3 taps are in the same place, but the last tap is well separated, the reverberation time is greatly increased. The Early Reflections will be cycled to produce an increasing numbers of reflections, but as the reverberation time is much longer, the repeated Reflections will tend to die away before the next echo is expected, thus producing more discrete, yet complex Echoes.

Note: When designing Reverb effects, do so with a short transient sound.

Noise Reduction

The Noise Reduction algorithm implemented is designed to reduce Broadband noise with minimal reduction in signal quality. Broadband noise, is noise which is present at many frequencies, typically the same frequencies as the signal. This type of noise can not be removed by simple filtering alone; eg tape hiss, noise from fans, machines, sound cards etc.

The **Noise Reduction** is a two step process.
First a 'Noise print' that is representative of the noise to be removed must be taken, a one second sample will do.This noise region is analysed and its spectral make-up established.
Select **Special|Noise Reduction|Sample Noise** to take the 'Noise print'.

Reducing the Noise
Select the Section to be processed, then select **Special|Noise Reduction|Reduce Noise**.
A Noise Reduction dialog box will appear, three parameters can be modified:

Noise Reduction: Set how much to reduce the Noise by.

Threshold Gain: Set the threshold below which the Noise Reduction will be activated.
This adds an offset to the 'Noise print', anything below which will be treated as noise.

Overlap: Controls the precision of the Noise Reduction.

Notes.
Increased Overlap = Increased precision = Increased processing time.
The Noise Reduction itself is computationally very intensive and slow even on the fastest machines.

Hint: When making a recording, record a few seconds of silence before or after whatever you want to record and use this as the 'Noise print'. The Noise reduction also works best if the DC offset is removed first, the algorithm hasn't got infinite resolution either, so sub-sonic noise may still be present. It is recommended that you High Pass filter the signal with a cut-off of ~10Hz.

<u>Digital Mixing Desk</u>

The Digital Mixing Desk is for mixing Audio files in the same manner as you would with an analogue 4 or 16 track mixer. Mixing files can be done on any waveform editor, whilst this is fine for mixing a few short files together it is woefully inadequate for mixing several large files, where each minor adjustment requires manual mixing each file again, this is incredibly tedious.

The Digital Mixing Desk was designed as a fast, effective solution to this problem, allowing infinite fine tuning until the desired mix is achieved.

The Digital Mixing Desk will work on any Windows 95 system with a sound card.

<u>Over View</u>

A lot of effort has gone in to making the interface as intuitive and flexible as possible, so don't be fooled in to thinking 'It's too easy to learn so it can't be any good', that was our aim; minimal learning curve, maximum power.
The emphasis throughout is on maximum quality, both in sound quality and code quality. The Mixes produced are superior to those that can be achieved on most of our competitors software.

<u>The main features:</u>
The Digital Mixing Desk can be used in **Real Time** like a conventional Mixing Desk.

Loading a Track is very fast as all you are doing is providing a pointer to it. No intermediate file is created prior to Mixing, the data is just streamed from it's source when it is needed. For this reason the data should be located on your hard drive.

16 bit, 8 bit, stereo and mono files can all be mixed together, but they must all be of the same Sampling Frequency.

Left and right channels are mixed independently.

A **Delay** unit is available for each track.
If the tracks are slightly out of sync, they can be delayed to realign them, no expensive equipment required. Delays and simple surround effects can be quickly generated, by adding multiple copies of the same track and by adjusting the delays and amplitude of each.

Harmonic **Enhance**ments and **EQ**'ing can be applied to each track.

All processing is done to 32 bit accuracy.

You can perform re-**Dithering** to maintain the maximum Dynamic Range of the signal.

The mixed file can be saved. The mixer settings can also be saved for re-mixing at a later date.

If 15 tracks aren't enough, then the mixing session can be split to create intermediate files which can be further mixed.

Multiple copies of the Mixing Desk can be active.

Note.
The only restriction is that all tracks to be mixed must be PCM formatted *.WAV files of the same Sampling Frequency.

