Cool Edit Pro
Version 1.2 User Guide Addendum

Copyright 1999 Syntrillium Software Corporation.

We haven't been sitting on our hands since we released *Cool Edit Pro* 1.1. Version 1.2 has lots of new major features, minor features, bells and whistles, tweaks, and optimizations. See below for a list of the most important changes.

New Effect: Full Reverb

Cool Edit Pro's Reverb has long been recognized for its natural sound. Our new Full Reverb takes reverb to the next level with unique features like "Early Reflections" and "Coloration" to better simulate custom acoustic environments. You can simulate just about any wall surface materials and other sound-affecting factors by changing the reverb's frequency absorption with a 3-band equalizer in the style of a Parametric EQ.

New Effect: Hard Limiter

Try this one out just before you're ready to burn a CD. The Hard Limiter enables you to drastically attenuate audio above a certain threshold, while leaving all audio samples below that threshold alone. You can also boost amplitude without clipping. For example, you can amplify a piece of audio well beyond where it would normally clip, and the Hard Limiter will soften the regions that would otherwise be clipping, ensuring that the maximum sample amplitude does not go above the given limit. It's a great way to get the most volume possible from your digital audio without reducing perceived dynamic range.

New Effect: Pitch Bender

Wanna play with the turntable speed? Use the new Pitch Bender effect. Just draw the graph the way you want the speed and pitch to vary. You can simulate a record player slowing down, a power-up sound effect, and much more.

RealMedia G2 Support

Cool Edit Pro now supports RealNetworks's latest streaming audio formats with complete RealMedia G2 export capability.

New Effect: DTMF/Notch Filter

Creating custom notch filters is quick and easy with the new DTMF/Notch Filter. You can filter up to six definable frequencies in addition to standard telephone DTMF tones. Use this notch filter when you wish to remove very narrow band frequencies. For example, it can remove a 60Hz hum and associated harmonics while leaving all 59Hz and 61Hz audio untouched.

Better and Faster Click and Pop Eliminator!

The Click and Pop Eliminator is now even easier to use. The new "Auto Find All Levels" button scans your selection for clicks and determines optimum values for all of the Threshold, Detect and Reject settings. You can customize

the Auto Find feature with the "Sensitivity" and "Discrimination" settings. The new "Pulse Train Verification" and "Link Channels" settings provide even more control in finding and eliminating clicks and pops. The Multiple Pass feature is also dramatically faster than in version 1.1.

More Preview!

Reverb, Full Reverb, Noise Reduction, Hiss Reduction, Echo Chamber, and Convolution now have Preview so that you can listen to the affect before applying it to the wave.

New Feature: Delete Silence

You asked for it! Now *Cool Edit Pro* can automatically detect and remove silence between words or other audio. It's ideal for cleaning up voice prompts and speeding up narratives without affecting the foreground audio.

New Feature: Auto-Cue

The new Auto-Cue sub-menu contains features for automatically adjusting the current selection to define a phrase and for marking phrases and beats in the Cue List.

New "Traditional" controls for Dynamics Processing

The Dynamics Processing dialog now has tabs for easier access to controls. If you prefer traditional ratio and threshold parameters to *Cool Edit Pro*'s freeform graph-based controls, click on the "Traditional" tab.

Other changes…

Enhanced Timed Record feature: Now you can record your favorite radio programs automatically! Just enable "Timed Record" and press the Record button in single-track view, and then enter the start time and recording length of the desired program.

Cool Edit Pro loads faster! Version 1.2 scans for DirectX plug-ins only when you tell it to, whereas previous versions scanned each time you loaded the program. This means version 1.2 loads much faster.

Improved Chorus, Flanger, Sweeping Phaser, Brainwave Synchronizer, and Clip Restoration effects with new settings.

Presets for sample rates up to 192KHz: Actually, *Cool Edit Pro* has been able to record and play 192KHz samples since version 1.0, but now you can easily select 96Kz and 192KHz rates with one click when creating a new waveform or session. *Cool Edit Pro* 1.2 can even go as high as 10MHz if your hardware can handle it!

Histogram in Analyze/Statistics: Now you can view a histogram of the RMS amplitude values in the highlighted selection of audio. This can be useful for determining future amplification amounts, dynamics compression amounts, etc.

Support for 29.97 non-drop SMPTE format: The Time Format list now contains: Decimal, Samples, Compact Disc 75fps, SMPTE 30fps, SMPTE Drop (29.97 fps), SMPTE 29.97 fps, SMPTE 25 fps (EBU), SMPTE 24fps (Film), Bars and Beats, Custom. SMPTE 29.97fps is new in version 1.2.

Definable colors for waves in Multitrack View: The new "Hue" setting in the Wave Block Info dialog enables you to organize your waves by color for easier navigation inside a session. Use Ctrl+Shift+Up/Down Arrows for quicker access.

"Sticky" Pan and Volume controls in Multitrack View: Now the Pan and Volume sliders for individual tracks stay visible until you close them.

Multitrack Insert/Delete Time features: You can make or remove empty space easily in your sessions with these options under the Edit menu.

Multitrack Destroy Wave: Don't like the track you just recorded and want to throw it away—for good? Use Destroy Wave and it's gone from both Multitrack view and Edit view.

Solo This Track Only feature: If you want to hear only one track when you press Play, right-click on the track and select Solo This Track Only.

EBU Extensions: The Wave Information dialog now contains European Broadcast Union extensions with these special broadcast fields: Description, Originator, Originator Reference, Origination Date, Origination Time, Time Reference, and Coding History.

Send 24-bit Using 4-byte PCM Format: Some sound cards, like the Yamaha DSP Factory, require 4-byte PCM Format to receive 24-bit data from an application. *Cool Edit Pro* now supports this format in /Options/Settings/System.

Options/Settings/Multitrack – Panning: Now you can select whether to use *Cool Edit Pro*'s classic panning behavior, or new "Equal-power Sinusoidal" behavior, which mimics the style of most hardware mixers by enabling you to make panning changes without changing absolute amplitude of the audio being panned.

Auto Zero-Cross Edits in Multitrack View: *Cool Edit Pro* can automatically zero-cross-adjust edits in Multitrack view to eliminate pops. See Options/Settings/Multitrack.

Lag time, drift correction, and slack time settings for correcting sound card timing problems: In /Options/Settings/SMPTE, you can adjust for small variances in the timing behavior of your sound card.

Drag and Drop from other applications and windows: - any waveforms dragged from an explorer style dialog (or File Open type dialog) can be dropped into the waveforms list to add the waveform to list of currently opened files.

Log Sweep and Overlap options in Generate Tones: Check the log sweep box to sweep the base frequency logarithmically (rather than in linear fashion) when transitioning from one frequency to another. Check the Overlap box in the Source Modulation section to mix the tones on top of the selected audio.

New "Yes to All" and "No to All" buttons save time when saving or closing multiple files.

Now you can downsample to any bit resolution from 4 to 31 bits with the Convert Sample Type feature (Edit menu).

Improved file format support: Raw PCM Waveforms have 24/32-bit file support. Text files have 32-bit file support. AIF files have 24-bit file support.

See below for updates from version 1.0 to version 1.1.

Faster DSP

Up to 300% faster for many of the effects in the Transform menu in Edit view.

New effects

Try out *Cool Edit Pro*'s new Pan/Expand and Clip Restoration effects in Edit view, and Vocoder and Envelope Follower in Multitrack view.

Enhanced DirectX® support

Cool Edit Pro now supports Preview for DirectX/ActiveMovie plug-ins effects using Microsoft DirectX 5.1.

Crash recovery

Now if your computer goes down in the middle of a recording or editing session due to power loss or another cause, *Cool Edit Pro* can automatically restore your session and recover from the crash so that you don't lose your work. The next time you load *Cool Edit Pro* , it will look for your work in progress and ask you if you would like to continue with your work, delete it, or keep it for later restoration.

Better full-duplex operation on different sound cards

Some cards (notably the Sound Blaster® from Creative Labs and the Card D® from Digital Audio Labs) require devices to be opened in a specific order for full-duplex operation. *Cool Edit Pro* 1.1 automatically detects the installed sound card and selects the optimum settings for most cards. You can also change the settings manually in /Options/Settings/Multitrack.

Enhanced Hiss Reduction

Now *Cool Edit Pro* can analyze the sample to be processed and automatically recommend optimum settings.

Wave Grouping in Multitrack

Now you can select multiple waves and right-click on one of them or pull down the Edit menu to define them as a group. When you select one member of a group, the others are automatically selected, and most editing functions (like moving waves) will affect the entire group. You can also change the group color via the right-click menu.

Other changes…

Now you can **pause background mixing** by double-clicking on the background mix progress bar in the lower left corner of the Multitrack view, or by selecting Pause Background Mixing from the Options menu.

There's a new **master volume control** for Multitrack playbacks and mix-downs. Look for it in the upper left-hand corner of the Multitrack view. It works like the other volume controls—left-click and drag up or down for quick changes, or right-click to pop up a volume slider.

Now you can **mix down to Mono** in /Edit/Mix Down in Multitrack view.

Now *Cool Edit Pro* can handle up to 4,095 simultaneously-loaded files, and the aggregate size of all files loaded and in a session is limited only by available hard disk space. Note: the maximum size for a single file is 2 gigabytes. Also, you must have at least as much free hard disk space as the aggregate size of all files before loading. For example, to load a 30-megabyte session, you must have at least 30 megabytes free. More space is required for editing and for use of the Undo feature.

The new **"Radio Industry" option** in /View/Info in Edit view changes the titles of the text edit fields from standard RIFF specifications to Radio Industry titles. Note that this option affects only *Cool Edit Pro*'s user interface; it does not affect the information stored in those fields.

Cool Edit Pro's outstanding dithering just got better! There are five new noise shaping curves (E, E2, and optimized curves for 44.1KHz, 48KHz, and 96KHz audio) in /Edit/Convert Sample Type (Edit view). The 44.1KHz optimized curve in particular can give fantastic results! Try any of these curves when converting down to 16-bit or 8-bit audio—particularly when going from 32-bit 44.1Khz audio to a final production 16-bit version that will be placed on CD. Use the 96KHz optimized curve when mastering to DVD at 16-bit 96KHz.

There are two new **Logarithmic Crossfade curves** in Multitrack view. Also, crossfade curves can span more than one waveform block.

Cool Edit Pro now uses the wheel (middle button) on Intellipoint compatible mice. When you rotate the wheel, *Cool Edit Pro* zooms in when in Edit view and scrolls vertically when in Multitrack view. You can also scroll through the waveform by clicking and dragging the wheel or middle button.

