


Chapter 7

DSP



Chapter 7:

DSP

Introduction

Peak allows you to transform your audio with a variety of powerful Digital Signal Processing (DSP) tools. You can apply these tools at any time by first making a selection in an audio document and then choosing the desired menu command from the DSP menu.

Processing Audio with Peak's DSP tools

The following general procedure describes how to process a selection in an audio document, or the entire document, with a particular DSP function. The specific capabilities and parameters of the DSP function will vary.

To process audio with a DSP function:

1. Select the portion of the audio that you wish to process with the DSP function. If no selection is made, the entire document will be processed.
2. Select the type of process you wish to use from the DSP menu.
3. A dialog appears allowing you to set the parameters for the DSP function.
4. Set the parameters for the DSP function as desired and click OK. Peak processes the selection with the DSP function or plug-in.



Note that if no selection is made, Peak will apply processing to the entire audio document.

Peak's Audio Processing Tools


Peak's DSP capabilities provide composers and professional sound designers with many interesting audio effects and processing tools. Peak's DSP functions include Add, Change Duration, Change Gain, Change Pitch, Convert Sample Rate, Convolve, Crossfade Loop, Invert, Fade In, Fade Out, Find Peak, Gain Envelope, Loop Tuner, Mono to Stereo, Stereo To Mono, Mix, Modulate, Normalize, Panner, Phase Vocoder, Rappify™, Repair Click, Repair Clicks, Remove DC Offset, Reverse Boomerang™, Reverse, and Threshold. The following sections explain how to use each of these functions.

Add

The Add command adds any selection of audio copied to the clipboard into the audio document at the selection point. To use the Add command, you must first copy a selection of audio. The copied material can then be mixed into the target audio material.

To use the Add command:


1. Select the audio that you wish to Add to another audio document and choose Copy from the Edit menu (⌘-C) or Toolbar.
2. Select the audio that you wish to mix the copied material into.
3. Choose Add from the DSP menu or Toolbar.
4. In the dialog that appears, use the slider to adjust the amount of the copied signal that you wish to add into the target audio document. Be careful not to adjust too high an amount which can potentially clip the signal.



The Add dialog

5. Click OK. Peak mixes the two signals together.
6. To hear the results, press the Space bar.

LE Add is not available in Peak LE.



The Amplitude Fit Envelope editor

Amplitude Fit

Amplitude Fit provides granular normalization of an audio selection on a grain-by-grain basis. Grains are small groups of samples, often around 30ms. As each grain is read in, it is normalized according to the Amplitude Fit Envelope—each normalized grain crossfaded with the previous grain and written out as the result. Amplitude Fit can be used to maximize the volume level of an audio selection, or to make quiet passages as loud as louder passages.


To apply the Amplitude Envelope to an audio selection:

1. Select the audio material you wish to process.
2. Choose Amplitude Envelope... from the DSP menu.
3. Draw the amplitude envelope you wish to apply to the audio selection in the envelope editor. Points above and below the 0% line will normalize the selected audio using the grain-by-grain normalization technique.

LE Amplitude Fit is not available in Peak LE.

Change Duration

The Change Duration command allows you to slow down or speed up the selected material by a specified amount *without* changing its pitch. You can specify the change in duration by a value in seconds, a percentage of the original, or, for rhythmically-oriented material, beats per minute. A change in duration by a reasonable amount, about 85% to 115%, can be very convincing. Exaggerated time stretching, 200% or more, can result in some very interesting granular textures. Try experimenting with the Change Duration function on drums, rhythm loops, speech, sampled instruments or sound effects to achieve a wide variety of useful effects.



The Change Duration dialog


To change the duration of a selection:

1. Select the portion of the audio that you wish to process.

2. Choose Change Duration from DSP menu. The Change Duration dialog appears.

3. Click the radio button for one of the following fields, and enter the value that you wish for the change in duration:

- In the Seconds field, type the new duration in seconds.
- In the Percentage field, type the percentage you wish to slow down or speed up the selected audio. For example, typing “50%” will speed up the selection to half its original duration, typing “200%” will slow down the selection to twice its original duration.
- In the Beats per minute field, type the old tempo for the selected audio and then the desired new tempo, and Peak will compute the correct new duration. Use this field to change the duration of rhythmically-oriented material.