If you need to mix a file with a different Sampling Frequency then you must equalise all the formats first, the GCT Audio File Converter will do this for you, it also provides anti-aliasing and filtering.

Getting Started

The Digital Mixing Desk can support a maximum of 15 track bars (15 channels if all mono, 30 channels if all stereo), when in use the Mixing Desk will look something like a normal analogue mixing desk, with each track having its own volume sliders, a delay setting, looping option, filtering and Enhance effect.

The majority of the controls are visible on screen all the time.

You will notice a **Master** track bar on the right of the screen, this behaves similarly to an individual track bar, but scales the amplitude of the entire group of tracks.

Let's Add a Track.

Click the **Add Track** button, a standard Open dialog will appear, select a *.wav file you have available, a Track bar will now appear (Track 1). Please use a mono track.

This Track bar looks very similar to what you would find on a normal Mixing Desk.

Click the **Solo** button, the raw track will play. Pull the volume slider up and down to adjust the volume. This is just playing the original track, and adjusting the volume of the soundcard. Only when Mixing is activated do we physically alter the data sent to the Sound Card.

Track bar controls.

With the cursor over a Track, it's format and length are visible in the Information panel at the bottom of the screen. If the track is a mono file one volume slider will be present, if stereo then two. The panel to the right of the information panel indicates the format of the destination file, the default is 16 bit mono. This can be changed from the Characteristics menu.

The **Volume** sliders allow each track to have independent volume control.
The **Gang** button effectively ties the stereo volume sliders together so they track each other.

The **Delay** allows the start of each track to be delayed, this can be used for track alignment, to run a sequence of tracks together, to create echo effects etc.

The **Loop** button enables the track to be looped continuously, or a specified number of times. Whilst over a Track right click the mouse to bring up the Loop menu. The default mode is continuous looping.

The **Mute** button allows the track to be muted without effecting the volume sliders.

The **Solo** button allows the track to be played on its own, this behaves differently depending whether mixing is taking place or not. Mixing Off - the raw unaltered file is played; Mixing On - the track is played with the settings as defined by the Track panel details, all the other tracks being kept in sync.

The **EQ** button switches in the Equalisation.
The **Enhance** button applies harmonic enhancement to the track.
Note. These two functions are computationally intensive.

Real Time Mixing.
This is the default mode, the Mixer behaves like a conventional Mixing Desk.

Click the **Mix Tracks** button. Now we are mixing, all the feature on the track bar can now be activated.

Push the volume slider up to +10dB, this may cause the output to clip. You can correct this either by reducing the volume on this track alone, or you can reduce the gain on the Master track bar. Bring the gain back to 0dB.

Note. Internally a Track's data never clips even with too much gain and EQ. Clipping is only caused when the data must be quantised. In the above example +10dB gain could overload the 16 bit range, however since we are using 32 bits we have plenty of headroom to spare.

Click the **Enhance** button, in and out listening to the difference, it sounds louder due to increased harmonic content.
Similarly, you can dial in some **EQ** parameters and switch in and out the EQ function.

To investigate the **Delay** function we need to add another track.
Click **Add Track** button again, and Open the same track again. Track 1 and Track 2 are identical.

Select Characteristics|Stereo, the mono track bars will now become stereo pairs.
Adjust the left/right volume controls, of each track

	Left	Right
Track1	0dB	-120dB
Track2	-120dB	0dB

Click the **Mix Tracks** button.
Clicking **Solo** on track 1, will play to the Left only, and **Solo** on track 2, will play to the Right only, this is how you have set up the track bars.

Now with both tracks active, type 0.1 in the delay of track 2 (100 milliseconds). The effect is significant.

Remember.
The Volume, Delay, EQ and Enhance can all be adjusted in Real time. They are available on each and every track, so each track can be individually tweaked until you have the effect you require.

Note.
If your system will not run in **Real Time** mode without 'stuttering' increase the **Input Block Size**,

(Options|Input Block Size) this helps reduce the disk overheads.

<u>Off Line Mixing.</u>
In the Off Line mode the mixing isn't done in 'real time', you set the track volumes independently then mix the file. This isn't as big a problem as it first appears since the option to create a test file is the default mode, that is you only mix a small section together, once you're satisfied with your set-up, you then do a full mix.