There's a new setting in /Options/Settings/General called "Mouse Wheel Zoom Factor". This represents the amount *Cool Edit Pro* zooms in when you rotate the wheel on Intellipoint compatible mouses. Values from 10% to 80% typically work well. Higher values will zoom in further when you roll the mouse wheel. Roll the wheel forward (away from you) to zoom in, and backward (toward you) to zoom out. The waveform zooms at the point where your mouse is over. If your mouse is outside the window, it zooms in and out from the center of the display.

Another new editing feature: In Multitrack view, if you right-click on a wave and then hold down the Ctrl key before moving the mouse, the movement locks to vertical only, keeping your wave at the same point in time. (Note: if you press Ctrl *before* right-clicking, *Cool Edit Pro* will create a copy of your wave.) Lift up on Ctrl at any time for free movement again, or press it again to lock it vertically again.

DirectX effects now show up in the Favorites dialog so you can add them to your Favorites and make hot key assignments to them.

Favorite hot key assignments now take precedence over general hot key assignments, so if you assign the same key to a favorite and the same key is assigned through the Keyboard Shortcuts, the favorite will execute when you hit that key. (Precedence was the other way around in version 1.0.)

New setting: "Limit Playback to 16-bit" in /Options/Settings/Devices (so that 32-bit Multitrack sessions play correctly on cards that cannot handle 32-bit audio correctly).

If you load a session that contains waves that are hidden or obscured by other waves, *Cool Edit Pro* now asks you if you want to automatically move the hidden waves to new tracks so that you can see them.

![icon] Full Reverb (Transform/Delay Effects Menu, Edit View)

Full Reverb, like our standard Reverb, is used to simulate acoustic space. It is also convolution-impulse-based (like the standard Reverb), which means there are no ringing, metallic, or other artificial sounding artifacts, although specific resonating can be achieved if desired (see Coloration). Full Reverb has some unique features not found in other reverbs, like Perception to simulate room irregularities, and source location to place the " singer" off-center, and have the early reflections realistically model their position within rooms that have acoustically desirable dimensions that you can customize. Practically any wall surface material or other sound-affecting factors can be simulated by changing the reverb's frequency absorption using a 3 band "parametric EQ" style interface (in the Coloration tab).
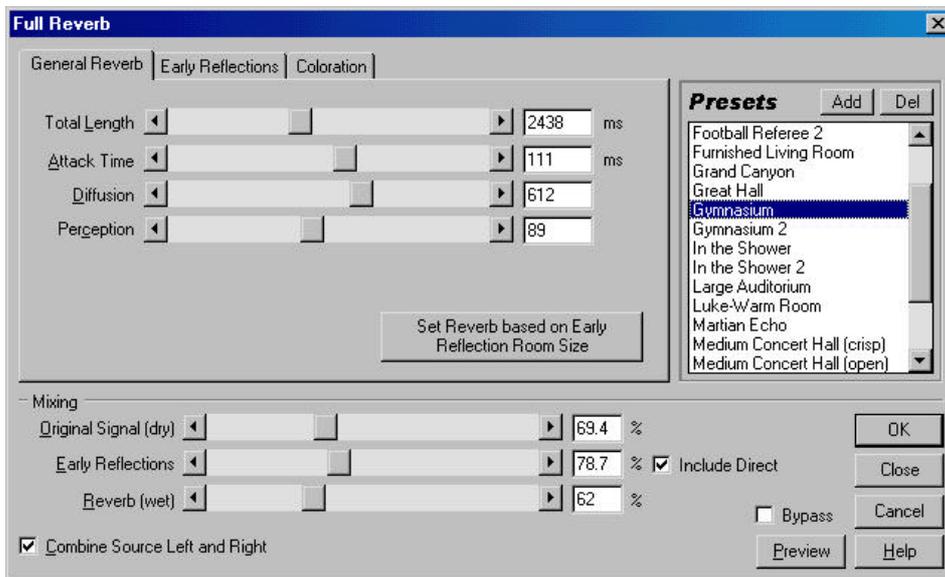
Full Reverb

![note] Whenever any of the reverb characteristics are modified, a new impulse is computed (which can be several megabytes in size) to simulate the environment you are specifying. This means it takes more CPU power to perform the reverbs, but they are much more natural sounding and easier to tailor.

![note] A note on preview: Because every parameter change (other than a mix level change) requires a new reverb 'impulse' to be built, you may have to wait a few seconds after pressing Preview for the reverb generation process. Once built, the preview should generally run in real time, and subsequent previews or pressing OK will not require rebuilding the impulse, nor will adjusting any of the mix level controls or the Include Direct checkbox.

GENERAL REVERB TAB

Total Length

>   This is essentially the time it takes for the reverb to decay 60dB. Depending on the Coloration parameters however, certain frequencies may take longer to decay to 60dB, while other frequencies may decay much faster. Longer values give longer reverb tails, but also take longer to

compute. The effective limit is about 6000ms (a 6-second tail). The actual tail generated is much longer than this in order to allow for decaying into the background noise level.

Attack Time

This is the time it takes for the reverb to build to its maximum amplitude. Generally, reverbs tend to build up over a short time span, then decay at a much slower rate. Interesting effects can be heard with extremely long attack times (like 400ms or more).

Diffusion

This setting controls the rate of echo buildup and how diffuse the echoes are. High diffusion values (above 900) give very smooth reverbs, without distinct 'echoes' heard in them. Lower values will produce more distinct echoes since the initial echo density will be lighter, but the density will build over the life of the reverb tail. Very interesting "bouncy" effects can be obtained by using low Diffusion values and high Perception values. Using low Diffusion values, and somewhat low Perception values with long reverb tails give the effect of a football stadium or similar arena.

Perception

This value models irregularities in the environment (objects, walls, connecting rooms, etc.). Low values will do a smoothly decaying reverb without any frills. Larger values will give more distinct echoes (coming from different locations). If a reverb is too smooth, it may not sound natural, so values up to about 40 will give just enough variation to the reverb to simulate small room variations.

Set Reverb based on Early Reflection

For convincing sounding reverbs, the tail length should be correlated with the feeling of room size given by the Early Reflections. Pressing this button will adjust the Total Length and Attack Time to roughly match the room size given in the Early Reflections tab. After these are set, you are free to fine-tune the length and attack time to change the effect.

Room Size

For convincing sounding reverbs, the tail length should be correlated with the feeling of room size given by the Early Reflections. Pressing this button will adjust the Total Length and Attack Time to roughly match the room size given in the Early Reflections tab. After these are set, you are free to fine-tune the length and attack time to change the effect.

EARLY REFLECTIONS TAB:

Room Size

This is the actual volume of the room in cubic meters. The more volume a room has, the longer the reverbs are for that room. Rooms from only a few meters to giant coliseums can be chosen.

Dimension

The ratio between the room's width (left to right) and depth (front to back) can be chosen. A sonically appropriate height is calculated and reported

below after "Actual Room Dimensions". Very narrow rooms simulate a person being at the other end of a long hallway (and generally sound louder because of all the echoes reflected off the narrow walls). Extremely wide rooms are not very deep, and you can then simulate a person being at one end of the hall or the other using the Location slider. The wider a room, the easier it is to hear the left/right location of the source. Generally, rooms between 0.25 and 4 sound best for nice sounding reverbs.

Left/Right Location

The source can be placed off-center to produce a different set of incoming early reflection echoes. Checking the "Include Direct" box in the Mixing section below will also adjust the Original Signal to sound as if it is coming from the same location by delaying one of the channels. Very nice effects are possible with singers being very slightly off center (between 5% and 10% left or right).

Set Reverb Based on Early Reflection Room Size

Once a room size and dimension are chosen, it is advisable to click this button to set the reverb length and attack time to match the current room size. This will make for a more convincing listening experience. While the reverb length and attack can be adjusted later, this provides a starting point for experimentation.

COLORATION TAB:

Amplitude Sliders

The three amplitude sliders to the right of the graph correspond to the low shelf, mid band, and high shelf bands, from left to right. A quick way to get a different sounding reverb is to just adjust these sliders up or down.

Frequency Sliders (Low/Mid/High)

There are three frequency sliders beneath the graph. The cornering frequency for the shelves, or the center frequency for the mid band, can all be set to any frequency desired. For example, to warm up the reverb, you might want to lower the High Shelf frequency while also reducing its amplitude.

Mid Band Q

This value describes how wide the mid band's affective area is frequency-wise. Lower values affect a wider range of frequencies (hence higher values will narrow the effective range). For distinct resonance, use values like 10 or higher. For just general boosting or cutting a wide range of frequencies, use lower values like 2 or 3. Sometimes it may be desirable to enhance some quality of the audio being processed, like to bring out a singer's voice. You could boost the frequencies just around the natural frequency of their voice to enhance resonance in that range (say in the 200Hz to 800Hz range).

MS (milliseconds) Box

This is the time it will take for the reverb to decay at each frequency following the coloration curve. For more colored reverbs, use lower

settings (like 100 to 250). Values up to 700 work fine. Basically, the lower the value, the more the graph affects the resultant reverb.

MIXING SECTION:

Original Signal (dry)

The amount of Original Signal with respect to the other levels will give the feeling of distance between the listener and the source. Use a low signal level to sound far away, or a high level (near 100%) along with low levels for the other sliders to give the feeling of being close to the source.

Early Reflections

Since these are the echoes that first reach the ear, they give the listener a feeling for the overall room size and dimension. If the mixing for this is set too high though, the result will sound somewhat artificial, but if too low the audio cues for room size will be lost. Generally about half the volume of the original signal is a good starting point.

Reverb (wet)

This controls the volume of the dense layer of sound usually associated with reverb. Like the Early Reflections volume, this slider too should be fine tuned to give a pleasing mix that gives the impression you are after. Too loud, and the reverb sounds very unnatural. The balance between the reverb and the original signal give the impression of distance, so increase the reverb volume with respect to the original signal to make the source sound further away.

Include Direct

With a stereo file, checking this box will cause the original signal's left and right channels to be phase shifted (delayed) slightly to match the direction of the incoming early reflections. The Left/Right Location in Early Reflections determines the direction with which the audio will sound like it's coming from when listened to with stereo headphones.

Combine Source Left and Right

When working with stereo audio, there are actually two source signals (left and right), that can be placed independently. This does, however, require twice as many calculations as a single audio source, so you can check this option for faster processing. When it is checked, Cool Edit Pro sums the left and right channels of a stereo source file before routing to the effects processor. When unchecked, the original signal and its respective reverbed signal remain in its own channel.

If you know that both channels are identical (that is, if they originated from a monophonic sample), you should definitely check this option to minimize processing time.