The DSP Preferences dialog

4. If you wish to adjust the quality of the pitch change, click on the Prefs button. The DSP Preferences dialog will appear, allowing you to choose the size of the Time Shifting Window that will be used in processing. A lower value is best for simpler, monophonic sounds, while a higher value yields the best results for more complex, polyrhythmic sounds. One you have set these

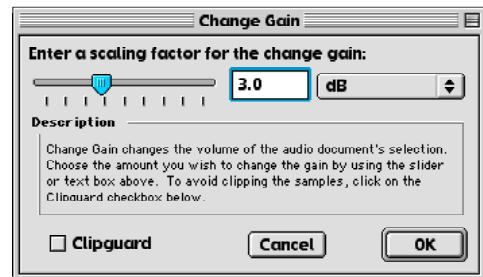
preferences to your liking, click OK to exit the DSP Preferences dialog.

5. Click OK when you have finished. Peak changes the duration of the selection according to the settings that you chose.

 *Change Duration is not available in Peak LE.*

Change Gain

The Change Gain function changes the gain (i.e., amplitude) of a selection. You can specify the amount of gain change either in decibels (dB) or as a percentage. If you wish to double the volume of a sound, you must apply approximately 6dB of gain change, or add 200%. Enable the Clipguard checkbox in the Change Gain dialog to protect against the possibility of clipping. Clipguard will search through the audio document or selection for the maximum peak in amplitude, and then limit the Change Gain slider's range based on the maximum peak it finds in the audio document or selection.



The Change Gain dialog

To change the gain of a selection:

1. Select the portion of the audio that you wish to process.
2. Choose the Change Gain command from the DSP menu or from the Toolbar. The Change Gain dialog appears.
3. Enter the number of decibels or percentage by which you wish to change the amplitude of the selection by.
4. If you wish to protect against the possibility of

clipping, enable Clipguard by checking the Clipguard checkbox.

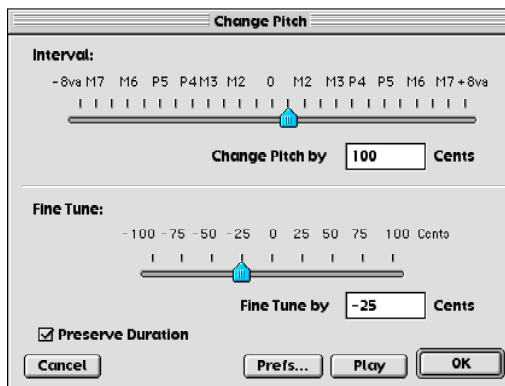
- Click OK when you have finished. Peak will change the gain of the signal by the amount you specified.

Change Pitch

Peak's Change Pitch function allows you to alter the pitch of an audio selection by as much as an octave. The Change Pitch dialog uses a pitch slider that allows you to choose a new pitch by musical interval, and "fine tune" the pitch change by smaller increments called "cents." (Cents are divisions of a musical octave—one octave is equivalent to 1200 cents—thus, 100 cents is a semi-tone, 50 cents a quarter-tone, etc.) You can also choose to alter the length, or duration, of the selection just as you would by slowing down or speeding up analog tape, or you can choose to preserve the duration of the selection (something not possible with analog tape!). You can even preview the pitch change by clicking on the Play button at the bottom of the Change Pitch dialog.

To change the pitch of an audio selection:

- Select the portion of the audio that you wish to process.
- Choose Change Pitch from the DSP menu or from the Toolbar. The Change Pitch dialog appears.
- Select the interval of transposition up or down by entering a positive or negative value in cents in the "Change Pitch" field or by using the pitch slider. Fine tune the interval of transposition by entering a positive or negative value in cents in the "Fine Tune by" field or by using the Fine Tune slider. If you wish to audition the transposition, click on the Play button to preview the results of the pitch change in real time. Check on the Preserve Duration checkbox to retain the original duration of the selected audio.



The Change Pitch dialog

- If you wish to adjust the quality of the pitch change, click on the Prefs button. The DSP Preferences dialog will appear, allowing you to choose the size of the Time Shifting Window that will be used in processing. A lower value is best for simpler, monophonic sounds, while a higher value yields the best results for more complex, polyrhythmic sounds. Once you have set these preferences to your liking, click OK to exit the DSP Preferences dialog.
- Click OK when you have finished. Peak transposes the pitch of the selected audio up or down by the amount you specified.



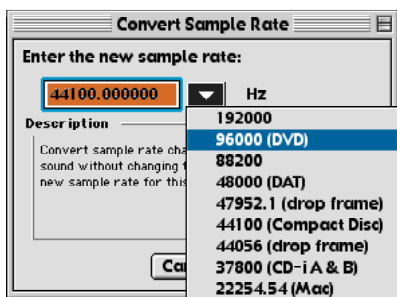
Change Pitch is not available in Peak LE.