Well, that's basically it.
A fast, powerful mixing tool, with delay lines, EQ, and Enhancing features.

If you are mixing with high sampling rates (>32kHz) then **Dithering**, (Characteristics|Dithering) can be used to extend the Dynamic Range, you will find more information on this elsewhere in the Manual.

Note.
All tracks to be mixed must be of the same **Sampling Frequency.**
If you need to mix a file with a different Sampling Frequency then you must equalise all the formats first, the GCT Audio File Converter will do this for you, it also provides anti-aliasing and filtering.

GCT Audio Wave Editor provides a tool for splitting a Stereo file in to independent mono left and right files, this gives you more control of the mixing process.

<u>Add a Track</u>
This creates a Track bar, and assigns the specified file to it.

EQ - Equalisation, two bands each with 1.5 octave bandwidth.
- Gain adjustment limited to +/- 20dB.
Enhance - Enhances the track by adding more harmonics.

Volume Sliders - control the Track volume Left and Right.
Gang - ties the volume sliders together so they track each other.
Delay - track offset, to allow alignment of files.

Loop - enables track to be looped.
Mute - mutes the track.
Solo - plays only this track.

With the cursor over a Track, it's format and length are visible in the Information panel at the bottom of the screen.

Note.
The EQ and Enhance functions are computationally intensive.

When playing the raw track, if your Sound card doesn't support independent left/right volume control (most don't), then the volume played will default to that of the left channel, if the volume setting is initially zero then the right channel volume is selected.

Mix the Tracks
Mixes the tracks as specified by the Track bar settings.
The **Master** Track to the right of the screen provides a global gain adjustment to the mixed tracks.

In the **Real Time** mode the Mixer behave like a conventional Mixing Desk.

In the **Off Line** mode the mixing isn't done in 'real time', you set the track volumes independently then mix the file. This isn't as big a problem as it first appears since the option to create a test file is the default mode, that is you only mix a small section together, once you're satisfied with your set-up, you then do a full mix.

The results of the Mixing is written to your Hard Disk.

If your system is not fast enough to run the mixing in Real Time, ie you get excessive dropout the file written to the Hard Disk will still be correct.

All tracks to be mixed must be of the same **Sampling Frequency.**

Remove a Track
Remove the specified Track from the Mixing Desk.

The Master Track
The Master Track behaves like any of the other Tracks.
Hitting the **Play** button just plays the Mixed track if one has been created.

The **Echo** button provides a means of adding echo to the grouped tracks.
This tool can be used to create Delay, Echoing and Reverb - type effects.

The **Delay** parameters.
Delay: This is the time between each echo.
A Delay of 50 - 100 milliseconds creates a doubling effect, this can be used to thicken the sound.
With Delay settings of greater than 100 milliseconds, each echo can now be heard as seperate.

Gain: This sets the amplitude of the initial Echo.
Bandwidth: This sets the cutoff frequency for the the Delayed signal, so a real echo can more faithfully be simulated. In a real Echo the higher frequencies are absorbed by the surroundings.

The **Regeneration Gain:** This sets the Gain of each successive echo.
Each Echo will be a percentage of the previous one.
Setting a Regeneration Gain of 0 produces no echo at all, just the primary Delay.
Setting a Regeneration Gain of 100 produces an echo that never dies away.

By Setting a short Delay, high Gain and Regeneration Gain, and maximum Bandwidth, some interesting effects can be produced.

File Menu Commands

Add a Track.
This creates a Track bar, and assigns the specified file to it.

Save As.
Saves the mixed file.

Open mixer settings.
Clear the mixer setup and Load the selected mixer settings.

Save mixer settings.
Saves the mixer settings, for this session.
If you move the Track data off your computer on to a backup medium, then place the associated *.DMD files in the same directory.
The original location of each track is stored in the *.DMD file, but if the Tracks are no longer there the computer will look in the same directory as the *.DMD file.