![icon] Hard Limiter (Transform/Amplitude Menu, Edit View)

The Hard Limiter enables you to drastically attenuate audio above a certain threshold, while leaving all audio samples below that threshold alone. For example, you can amplify a piece of audio well beyond where it would normally clip, and the Hard Limiter will soften the regions that would otherwise be clipping, ensuring that the maximum sample amplitude does not go above the given limit.
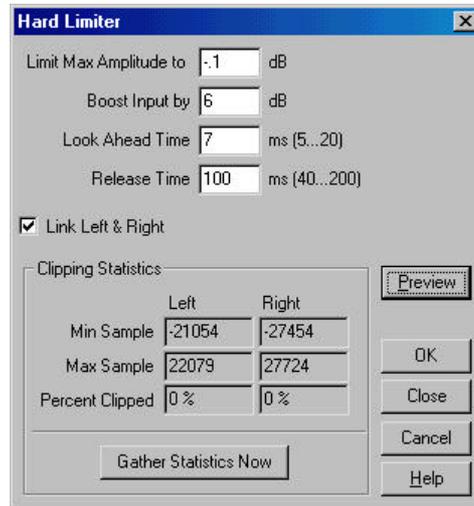
Hard Limiter

This can be extremely handy if, for example, you did a Mixdown to a 32-bit audio file, and only have a few places where the audio is clipping, but do not want to Normalize down the entire waveform (which would reduce the volume of it). Hard Limiting will bring down only those portions that are clipping (or would be clipping if converting to 16-bit).

Limit Max Amplitude to

This is the maximum sample amplitude that will be allowed. To avoid clipping in 16-bit, set this value to about -0.1dB, or -0.5 dB to give a little more headroom for any future edits.

Boost Input by

Before performing the limiting, the audio can be pre-amplified first. To make a selection louder and ensure that no clipping occurs, you can use this to amplify it and soften the regions that would otherwise clip (similar to what they do on TV commercials).

Look Ahead Time

This is the amount of time generally that it will take to attenuate the audio before the loudest peak will be hit. If this value is too small, audible distortion effects may be heard. Use at least 4ms, but up to 10ms is recommended (the default is 7ms).

Release Time

This is the time it takes for the attenuation to rebound back 12dB (or roughly the time it would take audio to resume normal volume if an extremely loud peak was encountered). Generally around 100ms works

very well, and will preserve very low bass frequencies. If this value is too long, audio may just stay very quiet, and not resume normal levels for a while. A value of 2000ms, for example, would give you an Automatic Gain Control effect.

Link Left & Right

If checked, both channels will vary in loudness together. This preserves the stereo image of the source. But it is possible to have the left and right channels vary independently by unchecking this options).

Gather Statistics Now

The Clipping Statistics show what percentage of the audio would clip if limiting were not performed. Click this button to update the statistics if you change any of the input parameters.

⏱ *Pitch Bender (Transform/Time\Pitch Menu, Edit View)*

This function varies the pitch of the source over time. You can use the graph to "draw" your tempo to create smooth tempo changes or effects, such as that of a record speeding up or slowing down.
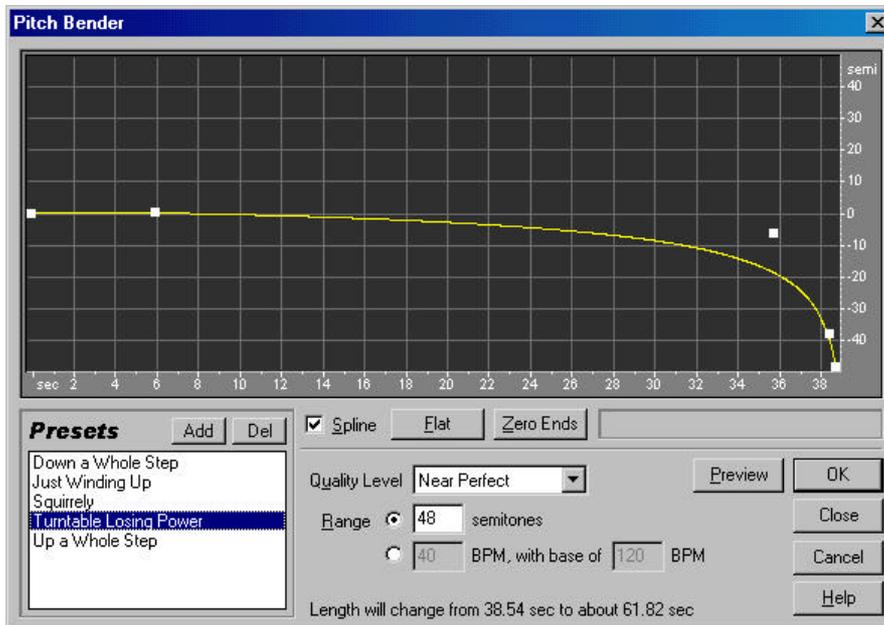
Pitch Bender

Pitch Graph

Use the graph to draw your tempo change over time. Higher points on the curve speed up the sound, lower points slow it down. Choose Spline to do perfectly smooth curves.

Spline

Check this option to generate a smoother, best fit curve instead of a straight line between control points on the graph. When you use spline curves, the line will not ordinarily travel directly through the control points; rather, the points control the shape of this curve. To get the curve closer to a control point, click to create more control points near the point in question. The more control points there are clustered together, the closer the spline curve will be to those points. Use Spline Curves when you want very smooth curves instead of straight lines (with their discontinuities at the control points).

Flat

This will reset the graph, with all points being set to no pitch bending.

Zero Ends

If pitch bending the middle of a selection, it is good to have the endpoints be at the same rate as the surrounding audio. Pressing this button will set the endpoints to no pitch shift.

Quality Level

The lower quality levels have more unwanted harmonic distortion. Usually levels from "Very Good" on up are extremely high quality and you will not

notice any ill effects from harmonic distortion of the signal. Aliasing will still occur however when pitch shifting up, but the higher quality levels will reduce the distortion effects tremendously when shifting the pitch down.

Range

This sets the vertical scale. The range can be set in terms of semitones (12 semitones to an octave) or in beats per minute. In semitones format, the pitch changes logarithmically up and down, while in beats per minute format, the pitch changes linearly up and down. Each format will give a different rate of pitch change for a perfect diagonal line on the graph. The number of semitones to shift up or down can be entered. If doing beats per minute, a range as well as the base tempo need to be entered. You can enter the exact tempo of your piece to change to different rates, but this is not required.

Length Report

At the bottom of the dialog is a report of what the new length of the file will be. When going down in pitch, it is easy to make very long files depending on how large a Range is selected.

DTMF/Notch Filter (Transform/Filters Menu, Edit View)

This filter removes up to six user defined frequencies, in addition to standard telephone DTMF tones. Use this filter when you wish to remove very narrow band frequencies. For example, you can remove a 60Hz hum while leaving all 59Hz and 61Hz audio untouched.
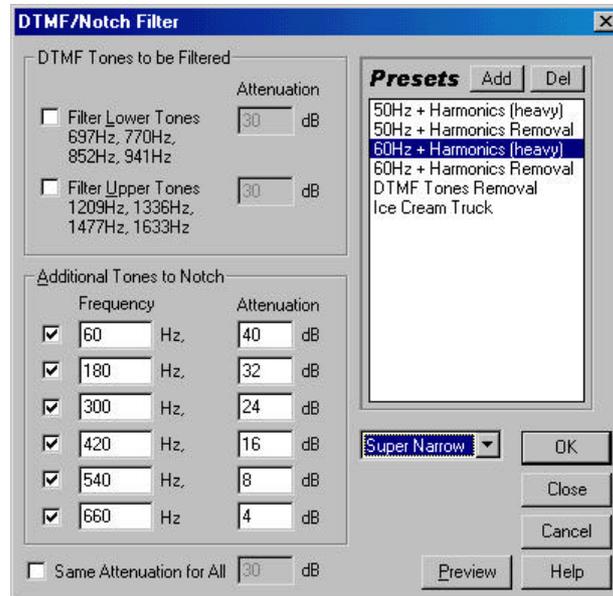
DTMF/Notch Filter

Tones to Notch

Here you may select and enter up to six frequencies to notch. Click in the checkbox to activate the particular frequency. Each frequency can be attenuated separately by un-checking 'Fix Attenuations to' at the bottom and entering the attenuation level desired. DTMF tone-sets (those found in the telephone system) can also be filtered as a group. This is sometimes handy when preparing audio for airing on the radio.

Fix Attenuations to

When checked, all attenuation levels are the same. Uncheck this to set individual attenuation levels for each of the frequencies.

Notch Width

Ranges from Narrow for a 2nd order notch filter (which can remove some side frequencies) to Super Narrow for a 6th order notch filter which is very specific. It is recommended that you use no larger than 30dB attenuations for a Narrow setting, no more than 60dB for Very Narrow, and no more than 90dB for Super Narrow. Any more attenuation for the given notch width will result in a wide rage of neighboring frequencies being removed as well.

⚡ *Click/Pop Eliminator (Transform/Amplitude Menu, Edit View)*

[Note: this description replaces the "Click/Pop Eliminator" topic on pages 98-102 in the *Cool Edit Pro User Guide.*]

The Click/Pop Eliminator works by searching for anomalies in the audio data that could be construed as clicks or pops (Detection), and then replacing or repairing the damaged location (Correction). Using the Click/Pop Eliminator is more accurate than just cutting out the click, or replacing the data with a straight line. You can correct an entire selection or instantly remove a single click if one is highlighted (hint: use a high zoom level to isolate the click). To speed up the repairing of single clicks, configure the F3 key to correct them when they are highlighted by first choosing the Click/Pop filter and pressing "Fill Single Click Now". The F3 key then repeats the last action without bringing up the dialog box. You can also create a quick key for filling in single clicks from the Favorites menu.
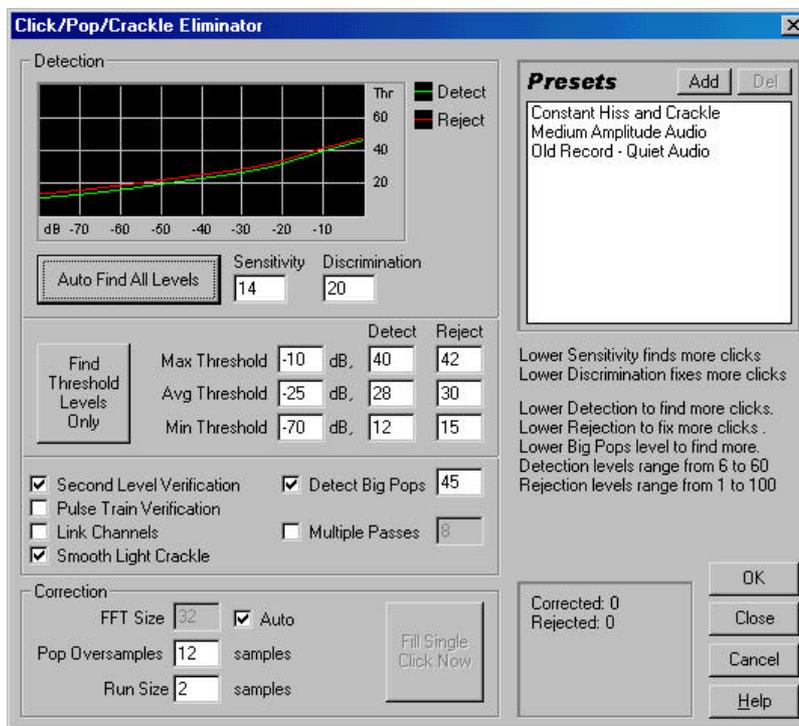
Click/Pop/Crackle Eliminator

🐾 Use the Spectral View feature with the spectral resolution set to 256 bands and a Window Width of 40% to see the clicks in a program. See /Options/Settings/Spectral to adjust these parameters. Clicks will ordinarily be visible as bright vertical bars that go all the way from the top to the bottom of the display.