Convert Sample Rate

The Convert Sample Rate command allows you to change the sample rate of a sound without changing its pitch. This feature is very useful for converting audio material into lower or higher sample rates as required by other applications. Please note that sample rate conversion is applied to an *entire* document. It cannot be applied to just a selection within a document. Refer to Chapters 3 and 4 for an explanation of commonly used sample rates.

To change the sample rate of a document:

1. Choose Convert Sample Rate from the DSP menu or Toolbar. The Convert Sample Rate dialog appears.
2. Type in the sample rate that you wish to convert the audio document to, or click the down arrow to select from a pop-up of commonly used sample rates.
3. Click OK. Peak converts the entire audio document to the selected sample rate.



The Convert Sample Rate dialog



The Apple Sound Manager is limited to sample rates from 11kHz to 64kHz. Allowable sample rates with ASIO will depend on the particular sound card and its ASIO driver. Digidesign's DAE is restricted to sample rates of 44.1kHz and 48kHz.

Convolve

The Convolve command is a unique and powerful sound design tool that allows you to apply the sonic (e.g., spectral) characteristics of one sound onto another. Convolution works by multiplying the frequency spectrum of the *impulse* contained in the clipboard and that of the target audio document, reinforcing the frequencies that are in common between the two. The results are always interesting and often quite unlike anything you've heard before. This is especially true when the character of the two sounds are very different, and when the clipboard impulse is harmonically rich (imagine, for example, convolving a rainfall sample with piano tinkling!). To

use the Convolve DSP command, you must first copy a selection of audio. The copied material will provide the spectral “character” that you will apply to the target audio material. Convolution can be very useful not only for creating new and unusual sound, but also for giving an audio selection a sense of space—try copying a small amount of room noise to the clipboard and then Convolve it with a selection of audio and the convolved audio will sound like is being played in that room.



Because the clipboard contents that provide the spectrum for this process must be held in RAM, small clipboard impulses should be used, unless a large amount of RAM has been allocated to Peak. This process can use a lot of RAM!

To use the Convolve:

1. Select the audio with the characteristics you wish to apply and choose Copy from the Edit menu (⌘-C) or Toolbar.
2. Select the audio that you wish to modify with the copied audio impulse.
3. Choose Convolve from the DSP menu. Peak applies the spectral character of the copied material to the selection.
4. To hear the results, press the Spacebar.



Convolve is not available in Peak LE.

Crossfade Loop


The Crossfade Loop function applies a “smoothing” effect to loops made in Peak audio documents. Crossfade Loop fades the end of the loop into the beginning of the loop to make the loop sound smoother. (It uses the Blending envelope you've set in Peak's Preference menu's Blending... dialog.) Use the Crossfade Loop dialog to select the length of the crossfade in milliseconds.

To create a crossfade loop:


1. Create a loop using one of the techniques

explained earlier in this chapter.

2. Choose Crossfade Loop from the DSP menu.
3. In the Crossfade Loop dialog that appears, enter a duration for the crossfade-in milliseconds and click OK.
4. To hear the completed crossfade, choose Select Loop from the Edit menu, select Use Loop in Playback from the Preference menu (⌘-L) or click the Loop button on the Toolbar, and press the Spacebar. You will hear the loop, complete with your crossfade.



The Crossfade Loop dialog

 *Crossfade Loop is not available in Peak LE.*

Dither

Dither... allows you to use Waves IDR™ noise shaping and dithering, built-in to Peak. Typically you would pick your destination dithering bit-depth, click OK, then save to that bit-depth. The IDR available in Peak contains a fixed type of dither (type 1), and a fixed amount of noise shaping (normal mode for sample rates over 32Khz, moderate mode for sample rates under 32kHz, no noise shaping below 22Khz).


To dither down the bit-depth of an audio file:

1. Choose Select All under the Edit menu (⌘-A).
2. Choose Dither from the DSP menu.
3. Choose the bit-depth you wish to dither to and click OK. For example, if you have a 24-bit file and you want to dither it down to 16-bit, choose 16-bit.



The Dither dialog

4. Choose Save As... (Shift-⌘-S) from the File menu and choose to Save As the new bit-depth. For example, if you have dithered a 24-bit file down to 16-bit, choose to Save As a 16-bit.

 *Dither is not available in Peak LE.*

Invert

The Invert function allows you to invert the phase of a selection or an entire audio document.

To invert the phase of a selection:

1. Select the portion of the audio that you wish to invert.
2. Choose Invert from the DSP menu. Peak inverts the phase of the selected audio.

Fade In & Fade Out

The Fade In and Fade Out commands allow you to apply an amplitude envelope to an audio selection. The Fade In and Fade Out DSP functions, and the Fade Envelope Editor dialog are described at length in Chapter 5.

