

Cool Edit Pro

Version 1.1 User Guide Addendum

Copyright 1998 Syntrillium Software Corporation.

We've been working hard on *Cool Edit Pro* since we released version 1.0! Here's what's new in 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 mice. 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.

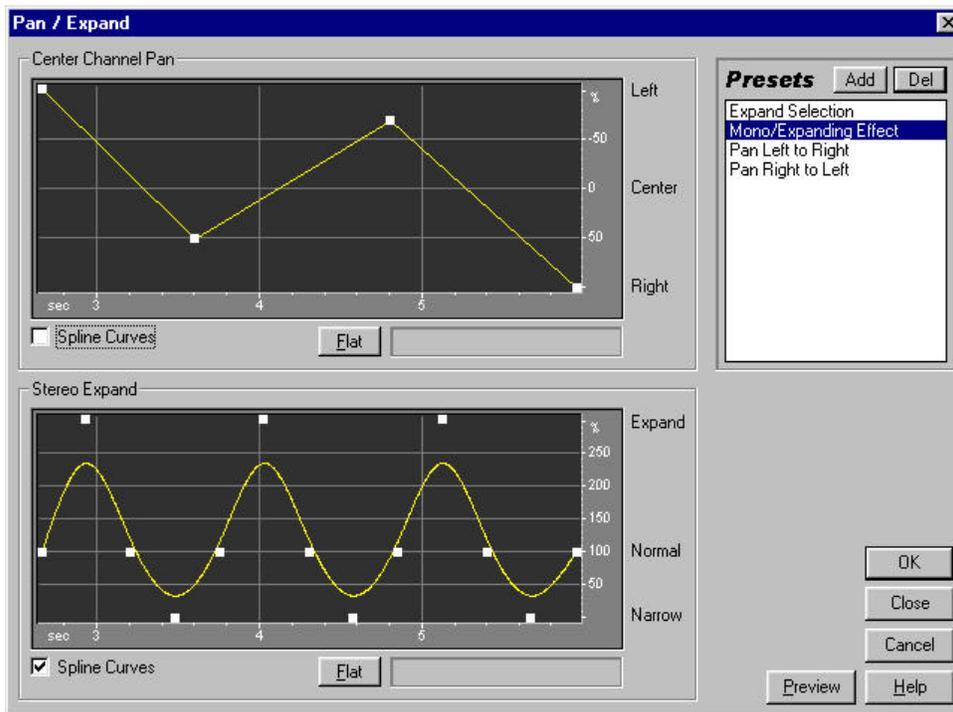
Favorite hot key assignments now take precedence over general hot key assignments, so if you assign the same key to a favorite and the same key is assigned through the Keyboard Shortcuts, the favorite will execute when you hit that key. (Precedence was the other way around in version 1.0.)

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.

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

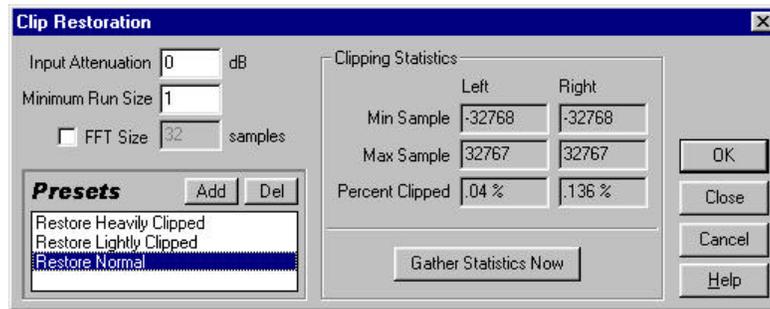
Check this option to generate a smoother, best fit curve instead of a straight line between control points on the graph. When you use spline curves, the line will not ordinarily travel directly through the control points; rather, the points control the shape of this curve. To get the curve closer to 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 current 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 be 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.

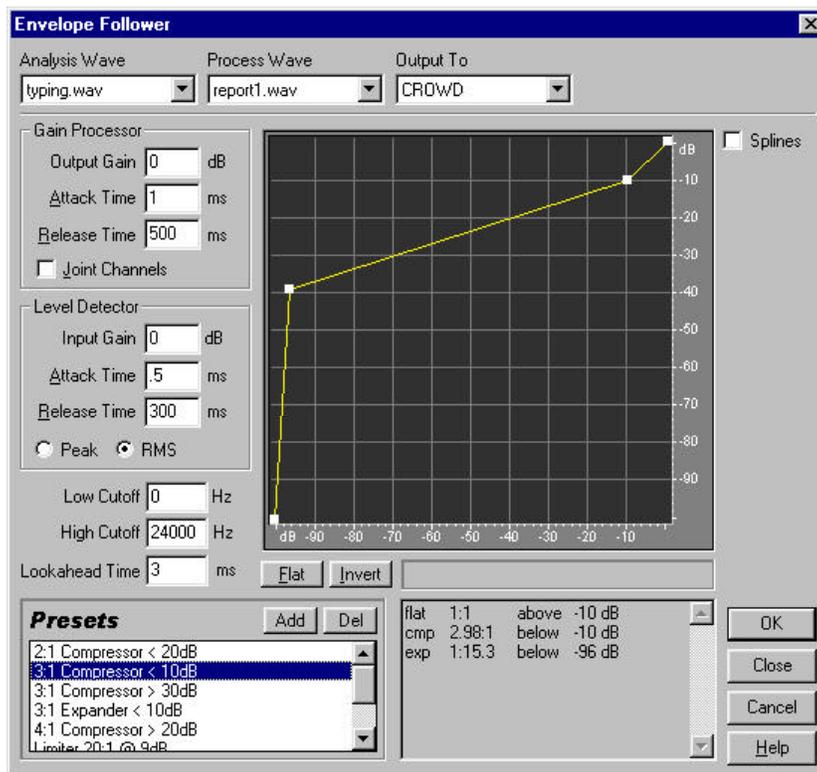
Remove Hiss/Keep Only Hiss

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 invertible, 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).

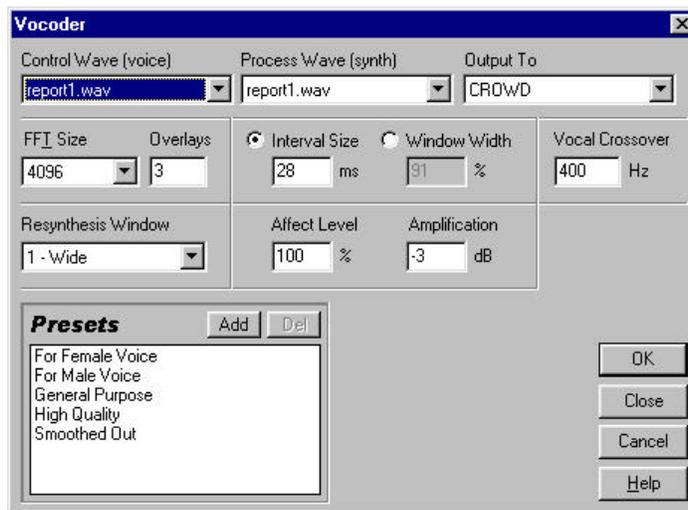
Splines

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 synth 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 choppy, 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 underlying 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.