🐾 To hear all the clicks that were removed, save a copy of the original file somewhere, then Mix Paste it (overlap it) over the corrected audio with a setting of 100% and Invert enabled.

🐾 It may take a little trial and error to find the right settings, but the results are well worth it—much better than searching for and replacing each click individually. The parameters that make the most difference in determining how many clicks are repaired are the Detection and Rejection thresholds (the latter

of which requires Second Level Verification). Make adjustments to these will have the greatest effects; you might try settings from 10 for a lot of correction, 50 for very little correction on the detection threshold, or 5 to 40 on the rejection threshold. The next parameter that affects the output most is the Run Size. A setting of about 25 is best for high-quality work. If you have the time, running at least 3 passes will improve the output even more. Each successive pass will be faster than the previous one.

Thresholds Graph

This graph shows the exact threshold levels to be used at each amplitude, with amplitude along the x-axis (left and right) and the threshold level along the y-axis (up and down). Cool Edit Pro will use values on the curve to the right (above -20dB or so) in processing louder audio, and settings towards the left in processing software sections. There are separate curves for detection and rejection.

Auto Find All Levels

Using the Sensitivity and Discrimination values as guides, pressing this button will scan the highlighted area for clicks, and determine values for all of the Threshold, Detect and Reject values. Five areas of audio will be checked, looking at the quietest regions up to the loudest regions.

Sensitivity: lower values will find more subtle clicks (Detect levels found when using Auto-Find will always be higher than this value). Try a value of 10 for lots of subtle clicks to 20 to set levels for fixing a few louder clicks.

Discrimination: higher values will fix fewer clicks, but leave original audio alone more (by not repairing clicks when there really aren't any). Use values of 20 to 40 if there are relatively few clicks in the audio. Go as low as 4 if there are constant clicks.

Find Threshold Levels

Find Threshold Levels will automatically fill in the Max, Avg, and Min threshold level settings. Use this button to auto-set the threshold levels, before setting the Max and Min detection and rejection levels (it is a good idea to settle on Max and Min settings first, because once they are in place, they usually need not be adjusted much). After Max and Min, set the Avg level to about three quarters the way up between the min and max settings (for example, Max of 30, Min of 10, then try about 25 for Avg). After auditioning a small piece of repaired audio, you can go back and adjust the settings accordingly. For example, if a quiet part still has a lot of clicks, lower the min thresholds a bit. If a loud piece still has clicks, lower the Avg or Max settings. The loud, average, and quiet portions of the audio are handled differently based on the detection and rejection settings.

Threshold Levels (Max, Avg, Min dB)

There are three definable levels (Max, Avg, Min; given in dB) for which you can specify unique detection and rejection threshold settings. These levels are ordinarily set to the maximum amplitude, average amplitude, and minimum amplitude of the audio. For example, if your audio has a maximum RMS amplitude of -10dB then you may set the Max Threshold to -10dB. If the minimum RMS amplitude was -55dB then set Min

Threshold to -55, etc. Press the Find Threshold Levels button to fill in these levels automatically with the maximum, average, and minimum levels. Generally, less correction is required for louder audio, as the audio itself will mask many of the clicks so that their repair is not necessary. Clicks are very noticeable in very quiet audio, so quiet audio will tend to require lower detection and rejection threshold levels. Once the levels are in place, assign appropriate detection and rejection thresholds to be used at each of the levels.

Detection Threshold (1 to 150)

This setting determines how sensitive the filter is in finding clicks and pops, where lower detection thresholds result in more clicks being found. Start with a threshold to 35 for high-amplitude audio (above -15dB), 25 for average amplitudes, and 10 for low-amplitude audio (below -50dB). These settings will find most clicks, and often all of the louder clicks. If there is a constant crackle in the background of your source audio, try lowering the Minimum threshold level more, or increasing the dB level to which the threshold is assigned. This value can be as low as 6, but if the setting is placed any lower, then the filter may removed sounds that are not actually clicks. The more clicks that are detected, the more repair processing there is, which increases the possibility of distortion. With too much distortion of this type, the audio begins to sound flat and lifeless. If this happens, try setting the detection threshold rather low and enabling the Second Level Verification, which will re-analyze the detected clicks and throw out (that is, leave unfiltered) the ones that are less likely to be clicks. Generally, if you still hear clicks after filtering your audio, try lowering the detection threshold, but if this distorts the audio too much, try either increasing the threshold or enabling Second Level Verification.

Rejection Threshold (1 to 150)

If Second Level Verification is enabled, the rejection threshold determines how many potential clicks (found using the Detection Threshold) are rejected. As with detection, lower settings result in more clicks being repaired. Increasing the rejection threshold will prevent more clicks from being repaired, as they may not be actual valid clicks. The idea is to reject as many detected clicks as possible, but still remove all audible clicks. A good starting value for this setting is 30. If a trumpet-like horn sound has clicks in it, and the clicks are not being removed, try lowering this value to reject fewer potential clicks. If a particular sound or instrument is "fuzzing out" and getting distorted, then increase the rejection threshold to keep repairs at a minimum (the fewer repairs that are needed to get good results, the better). Cool Edit Pro may interpret some valid audio as clicks, depending on the nature of the waveform in question. Some waveforms have sharp, albeit periodic, spikes throughout. The rejection algorithm looks to see if this click is isolated (meaning it is a real click) or has neighbors (meaning it is probably part of a desired waveform).

Second Level Verification

Enable Second Level Verification to start rejecting some of the potential clicks that were found by the detection algorithm. In some types of audio,

such as trumpets, saxophones, female vocals, and snare drum hits, some of the peaks in the natural waveform may be detected as clicks. If these peaks are corrected as clicks, the resulting audio may sound muffled. With Second Level Verification enabled, these peaks will be rejected, and not corrected after all. True clicks, on the other hand, will not be rejected. Because this option slows down the restoration, you should first try correcting the audio with this option turned off, and only use it for sections that are very troublesome. See Rejection Threshold for more information.

Pulse Train Verification

Use this option if too many clicks were detected because of a periodic waveform (like trumpet waveform peaks being considered clicks). If this option is used, some valid clicks may not be repaired. Try cleaning up audio without using this option first, and only use it if the clicks are very stubborn.

Link Channels

If checked, audio from both channels is simultaneously analyzed for possible clicks. If a click is found in one channel, the likelihood that a click will be detected in the other channel is increased.

Smooth Light Crackle

This will smooth out one-sample errors when detected, and can have the effect of removing more background crackle. If this makes the resulting audio sound thinner, flatter, or more tinny, leave this option turned off.

Detect Big Pops (30 to 200)

With some audio, large unwanted events (those more than a few hundred samples wide) may not be interpreted as clicks by the normal detection algorithm. Check this option if you want Cool Edit Pro to remove these large clicks as well. Note that a sharp sound like a loud snare drum hit can have the same characteristic as a very large pop, so you should only enable this option if you know the audio has very large pops (like a recording from a record with a very big scratch in it). If you find that obvious, loud pops are not being fixed, then enable this option, and use settings from about 30 (find kind of quiet pops) to 70 or so (only fix the very loudest pops). If this option is enabled, and drum hits sound softer, then the onset of the drum is being considered a big pop and is being "repaired". If this happens, try increasing the threshold a bit until the obvious pops are fixed.

Multiple Pass (2 to 32)

Some clicks may be too close together to be repaired effectively, so going over the audio with another pass will often catch those clicks that weren't found the first time around. If this option is checked, then any number of passes will automatically be performed (up to 32). Cool Edit Pro will perform fewer passes if no more clicks are found and all detected clicks have been repaired. Generally, about half as many or fewer clicks are repaired on each successive pass. When using multiple pass mode, a higher detection threshold may lead to fewer interventions in repairing, and increase the quality of the result somewhat while still removing all clicks. In the same vein, higher rejection thresholds may lead to better results.

FFT Size (8 to 512)

This is the size of the FFT that will be used to repair the clicks, pops, and crackle. In general, it is a good idea to just let the program itself decide on good FFT size choices by checking Auto. You may wish to override the automatic setting and give a specific FFT size to use for click repair when processing some types of audio. A good starting point is an FFT Size of 32. If the clicks are still quite audible, then try increasing this value to 48, then 64, etc. Any number between about 8 and 512 should work fine. The higher this value, the slower the correction will be, but with potentially better results. However, if the value is too high, there may be some extra unwanted low frequency distortion, which sounds "rumbly". If you are repairing clicks one at a time using Fill Single Click Now, then a high FFT size (128 to 256) should work rather nicely.

Pop Oversamples (0 to 300)

This value is used to give the clicks a little extra buffer space when being repaired. When a potential click is found, its beginning and end points are marked as closely as possible. The oversamples value will expand that range, and consider more samples to the left and right of the click as being part of the click. If clicks aren't being corrected fully enough (for example, if clicks become quieter but are still evident), try increasing this value. Start with a value of 8, and increase to slowly as much as 30 or 40. Audio that does not contain a click should not change very much if it is corrected, so this buffer area should remain mostly untouched by the replacement algorithm. Increasing the oversamples will also force larger FFT sizes to be used when in Auto mode. This setting may remove clicks more cleanly, but if set too high, audio will start to distort where the clicks were removed.