To create a Fade In:

1. Click the cursor at the desired location in the audio document and drag to select the audio you want to fade. The Fade In will be applied to the audio within this selection.

2. Choose Fade In Envelope from the Preference menu.
3. In the Fade Envelope Editor dialog that appears, you can use the default envelope, edit the envelope, or load any envelopes included with Peak or that you have created yourself.
4. Choose Fade In from the DSP menu. Peak applies the Fade In to the selection you have made in the audio document.
5. To hear the completed Fade In, press Option-Spacebar. You will hear the selected audio complete with your Fade In.

To create a Fade Out:

1. Click the cursor at the desired location in the audio document and drag to select the audio you want to Fade. The Fade Out will be applied to the selected audio.
2. Choose Fade Out Envelope from the Preference menu.
3. In the dialog that appears, you can use the envelope, create your own, or load any envelopes that you have saved to your hard disk.
4. Choose Fade Out from the DSP menu. Peak applies the Fade Out to the selection you have made in the audio document.
5. To hear the Fade Out, press Option-Spacebar. You will hear the selected audio complete with your Fade Out.

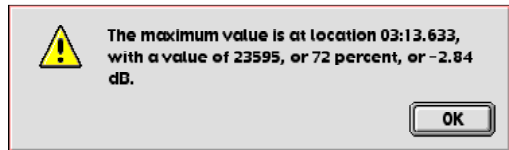
Find Peak


The Find Peak operation will place the insertion point at the sample with the maximum amplitude value that it locates in the audio selection.

To find the maximum amplitude point in an audio selection:

1. Select the audio in which you wish to locate the maximum amplitude.

2. Choose Find Peak from the DSP menu.
3. A dialog will appear telling you what the peak value is, and where it is located. The insertion point will be placed at the sample where the greatest amplitude was located.
4. Press the left arrow to bring the insertion point into view or the Shift key to the view of the insertion point at the sample level




 *Find Peak is not available in Peak LE.*

Gain Envelope

The Gain Envelope operation allows you to enter an amplitude envelope to be applied to an audio selection. The selected audio's amplitude will be boosted and/or attenuated according to the envelope you draw in the Gain Envelope editor. It is easy to cause samples to clip when using this feature, so use it carefully.

To apply variable gain and attenuation to an audio selection:

1. Select the audio material you wish to process.
2. Choose Gain Envelope... from the DSP menu.
3. Draw the gain envelope you wish to apply to the audio selection in the envelope editor. Points above the 0% line will amplify the selected audio. Points below the 0% line will attenuate the selected audio. Note that the waveform display in the Gain Envelope editor will change according to the envelope you draw.
4. To process the audio selection using the gain envelope, press Change.




The Gain Envelope dialog

Loop Tuner

Peak's Loop Tuner provides a way to visually line up the start and end points of your loop and listen to the effects of these adjustments as you make them. If you wish to “tune” a loop you’ve made, simply select Loop Tuner from the DSP menu or Toolbar, and a dialog will appear. The waveform display in the Loop Tuner dialog shows the Start and End points of the loop, which you can visually adjust with the scroll bars at the bottom of the window to achieve a natural transition at the loop point by carefully adjusting the slope alignment. The arrows of the slider will move the loop markers sample by sample and clicking in the body of the slider will move the loop markers to the next zero crossing. The two zoom buttons—magnifying glass icons—in the upper left of the Loop Tuner dialog allow you to adjust the vertical zoom up of the waveform. The two zoom buttons in the lower left hand corner of the Loop Tuner dialog allow you to adjust the zoom view in and out all the way down to the sample level. You can listen to the effects of the adjustments as you make them by clicking on the Play button. To exit this dialog, click on OK to accept the changes, or Cancel to leave the original loop unaffected.



The Loop Tuner dialog with a good, smooth transition



The Loop Tuner dialog with a bad, disjoint transition


LE Loop Tuner is not available in Peak LE.

Mono To Stereo/Stereo To Mono

These two DSP commands may be used to easily convert an audio document between one and two channel formats.

To change an audio document from mono to stereo

1. Select the entire audio document with the Select All command from the Edit menu (⌘-A).
2. Choose Mono To Stereo from the DSP menu or Toolbar.
3. In the dialog that appears, adjust the slider to adjust the left and right-channel balance in the mix.



Mono to Stereo Conversion dialog

4. Click OK. Peak converts the mono document to a stereo document.

To change an audio document from stereo to mono

1. Select the entire audio document with the Select All command from the Edit menu (⌘-A).
2. Choose Stereo To Mono from the DSP menu or Toolbar.
3. In the dialog that appears, adjust the slider to adjust the left- and right-channel balance in the mix.