Close all Tracks.
Closes all the tracks you have been working with.
To remove a single track use the 'Remove Track' button.

Exit.
Closes all the tracks you have been working with and quits the Mixing Desk.

Characteristics Menu Commands

Stereo.
Create a Stereo file.

Mono.
Create a Mono file. Default.

16 Bit.
Create a 16 Bit file. Default.

8 Bit.
Create an 8 Bit file.

Dithering.
Apply re-dithering to the output file.

Options Menu Commands

Test File.
Create a Test File, otherwise do a Full mix.

Set the Test Time.
Defines the length of the Test file.

Delayed Test File.
Create a Delayed Test File.
Lets you test a mixed section starting ss.mmm seconds in to the file.

Set the Delay Test Time.
Set the start point for the Test file, ss.mmm seconds in to the file.
For example, the files may be 5 minutes long, and you want to test the mixing at 2 minutes without mixing the whole lot. Also useful for testing the overlap section when running one file into another.

Input Block Size.
This determines the amount of information read from the disk, it also sets the response time of actions applied to the mixer, default of 1/4 second.
Note. The greater the Block Size the less time is wasted on disk-head movement.
If your system will not run in Real Time mode without 'stuttering' increase the block size.

VU Meter.
This provides a visual representation of the peak value in the Mix, it is inactive while just playing the files.
In the **Real Time** mode it provides a measure of the peak within each Mixed Block.
In the **Off Line** mode it provides a measure of the peak within the Mixed section.
Note. The Mixer clips at 0dB (when the red block is highlighted).

Creating Surround Sound effects

Example 1.
Add a Mono *.wav file.
Add the same track again. Two track bars will now be visible.

Adjust the volume controls. Trk1 (-3dB) and Trk2 (-10dB).

Set the delay of Track 2, to 0.1 seconds.

Mix the tracks.
In the Real Time Mode, you can compare the mixed track against the original by simply depressing the **Mute** button on Track 2 thus removing it from the mix.

Example 2.
Set Characteristics|Stereo. Stereo volume controls will appear.

Adjust the left/right volume controls, of each track.
Trk1 (-3dB, -10dB) - Trk2 (-10dB, -3dB).

Mix the tracks.

Note. The Delay, Volumes and Effects can all be adjusted in Real time, so they can be tweaked until you have the effect you are after.

Appending Tracks

Add two Mono *.wav files. Two track bars will be visible.
Set Characteristics |Mono. A single volume control will be available.
Set then volume control of both tracks to 0dB.

Set the delay of Track 2, to the length of Track 1.
With the cursor over Track 1, it's length is visible in the information panel at the bottom of the screen.

Mix the tracks.

To run the files together even more smoothly, reduce the delay of Track 2 slightly as required, and mix again.

Testing the mixing of the crossover section.
With Options|Test file set.
Set Options|Set Test Time. to approx. 20 seconds.
Set Options|Delayed Test.
Set Options|Set Delay Time. to approx. 10 seconds before the merge point.

Mix the tracks
The delay of Track 2 can be quickly adjusted, and a re-mixed to achieve the required effect.

Alignment

If you need **precision alignment** between the tracks create a Stereo mix so the offset can be accurately measured.

Trk1 Volumes(0dB, -120dB).
Trk2 Volumes(-120dB, 0dB).

Align the tracks by adjusting the delays of the tracks, mix the tracks and save the output file.
Open this file with a suitable Waveform editor such as the **Audio Wave Editor** and precisely measure the offset between the left and right channels. Add this offset to the appropriate track's delay and precision alignment will be achieved.

Maximising Sound Quality

We recommend you use the Digital Mixing Desk as the final mixing tool as this will allow full control of the resultant output file. All the internal processing is done to the highest quality, the resolution is only reduced in the final quantisation process when some of the information must be discarded. If the Sampling frequency is >32kHz, it is recommended that you enable the **Dithering** function Characteristics|Dithering so you retain as wide a Dynamic Range as possible.

Always use 16 bit data, and as high a sampling rate as possible.
If the destination media is CD then use a 44100Hz sampling rate.

Mixing Audio and Midi files.