Run Size (0 to 1000)

The repair algorithm may not work well on two consecutive clicks that are extremely close together. The Run Size setting sets a ceiling on the potential distance between two clicks—that is, any two clicks closer together than the Run Size will be treated as one. A good starting point would be to set this value somewhere around 25 (or half the FFT size if you are not using the Auto setting). If Run Size is too large (over 100 or so) then the corrections may become more noticeable, as very large blocks of data are repaired at once. If you set the Run Size too small, then clicks that are very close together may not be repaired completely on the first pass.

Fill Single Click Now

You can correct a single click by highlighting it and pushing the Fill Single Click Now button. This is the manual override for click restoration. If FFT Size is set to Auto, then an appropriate FFT size will be used for the restoration based on the size of the area being restored. Otherwise, settings of 128 to 256 work very well when filling in single clicks. Once a single click is filled, the F3 (Repeat Last Command key) key will be set to repeat this action, so filling in future clicks requires only pressing F3.

Previously Corrected/Rejected

After running a click repair, re-enter the Click/Pop Eliminator to display the number of previously corrected clicks. The dialog will also show the number of rejected clicks that would have been corrected if Second Level Verification were not enabled.

Dynamics Processing (Transform/Amplitude Menu, Edit View)

[Note: this description replaces the "Dynamics Processing" topic on pages 65-69 in the *Cool Edit Pro User Guide.*]

The dynamics processor varies the output level of a waveform, based on its input level. This lets you limit or compress the dynamic range of a sample so that the perceived loudness is kept below a defined limit, or so that the waveform's overall dynamic range is kept at roughly the same level. You can also expand or gate the signal so that low-level signals are reduced in level, thereby increasing the perceived dynamic range, or so that signals that fall below a certain threshold [i.e. noise] are eliminated. Control of the Dynamics Processor is done through any of the four tabbed sheets:
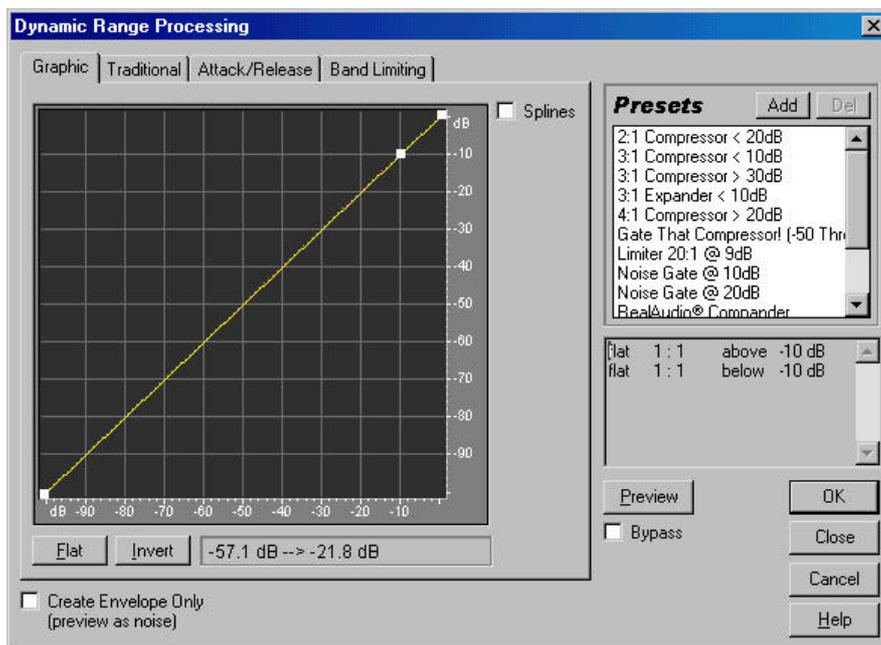
Dynamic Range Processing

GRAPHIC TAB:

Graph

The graph depicts input level along the x-axis (left and right) and the new output level along the y-axis (up and down). A line that flows directly from the lower-left to the upper-right (default) depicts a signal that has been left untouched, since every input value goes to the exact matching output value. Adjusting the shape of this line will adjust the input or output assignments, thereby altering the dynamic range. For example, you can boost all input that has a level of around -20dB, leaving everything else unchanged. You can also draw an inverse line (a line from upper-left to lower-right) that will dramatically boost low amplitudes while dramatically suppressing high amplitudes (that is, all quiet sounds will be loud, and all loud sounds will be quiet.)

🔦 To add a point to the graph, click in the grid at the location where you would like the point placed.To move a point on the graph, click and hold on the point and drag it to a new location. When the mouse cursor is located over a point, you will see it change from an arrow to a hand.To

numerically enter input and output signal levels for an Edit Point, right click on the point to bring up the edit box

Splines

Check this option to generate a smoother, best fit curve instead of a straight line between control points on the graph. When you use spline curves, the line will not ordinarily travel directly through the control points; rather, the points control the shape of this curve. To get the curve closer to a control point, click to create more control points near the point in question. The more control points there are clustered together, the closer the spline curve will be to those points. Use Spline Curves when you want very smooth curves instead of straight lines (with their discontinuities at the control points).

Flat

The Flat button resets the graph to its default state of an unchanged signal, removing all control points.

Invert

The invert button changes the graph to one that will function as the exact opposite. For example, if a transfer function with a compressor characteristic is being displayed, pressing Invert will change the graph to one with the corresponding expander characteristic. For a graph to be invertable, it must have points in the two default corners (-100,-100 and 0,0) and its output level must always increase from left to right (i.e. each Edit Point must be higher than the one to its left).

TRADITIONAL TAB:

The Traditional Tab lets you enter ratios and thresholds numerically. You can enter compressor settings for up to six sections or stages, each with its own ratio and threshold setting. The threshold settings must get lower in value down the list. Simply enter the data to match the compressor you want. For example, to create a 3:1 compressor above -20dB, choose 'Compress', enter 3 (for 3:1) in the 'Above' line, and -20 dB for the threshold. If you want to expand 2:1 below -20dB as well, choose 'Expand' and enter 2 in the first 'Below' line. Additional ratios at thresholds can be entered farther down the list.

ATTACK/RELEASE TAB:

Output Gain

This is a gain applied to the output signal and is the last step performed on the audio.

Attack Time (Out)

This is the attack time applied just before output. Attack time determines the time in milliseconds that it takes for the processed output signal to reach its specified output volume. If there is suddenly a quiet portion that drops 30dB, it will take this much time before the output actually drops to its corresponding volume level. If the sum of Attack and Release times is too short (less than about 20 ms total), audible effects, such as a "vibrating" sound, can be heard at around 1000 Hz/ms total. For example,

if the Attack and Release times are each set to 5 ms (making 10 ms total), then a vibrating sound at 100Hz can be heard. Thus, a total value of about 30 ms is about as low as you can go without introducing these effects.

Release Time (Out)

This is the release time applied just before output. Release Time is the time it takes the end of a previous output level to reach the specified output volume. For example, where the Attack is the time it takes for the start of a pulse to reach the desired output volume, the Release is the time it takes for the end of the pulse to reach the desired level.

Joint Channels

In Stereo, each channel can be compressed independently, sometimes causing the surrounding background noise to get louder on one channel at a time. This can sound strange. For example, a loud drum beat in the left channel will make the background noise sound louder in the right than in the left. If Joint Channels is checked, both channels are used to find a single input dB value, and both channels will be amplified together by the same amount (thus preserving the stereo center-channel image). For example, a loud drum beat on the left channel will also cause the right channel to be reduced in level by an equal amount.

Input Gain

This is the gain added to the signal before it goes into the Level Detector (the section that detects the current level). This essentially "pushes" the graph up or down by the gain given.

Attack Time (In)

This is the attack time applied when retrieving the current amplitude information. Attack time determines the time in milliseconds that it takes for the processed output signal to reach its specified output volume. If there is suddenly a quiet portion that drops 30dB, it will take this much time before the output actually drops to its corresponding volume level. If the sum of Attack and Release times is too short (less than about 20 ms total), audible effects, such as a "vibrating" sound, can be heard at around 1000 Hz/ms total. For example, if the Attack and Release times are each set to 5 ms (making 10 ms total), then a vibrating sound at 100Hz can be heard. Thus, a total value of about 30 ms is about as low as you can go without introducing these effects.

Release Time (In)

This is the release time applied when retrieving the current amplitude information. Release Time is the time it takes the end of a previous output level to reach the specified output volume. For example, where the Attack is the time it takes for the start of a pulse to reach the desired output volume, the Release is the time it takes for the end of the pulse to reach the desired level.

Peak

Peak mode is a graph interpretation method that is a little more difficult to use than RMS, and is a bit outdated. It equates to the RMS value times

two. That is, if the RMS value is -20dB, then the equivalent peak value will be -40dB. This occurs because the RMS value calculated was mapped to a peak sample value for output. This method is basically here for backward compatibility.

RMS

This is a new graph interpretation method that more closely matches the way people hear volume. This mode causes the output to be exactly the RMS amplitude that is specified in the graph. For example, a limiter (flat horizontal line) at -10dB causes the RMS amplitude of the result to average -10dB (where 0dB is a maximum amplitude sine wave without clipping).

Lookahead Time

Lookahead Time is used to handle sharp spikes that may occur at the onset of a louder signal. At times, and for brief instances, these onsets can go beyond the limits of your compressor settings, which may be desirable in certain compression scenarios since it can enhance the impact of, say, a drum hit. However, this is obviously not desirable if you are using limiting in order to reduce the maximum amplitude of the audio.

🎤 The spikes occur because it takes a little time to determine (the Level Detector's attack value) and react (the Gain Processor's attack value) to the current signal level, so Lookahead Time will actually cause the attacks to start before the audio gets loud, instead of right on top of the transient. Otherwise, with a Lookahead Time of 0, a spike will stay loud until all of the attack times have elapsed.

BAND LIMITING TAB:

Low Cutoff

This is the lowest frequency that dynamics processing will affect. You can define a band, or range, to which compression or expansion is applied, within the current frequency range. The Low Cutoff point is the bottom boundary in defining this band. For example, with values of 1000 for Low Cutoff and 5000 for High Cutoff, dynamics processing only affects audio in the frequency range of 1kHz to 5kHz. To use the entire frequency range of the source material, leave this setting at 0.

High Cutoff

This is the highest frequency that dynamics processing will affect. The High Cutoff point is the top boundary in defining this band. For example, with values of 1000 for Low Cutoff and 5000 for High Cutoff, dynamics processing only affects audio in the frequency range of 1kHz to 5kHz. To use the entire frequency range of the source material, leave this setting at 0. To use the entire frequency range of the source material, this setting should be at 1/2 the current sample rate (24000 for 48kHz, 11025 for 22kHz, etc.).

🎤 Leaving this setting at the default of 24000 will work in affecting the entire frequency range for all sample rates below 48kHz.

Create Envelope

You can use this option to apply any dynamics processing, and have the result returned as an amplitude envelope. This envelope can then be copied and used with /Edit/Mix Paste (modulate) to modulate it with another sound's amplitude. Of course, it is also useful for just seeing what the amplitude envelope of your audio is.

Delete Silence (Edit Menu, Edit View)

Use this to remove periods of silence between words or other sounds.
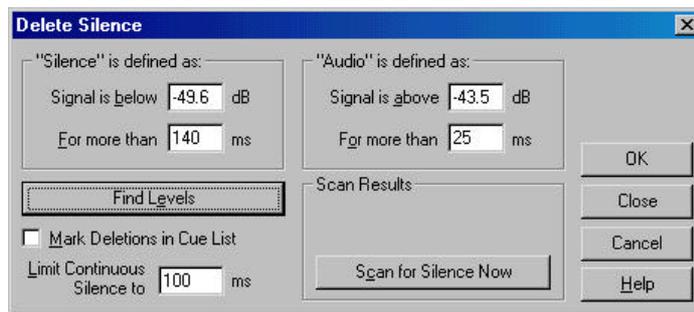
Delete Silence

Silence Definition

If audio is below the given threshold for more than the given time period, that spot in the audio will be considered silence. For very quiet high quality audio, the dB value will be lower (like -60dB). For noisier audio, the value may be much higher (like -30dB). Use higher values for the time to keep groups of words together, for example.

Audio Definition

If audio is above this given threshold for more than the number of milliseconds given, audio will be considered valid, and not silence. Use higher values for the time to ignore short periods of audio (like clicks, static, or other noise). If this value is too high however (above 200ms), short words may be skipped.

Find Levels

To make the choice of signal levels easier, press this button to scan the highlighted selection and determine a good starting point for signal levels. If words or phrases are getting chopped off, try lowering the signal level values. If not enough silence is getting removed, try increasing the signal level values.

Mark Deletions in Cue List

If checked, each location where silence was removed, will be marked by adding it the cue list.

Limit Continuous Silence

Sets a minimum amount of silence to keep at all times. If a silent region of less than this length is found, it will not be removed at all. If a silent region of more than this length is found, enough silence is removed so that there is exactly this length of silence left. Set this value to zero to remove as much silence as possible.

🔌 A setting of 150ms or so will leave a nice natural pause where the larger region of silence was found (making a speaker sound more natural). If too much silence is removed, the speaker will sound very artificial.

Scan for Silence Now

This pre-performs the silence deletion, but instead of actually deleting, it reports back how much silence was removed, and how many sections of silence were found.

Auto-Cue (Edit Menu, Edit View)

Use Auto-Cue to define spoken/musical phrases or beats, and add them to the Cue List automatically.

To use Auto-Cue, first select the general region in which you wish to find phrases or beats. Then choose /Edit/Auto-Cue and the appropriate submenu choice from among the following options:

Adjust Selection to Phrase

Selects phrase within the selected region by adjust the highlight inward, ignoring any silence before and after the audio. This does not add to the Cue List.

Find Phrases and Mark

Scan highlighted area, marking non-silent regions in the cue list.

Find Beats and Mark

Scan highlighted area, marking all beats found in the cue list.

Auto-Cue Settings

Select this option to customize the way the Auto-Cue function works.
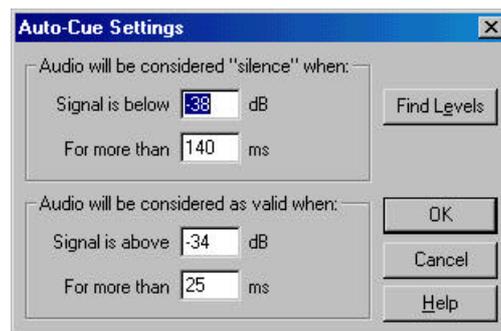
Auto-Cue Settings

Silence definition: if audio is below the given threshold for more than the given time period, that spot in the audio will be considered silence. For very quiet high quality audio, the dB value will be lower (like -60dB). For noisier audio, it may be much higher (like -30dB). use higher values for the time to keep groups of words together for example.

Audio definition: if audio is above this given threshold for more than the number of milliseconds given, audio will be considered valid, and not silence. Use higher values for the time to ignore short periods of audio (like clicks, static, or other noise). If this value is too high though (above 200ms), short words may be skipped.

Find Levels: to make the choice of signal levels easier, press this button to scan the highlighted selection and determine a good starting point for signal levels. If words are getting chopped off, try lowering the signal level values. If not enough silence is getting removed, try increasing the signal level values.

⋈ Pan/Expand (Transform/Amplitude Menu, Edit View)

The Pan/Expand effect is used to pan, or shift, the center channel of a stereo waveform. It also enables you to expand or narrow the stereo separation of the left and right channels.
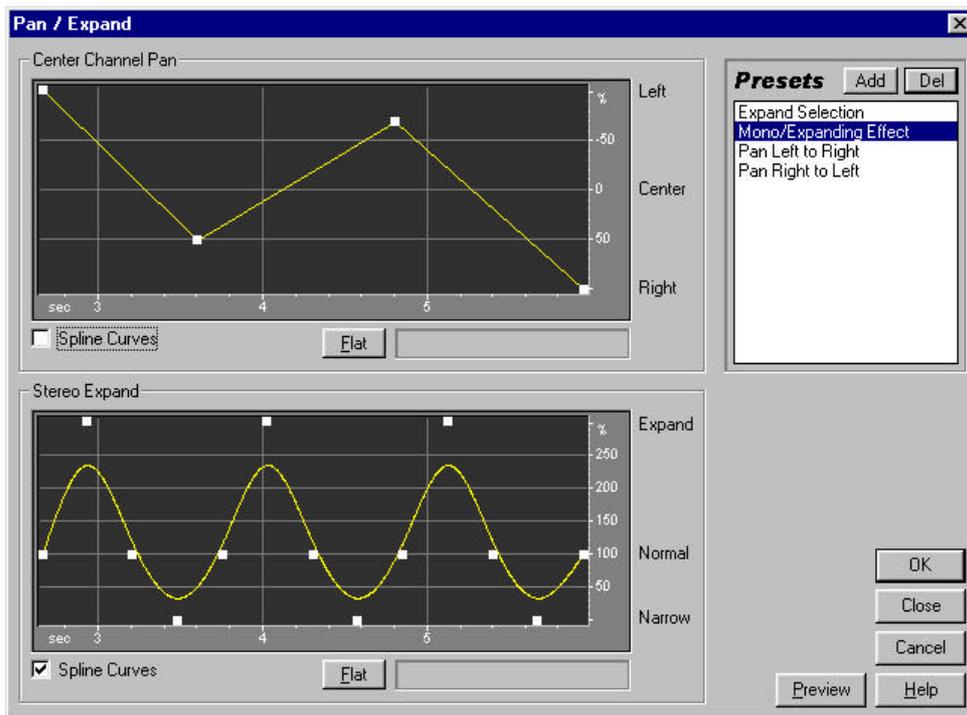
Pan/Expand

Center Channel Panning makes use of the "surround" and "center" channels of a stereo recording, where the "surround" channel is the difference (L-R) of the two original channels, and "center" is the sum (L+R) of them. In this case, you can think of a stereo recording as having 4 channels (left, right, middle, surround), and this effect will pan these channels around. For example, you can pan hard left to get the old "center" channel coming out the left speaker, and the old 'surround' channel coming out the right. This type of panning can provide realism to original stereo recordings.

Expanding works by subtracting out or adding differing amounts of left and right channel signal, so things occurring on the right or left are boosted or cut.

Both of these elements can be altered dynamically, over time, by using the respective graph.

Center Channel Pan

This graph represents the pan position of the center channel of a stereo waveform over time. You can use the graph to position the center channel anywhere from hard left (-100%) or right to hard right (100%), with the corresponding surround channel moving right to left in the opposite direction. Use this for panning original stereo data more realistically than amplitude panning.

To add a control point to the graph, click in the grid at the location where you want to place the point. To remove a control point, drag it off the

graph area.To enter time and panning for a control point numerically, right click on the point to bring up the edit box, or double-click on the curve.To move a point on the graph, click and hold on the point and drag to a new location. When the mouse cursor is located over a point, you will see it change from an arrow to a hand.

Stereo Expand

Stereo Expand shows the expand level over time. This either amplifies (>100%) or removes (<100%) the differences between channels. With some material, a stereo expanding effect can be achieved by increasing the differences between the left and right channels. The expansion level can vary over time for interesting effects (growing from a mono signal to a very wide signal for example).

To add a control point to the graph, click in the grid at the location where you want to place the point. To remove a control point, drag it off the graph area.To enter time and expansion for a control point numerically, right click on the point to bring up the edit box, or double-click on the curve.To move a point on the graph, click and hold on the point and drag to a new location. When the mouse cursor is located over a point, you will see it change from an arrow to a hand.

Spline Curves

Check this option to generate a smoother, best fit curve instead of a straight line between control points on the graph. When you use spline curves, the line will not ordinarily travel directly through the control points; rather, the points control the shape of this curve. To get the curve closer to a control point, click to create more control points near the point in question. The more control points there are clustered together, the closer the spline curve will be to those points. Use Spline Curves when you want very smooth curves instead of straight lines (with their discontinuities at the control points).

Flat

The Flat button resets the graph to its default state by removing all control points.

⚡ *Clip Restoration (Transform/Noise Reduction Menu, Edit View)*

The Clip Restoration filter can remove the evil occurrences of clipping from your waveform. It does this by actually "filling in" the clipped segments of the waveform. Clipping is a phenomenon that occurs in digital audio when the amplitude value of a signal exceeds the maximum level that can be represented by the current bit resolution (i.e. 256 steps in 8-bit audio). This can happen when your source recording levels are simply too high. Clipping causes the signal to distort, and appears in the waveform display as a "chopping-off" of the top or bottom of the waveform. The audible result is a static-like distortion.
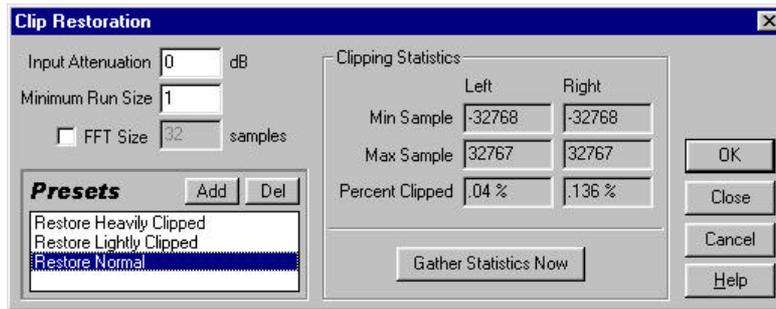
Clip Restoration

 Do not use DC Bias Adjust on any audio you wish to run the Clip Restoration filter on (you can do a DC adjust after de-clipping, of course), because DC Bias Adjusted material modifies the clipped area and makes it unrecognizable as being clipped.