Normally when you mix Audio and Midi files together it is the Sound Card which is doing the mixing, the results of which can vary significantly. If however you record the output of a Midi file in to a *.WAV file (using Digital loopback if possible) then mix it to the Audio files using the Digital Mixing Desk far superior results can be achieved.

Digitising Signals.

In general, do not use your sound card to record an analogue input, your computer is a very noisy environment. Performance figures for sound cards are usually measured in a clean environments so take the specifications with a pinch of salt.

If you must digitise using the sound card then record the original to tape, then digitise it.
Whilst recording you must try and minimise the controllable sources of computer noise. Here are some suggestions; leave the computer idle (doing nothing but recording), reduce the screen to a flat neutral colour, if your software has fancy VU bars bouncing up and down turn them off (they turn the Red and Green video guns hard on and off, and can add lots of noise to the input signal), it all helps.

If possible, record all files away from the computer, then digitise and transfer to the computer via S/PDIF.
Record the Midi file to *.WAV file by using Digital loopback (the digital data is routed back on the input for saving, without the loss of quality introduced by the D/A - A/D process), some sound cards support this feature.

Dithering

This function is accessible through the **Characteristics** menu.

Dithering is a mathematical technique used prior to quantisation in order to reduce signal distortion from correlated error components, in doing so, a small amount of noise is added to the final signal. With the application of a Noise-shaping filter the perceived degradation caused by the quantisation is reduced by concentrating the noise power at the higher frequencies where the ear is less sensitive.

What ?
Let's run that by you again.
Noise is added to the signal in such a way that it sounds better, even though slightly more noise is present. The additional noise is shaped so that it is shifted to where ear is less sensitive.

How much Noise is added ?
Less than 2 LSB's of noise is added, the exact amount depends on which frequency component of the noise you look at.

Using Dithering.
Dithering should be used on signals sampled at >32kHz since the added noise is moved to the top half of the frequency spectrum, for Fs = 32kHz the noise is increased between ~8kHz - 16kHz (Fs/4 - Fs/2). Since the Ear is most sensitive at lower frequencies, the higher the sampling rate the better the results.

We don't restrict the frequencies at which dithering can be used, as long as you understand the consequences you can make the decision of when, and when not to use dithering for yourself.

You shouldn't really use dithering with 8 bit signals as the addition of ~ 2 LSB's of noise is significant on an 8 bit signal.
However, the 8 bit mode is useful in testing the effect dithering has if you consider the 8 bit signal to be a 16 bit signal but 48dB down. Whilst in the 8 bit mode reduce the signal level by 20dB (effectively a 16 bit signal down 68dB), so the output range is reduced by 10x from (+/-127 down to +/- 12). Switch in and out the Dithering whilst mixing in real time and listen to the difference.

How do I increase the Performance on my system

The Mixing Desk places quite a lot of strain on all the components of your system (except video) . Probably the biggest bottleneck is the Hard Disk, trying to ensure that all the track data is available at the same time for mixing.
Here are some thoughts on reducing this bottleneck:.

I will assume your Hard Drive is in good condition and Defragmented regularly.

1. The Mixing process is not a CPU killer, so you don't need the latest Mega, Giga, Terra FLOP processor to run the Mixing Desk, however you do need at least one big, fast Hard Drive (the more the merrier).

It is the movement of the Disk Heads and accessing the data which is the big problem.
The Options|Block Size reduces Head movement by reading in bigger chunks of data on each read operation, unfortunately the penalty is slower response to the track bar adjustments.

2. Increase the RAM available to Windows to cache the Disk reads.
Reduce the number of programs running, Windows can then use the spare RAM to cache reads from your Hard Drive, this can significantly reduce Disk Head movement. Alternatively increase the RAM available, buy some more.

3. The mixed data is written to the Temp directory. If the Track Data can be brought from a separate disk would help. ie) write to disk C, with all reads from disk D.

4. Do not try this unless you really know what you are doing (I haven't tried this).
Partition your Hard Drive in to as many partitions as your Hard Drive has Heads, (4 heads - 4 partitions). Place your data on a separate partition, now when you read the data the head movement should be minimal.