 You can monitor clipping occurrence during playback by watching the level meters under the waveform display. If clipping occurs, the box(es) on the far right of the LEDs will light up.

Input Attenuation

Set this parameter to the amplitude threshold below which you want to remove all audio. Generally, values from -45dB (for very noisy audio) to -100 (for less noisy audio) work best. Try view the hiss level using Spectral Analysis to find the hiss level.

Minimum Run Size

Larger FFT sizes result in a lower detected hiss level, because the hiss will be spread across more frequency bands (approximately 3dB per doubling of FFT size). Use smaller FFT sizes for faster processing.

FFT Size

Precision Factor determines the accuracy of hiss reduction in the time domain, with larger values producing better results, but also at slower processing speeds. With lower precision factors, you may get a few milliseconds of hiss before and after the louder parts of audio; with higher values, this is reduced. Values higher than 40 or so will not ordinarily improve quality.

Clipping Statistics

These fields show the minimum and maximum sample values found in the currenlt selected regions, as well as the percent of samples clipped based on that data.

Gather Statistics Now

Pressing this button will fill in the *Clipping Statistics* fields with information for the current selection or file.

Hiss Reduction

This function removes all audio in all frequencies that are below a certain threshold (generally the noise floor or hiss level). If audio has a constant background hiss (noise can removed more effectively than hums or whines) then that hiss can be "noise gated" or removed completely. Any audio in any frequency band louder than the cutoff level will remain untouched. Use Hiss Reduction to literally carve out annoying hiss from cassette recordings, record albums (after click/pop removal), or microphone recordings.
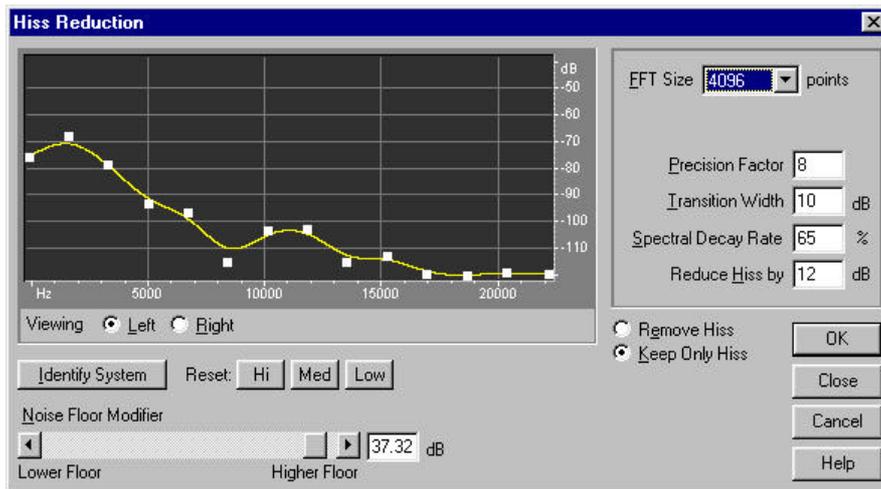
Hiss Reduction

To use, choose an FFT size first. Then either reset the noise floor graph to Hi, Med, or Low, depending on the estimated noise level of the recording or find a section relative free of audio that contains mostly hiss and click Identify System (see Identify System for details). Finally choose how much to reduce hiss by, and try a few times reducing hiss from a short sample with different noise floor modifier settings. When you are satisfied with the results, process the entire area you want to remove hiss from.

FFT Size

Lower FFT sizes (2048 and below) will have better time response (less swooshing before cymbal hits for example) but poorer frequency resolution (making things sound hollow or flanged). Higher FFT sizes (12000 and above) may cause swooshing effects, reverb, and long drawn out background tones, but be very accurate in the frequencies processed. In general, sizes from 3000 to 6000 work best.

Precision Factor

Precision Factor determines the accuracy of hiss reduction in the time domain and affects how quickly spectral components below the previous hiss level decay (see Spectral Decay Rate). With larger values generally producing better results, but also at slower processing speeds. With lower precision factors, you may get a few milliseconds of hiss before and after the louder parts of audio; with higher values, this is reduced. Values higher than 20 or so will not ordinarily improve quality any further. Usual settings range from 7 to 14.

Transition Width

This setting produces a slow transition in hiss reduction, with no reduction taking place at the level specified for Remove all hiss below a level of: above, to 100% reduction at this many dB over that setting. For example, a hiss level of -76dB and a transition width of 4dB will remove 100% of the hiss below 80dB, but only 50% of the hiss at -78dB. If the transition width is too small, other background artifacts may be heard, such as little "tinkles" in the background. If the transition width is set too high, some hiss may remain after processing.

Spectral Decay Rate

When audio above the noise floor is encountered, more audio in the same frequency band is assumed to follow. With low decay rates, less audio is assumed to follow, and the carving function will cut more closely in time to the frequencies being kept. With rates too low, background bubbly effects might be heard, and music may start to sound artificial. With rates too high (above 90%) unnaturally long tails and reverbs might be added. Values of 40% to 75% work best.

Reduce Hiss By

When audio components are found below the given noise floor settings, they are reduced up to this amount. With lower reduction amounts (below 6dB) not as much noise is removed, but the original audio signal stays relatively undisturbed. With higher levels (especially above 20dB) dramatic hiss reduction can be achieved, but the remaining audio may start to be distorted.

Identify System

This is the most powerful feature of Hiss Reduction. If a good estimate of what the noise floor looks like (i.e. how it is shaped) is graphed, then the algorithm can more effectively remove just hiss while leaving regular audio untouched. If a section of just hiss only is available, then clicking Identify System when just the hiss is highlighted will design a graph that most accurately reflects the noise floor. If no such area can be found, then just using an area with the least amount of music will work fine if some extra steps are taken afterwards. For best results then, choose an area of audio that has the least amount of high frequency information (if viewing in spectral view, the area will not have any activity in the top 3/4ths of the display). After pressing noise floor has been identified, you may need to skew the graph on the left hand side (representing the lower frequencies) downward some, so as to make the graph as flat as possible. If music was present at any frequency, the dots on the graph around that frequency will be higher than they should be and should be moved down.

The line on the graph represents where the estimated noise floor exists for each frequency in the source, and the information will be used to determine if audio is hiss or real audio data. The actual value used when doing the hiss reduction will be a combination of this graph and the Noise Floor Modifier which is used to basically shift the estimated noise floor reading up or down for fine tuning.

Reset (Hi, Med, Low)

For quick general purpose hiss reduction, a complete noise floor graph is not always necessary. To clear the graph, press Hi to assume the noise floor is at -50dB (if hiss is very loud). Press Med to set the floor to -70dB (average hiss) or Low to set the floor to -90dB (if there isn't much hiss at all). Then fine tune the noise floor setting using the Noise Floor Modifier slider until the appropriate amount of hiss reduction and quality level is achieved. The combination of the noise floor graph and the noise floor modifier replaces the older setting for "Remove all hiss below a level of" in previous versions.

Remove Hiss/Keep Only Hiss

For normal operation, choose Remove Hiss. If for some reason you want to extract the hiss only, choose Keep Only Hiss, and all the audio will be removed, leaving only hiss.

Envelope Follower (Transform Menu, Multitrack View)

The Envelope Follower varies the output level of one waveform, based on the input level of another. The amplitude map, or envelope of one waveform (the Analysis wave) is applied to the material of a second waveform (the Process wave), which results in the second waveform taking on the amplitude characteristics of the first waveform. This lets you, for example, have a bass guitar line which only sounds when a drum is being hit. In this example you would have the drum waveform as the Analysis wave, and the bass guitar waveform as the Process wave.

In addition to applying an amplitude envelope to a waveform, you can also alter the dynamic properties of the result with a variety of settings to otherwise expand, gate, compress, or limit the resulting signal.
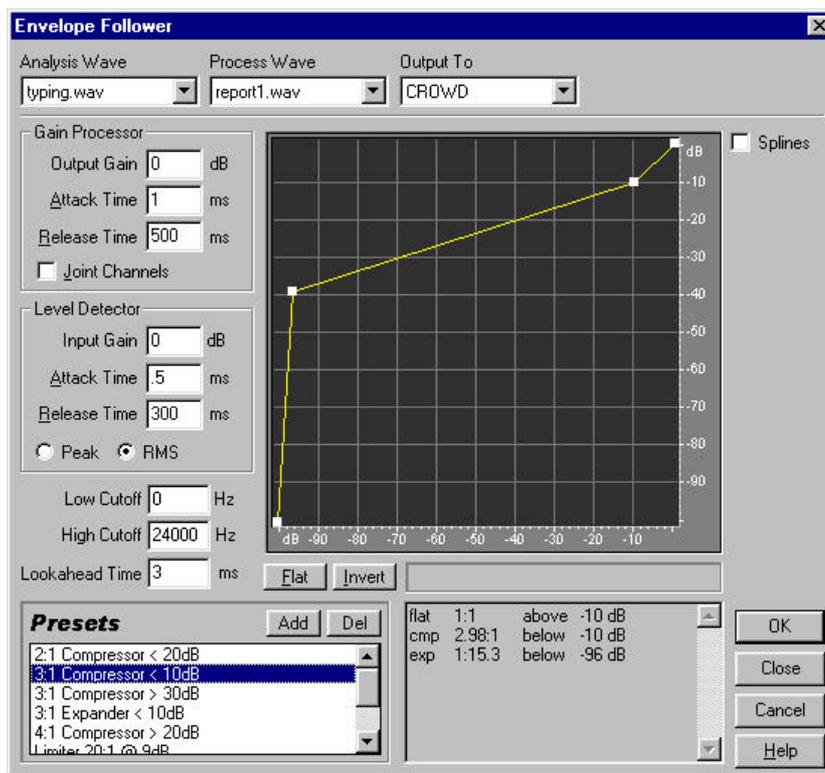
Envelope Follower

Analysis Wave

Choose the waveform from which to "read" the amplitude envelope.

Process Wave

Choose the waveform containing the material to apply the amplitude envelope to.

Output To

Choose the track to output the resulting waveform to. This defaults to the next available empty track.

Output Gain

This is a gain applied to the output signal and is the last step performed on the audio.

Attack Time (Out)

This is the attack time applied just before output. Attack time determines the time in milliseconds that it takes for the processed output signal to reach its specified output volume. If there is suddenly a quiet portion that drops 30dB, it will take this much time before the output actually drops to its corresponding volume level. If the sum of Attack and Release times is too short (less than about 20 ms total), audible effects, such as a "vibrating" sound, can be heard at around 1000 Hz/ms total. For example, if the Attack and Release times are each set to 5 ms (making 10 ms total), then a vibrating sound at 100Hz can be heard. Thus, a total value of about 30 ms is about as low as you can go without introducing these effects.

Release Time (Out)

This is the release time applied just before output. Release Time is the time it takes the end of a previous output level to reach the specified output volume. For example, where the Attack is the time it takes for the start of a pulse to reach the desired output volume, the Release is the time it takes for the end of the pulse to reach the desired level.

Joint Channels

In Stereo, each channel can be compressed independently, sometimes causing the surrounding background noise to get louder on one channel at a time. This can sound strange. For example, a loud drum beat in the left channel will make the background noise sound louder in the right than in the left. If Joint Channels is checked, both channels are used to find a single input dB value, and both channels will be amplified together by the same amount (thus preserving the stereo center-channel image). For example, a loud drum beat on the left channel will also cause the right channel to be reduced in level by an equal amount.

Input Gain

This is the gain added to the signal before it goes into the Level Detector (the section that detects the current level). This essentially "pushes" the graph up or down by the gain given.

Attack Time (In)

This is the attack time applied when retrieving the current amplitude information. Attack time determines the time in milliseconds that it takes for the processed output signal to reach its specified output volume. If there is suddenly a quiet portion that drops 30dB, it will take this much time before the output actually drops to its corresponding volume level. If the sum of Attack and Release times is too short (less than about 20 ms total), audible effects, such as a "vibrating" sound, can be heard at around 1000 Hz/ms total. For example, if the Attack and Release times are each set to 5 ms (making 10 ms total), then a vibrating sound at 100Hz can be heard. Thus, a total value of about 30 ms is about as low as you can go without introducing these effects.

Release Time (In)

This is the release time applied when retrieving the current amplitude information. Release Time is the time it takes the end of a previous output

level to reach the specified output volume. For example, where the Attack is the time it takes for the start of a pulse to reach the desired output volume, the Release is the time it takes for the end of the pulse to reach the desired level.

Peak

Peak mode is a graph interpretation method that is a little more difficult to use than RMS, and is a bit outdated. It equates to the RMS value times two. That is, if the RMS value is -20dB, then the equivalent peak value will be -40dB. This occurs because the RMS value calculated was mapped to a peak sample value for output. This method is basically here for backward compatibility.

RMS

This is a new graph interpretation method that more closely matches the way people hear volume. This mode causes the output to be exactly the RMS amplitude that is specified in the graph. For example, a limiter (flat horizontal line) at -10dB causes the RMS amplitude of the result to average -10dB (where 0dB is a maximum amplitude sine wave without clipping).

Low Cutoff

This is the lowest frequency that dynamics processing will affect. You can define a band, or range, to which compression or expansion is applied, within the current frequency range. The Low Cutoff point is the bottom boundary in defining this band. For example, with values of 1000 for Low Cutoff and 5000 for High Cutoff, dynamics processing only affects audio in the frequency range of 1kHz to 5kHz. To use the entire frequency range of the source material, leave this setting at 0.

High Cutoff

This is the highest frequency that dynamics processing will affect. The High Cutoff point is the top boundary in defining this band. For example, with values of 1000 for Low Cutoff and 5000 for High Cutoff, dynamics processing only affects audio in the frequency range of 1kHz to 5kHz. To use the entire frequency range of the source material, leave this setting at 0. To use the entire frequency range of the source material, this setting should be at 1/2 the current sample rate (24000 for 48kHz, 11025 for 22kHz, etc.).

Leaving this setting at the default of 24000 will work in affecting the entire frequency range for all sample rates below 48kHz.

Lookahead Time

Lookahead Time is used to handle sharp spikes that may occur at the onset of a louder signal. At times, and for brief instances, these onsets can go beyond the limits of your compressor settings, which may be desirable in certain compression scenarios since it can enhance the impact of, say, a drum hit. However, this is obviously not desirable if you are using limiting in order to reduce the maximum amplitude of the audio.

The spikes occur because it takes a little time to determine (the Level Detector's attack value) and react (the Gain Processor's attack value) to the current signal level, so Lookahead Time will actually cause the attacks to start before the audio gets loud, instead of right on top of the transient. Otherwise, with a Lookahead Time of 0, a spike will stay loud until all of the attack times have elapsed.

Graph

The graph depicts input level along the x-axis (left and right) and the new output level along the y-axis (up and down). A line that flows directly from the lower-left to the upper-right (default) depicts a signal that has been left untouched, since every input value goes to the exact matching output value. Adjusting the shape of this line will adjust the input or output assignments, thereby altering the dynamic range. For example, you can boost all input that has a level of around -20dB, leaving everything else unchanged. You can also draw an inverse line (a line from upper-left to lower-right) that will dramatically boost low amplitudes while dramatically suppressing high amplitudes (that is, all quiet sounds will be loud, and all loud sounds will be quiet.)

To add a point to the graph, click in the grid at the location where you would like the point placed. To move a point on the graph, click and hold on the point and drag it to a new location. When the mouse cursor is located over a point, you will see it change from an arrow to a hand. To numerically enter input and output signal levels for an Edit Point, right click on the point to bring up the edit box

Flat

The Flat button resets the graph to its default state of an unchanged signal, removing all control points.

Invert

The invert button changes the graph to one that will function as the exact opposite. For example, if a transfer function with a compressor characteristic is being displayed, pressing Invert will change the graph to one with the corresponding expander characteristic. For a graph to be invertable, it must have points in the two default corners (-100,-100 and 0,0) and its output level must always increase from left to right (i.e. each Edit Point must be higher than the one to its left).

Splines

Check this option to generate a smoother, best fit curve instead of a straight line between control points on the graph. When you use spline curves, the line will not ordinarily travel directly through the control points; rather, the points control the shape of this curve. To get the curve closer to a control point, click to create more control points near the point in question. The more control points there are clustered together, the closer the spline curve will be to those points. Use spline curves when you want very smooth curves instead of straight lines (with their discontinuities at the control points).

Vocoder (Transform Menu, Multitrack View)

A vocoder is a special effect which takes two inputs, usually an instrument of some kind and a voice, and modulates one signal (the process signal, usually the instrument) with the input of the other (the control signal, usually the voice). Specifically, the amplitude of given frequencies in the process signal are made to fluctuate in response to the same amplitude changes in the control signal, allowing for the sound of one instrument to be "controlled" with the other. In the example above, the instrument (the process sound) could be made to "sing" by affecting it with the voice (the control sound).

For ease of use, try checking "Window Width" and set it to about 90%, use 3 or 4 overlays, a resynthesis window of 1 or 2, and FFT sizes from 2048 to 6400.
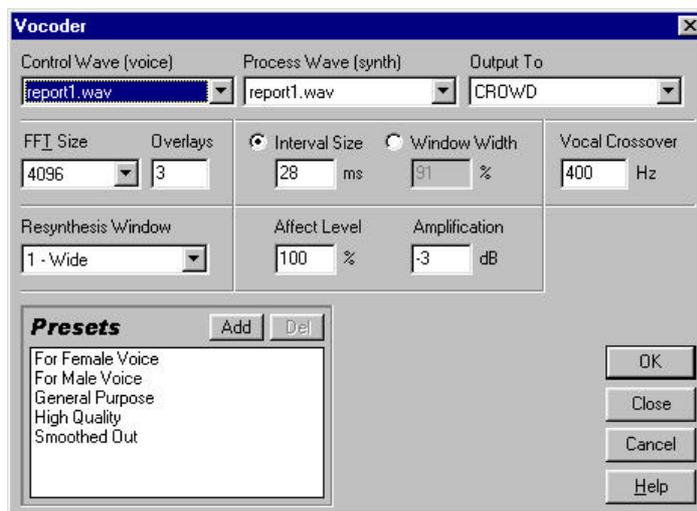
Vocoder

To use *Cool Edit Pro*'s vocoder:

1. Place one wave (a vocal, for example) in track 1, and another wave (perhaps some syth sounds) into track 2.

2. Highlight both waves by clicking the top wave first, and as you highlight drag to the lower right of the second wave.

3. Choose Vocoder from the Transform menu, and set the following parameters:

Control Wave

Choose the waveform to act as the control signal from this drop-down list. This is usually a vocal of some kind.

Process Wave

Choose the waveform to act as the process signal from this drop-down list. This is generally some kind of synthesized sound with which to "replace" the vocal chords.

Output To

Select the track to output the resulting waveform to. This defaults to the first available empty track.

FFT Size

The FFT Size parameter specifies the size of the FFT to use, which can affect processing speed and quality. For cleaner sounding filters, use higher values. A value of 4096 works nicely. Smaller numbers may sound choppier, while larger values will produce extra smooth results, but will take longer to process.

Overlays

This determines the number of FFT's that are overlapped in processing the vocoder. More overlays can produce smoother results (but will take longer to process). Values of 4 to 12 work well.

Interval Size

This sets the time interval per FFT taken. Values between 10ms and 30ms usually work best, but higher overlay settings may require you to vary the interval size used. Smaller values can produce a hum, while larger values can be too blocky sounding.

Window Width

As an alternative to setting the Interval Size, you can simply choose a percentage to use for Window Width in specifying the interval used per FFT taken. A value of 90% will generally produce good results.

Vocal Crossover

This frequency setting is used to filter out, or separate, the source wave's unerlying base frequency (voice) from the vocal formants. With higher values, more formants and the less of the actual vocal tone from the source is used. As this value gets lower, more of the speaker's voice characteristics get vocoded in (not to be confused with the speaker's formant characteristics). Ideally, you want none of the source voice itself carried over, but all of the formant information, so the synthesizer "talks".

Resynthesis Window

Choose the size of the window used for resynthesis of the vocoded signal. Narrower windows will make hard consonants sound clearer, and with higher overlay settings, narrow resynthesis windows will give better time resolution in the event that the vocoded signal sounds too smoothed out. The resynthesis window number should never be larger than or equal to the number of overlays.

Affect Level

This sets the amount of processed (vocoded) signal that ends up in the resulting waveform. Set to 100% for full vocoder, while 50% keeps more of the original wave. 15% will produce a subliminal effect as it will just barely affect the process wave with the voice.

Amplification

You can set an amplification level in dB to adjust the final waveform by. Usually this can be zero, but if the results are too quiet or loud, raise or lower the amplification level to compensate.