The **Equalisation** and **Enhancing** are processor intensive, the algorithms implemented are pretty much optimised so I can't help you there.

Testing has revealed that the performance on a P120 with 16MB RAM and a single WD AC31000H Hard Drive (with no fancy partitioning) is more than adequate for all but the most demanding user.

Audio File Converter

The Audio File Converter is a simple tool which lets you transform audio files in to, another format, be it through a change in sampling rate, resolution or number of channels.

Conversion from other common file formats to the Microsoft *.wav format are also provided.
Please read Conversion from other formats.

Anti-aliasing and filtering are provided to try and minimise any unwanted conversion artefacts.

Getting started

The Audio File Converter can transform any PCM encoded *.wav file to another format, be it through a change in sampling rate, resolution or number of channels.

The first thing to do is Open a *.wav file from menu File|Open.

The screen is split in to two halves.
The file's properties are indicated in the upper half of the screen; Sampling Rate, Channels, Resolution, File Size and File Length.

The lower half of the screen contains 12 buttons, each representing a different; Sampling Rate, Channels, Resolution combination, with the resultant (fully converted) file size indicated beside it. The button with the blue text indicates the present or nearest option.

Clicking one of these buttons will perform the file conversion and start playing the file.
Converted files will have the file size text greyed.

The filtering applied will be dependant on the settings in the Options menu.
If Test File is checked a partial conversion will be performed, the length of the test file is defined by Set Test Time. Once your happy with the conversion deselect Test File and do the conversion again, thus generating the complete converted file.

Anti-aliasing and filtering are provided to try and minimise any unwanted conversion artefacts.

File Menu Commands

Open....

Opens an existing file.
The file's properties are indicated in the upper half of the screen.

Save As...

Saves the currently selected file version to a new or existing filename.

Exit.

Closes the track you have been working with and quits the Audio Wave Converter.

Options Menu Commands

Anti-Aliasing.
Filtering only performed on file where a change in frequency takes place. Default On.
This filter should always be enabled, as it removes all frequencies above the Nyquist rate which cause distortion when down converting.

Noise Reduction.
Additional filtering performed after conversion. Default None.
Use this filter when the required output can not be achieved with the Anti-Aliasing filter alone.

Increasing amounts of filtering can be applied from None, Normal, Soft, Softer.
The Soft and Softer settings are overkill but on occasions can be useful.

Test File.
Create a Test File, otherwise create a full file. Default Test File.

Set the Test Time.
Defines the length of the Test file.

Conversion from other formats

A limited amount of support is provided for conversion of non-Microsoft file formats.
If possible the file will be converted and stored in your TEMP directory.

Formats supported.

- AIFF: Conversion from uncompressed PCM format only.
- AU: Conversion from uLaw and PCM formats only.
  (uLaw is common in this format hence the support).
- IFF: Conversion from uncompressed PCM format only.
- VOC: Conversion from uncompressed PCM only.
- WAV: Only standard uncompressed PCM supported.

Limitations of Conversion.

Some of these formats allow fragmented data, I don't.

The header is read and the format of the 1st sound block is assumed for all the data in the file.
If the data was fragmented glitches may appear in the data where the subsequent headers are, these can be removed later with a waveform editor. If data of different formats is stored in the file (a very stupid thing to do), then the output file although playable will probably sound a mess.

Other Information

All working files will be placed in your TEMP directory.

All the program components are placed in the same directory, to ease program removal.

Audio Wave Editor

There is something strange going on...

Other software which accesses the Sound Card can interfere with Recordings made with this software, for reasons which I don't yet understand...
The waveform WILL BE RECORDED CORRECTLY, but won't be allocated the correct Filename.

To minimise this problem (when making Recordings) we recommend:

1, DISABLE the Auto CD Play feature in Windows 95.
2, Load the Sound Card configuration software prior to loading the Audio Wave Editor
(This is only needed if you intend to switch the input source).

This Manual was created with Microsoft Word 6 and converted to Write 3.1 format for distribution.