Stereo to Mono Conversion dialog

4. Click OK. Peak converts the stereo document to a mono document.



Mono To Stereo/Stereo To Mono is not

available in Peak LE.

Mix

The Mix command allows you to mix material that you have copied to the clipboard with a target selection. This function can be used as a kind of “sound-on-sound” capability for mixing audio tracks together, or for blending sound elements. The Mix command is similar to the Add command, but it does not have the potential to clip because the target and clipboard contents are attenuated before mixing. To use the Mix command, you must first copy a selection of audio. The copied material can then be mixed into the target audio material.

To use the Mix command:

1. Select the audio you wish to mix into another audio document and choose Copy from the Edit menu or Toolbar (or press ⌘-C).
2. Select the audio that you wish to mix the copied material into.
3. Choose Mix from the DSP menu.
4. In the dialog that appears, use the slider to adjust the amount of the copied signal that you wish to mix into the target audio document.




The Mixer dialog

5. Click OK. Peak mixes the two signals together.
6. To hear the results, press Option-Space bar.

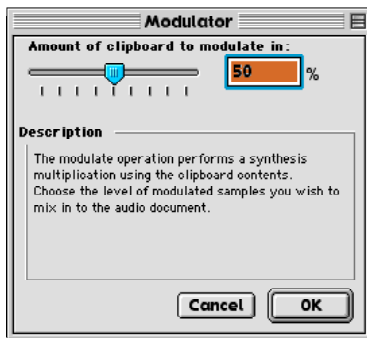
Modulate

This Modulate command functions as a “ring modulator” which multiplies two audio signals together (e.g., the material copied to the clipboard and the currently selected audio). The resulting audio includes the sum and difference tones of the frequency components of the modulated audio and the modulating audio. These are generally very complex timbres that often have a “metallic” (i.e., inharmonic) character to them.

 *Try using generated tones, like sine, swept sine, square, or saw-tooth waves with the Modulate command.*


To use the Modulate command:

1. Select the desired source audio and choose Copy from the Edit menu or Toolbar (or press ⌘-C).
2. Select the destination audio.
3. Choose Modulate from the DSP menu.
4. In the dialog that appears, use the slider to adjust the amount of the copied signal that you wish to use to modulate the destination audio document.




The Modulator dialog

5. Click OK. Peak processes the two signals.
6. To hear the results, press the Space bar.

 *Modulate is not available in Peak LE.*

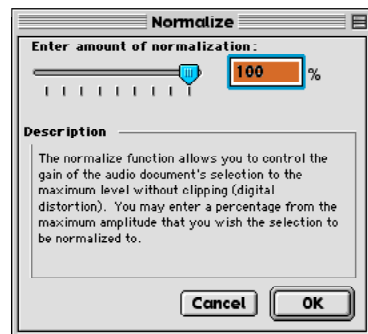
Normalize

This command allows you to optimize the volume of a selection or an entire audio document so that it is at its maximum possible amplitude without clipping. The normalize function is very useful for boosting the volume of material that was recorded at too low a level, or if used on multiple audio documents, for making sure that the amplitude of each of the documents is uniform.

 *Note that because normalization uniformly changes the amplitude of a selection (i.e., the proportions between loud and soft stay the same), it does not have the same effect as compression/limiting (which makes the soft parts louder and does not allow the loud parts to exceed a specified amplitude).*

To normalize a selection:

1. Select the audio that you wish to normalize. If you wish to normalize the entire audio document, choose Select All from the Edit menu (⌘-A).
2. Choose Normalize from the DSP menu.
3. In the dialog that appears, use the slider to adjust the percentage of normalization from the maximum level.



The Normalize dialog


3. Click OK. Peak normalizes the selected audio.

Panner


The Panner allows you to adjust the panning, or left-to-right movement, of a stereo document by drawing an envelope in the Panner dialog. Left is at the top of the graph, and right is at the bottom.

To adjust the panning of a selection:

1. Select the stereo document that you wish to adjust. If you wish to select the entire document, choose Select All from the Edit menu (⌘-A).
2. Choose Panner from the DSP menu.
3. In the Panner editor dialog that appears, use the envelope to “draw in” the panning you desire.
4. Click OK. Peak will change the panning of the document to reflect the changes you’ve made.



The Panner editor dialog

 Panner is not available in Peak LE.


Phase Vocoder

The Phase Vocoder is a type of audio spectrum analysis/resynthesis that allows you to modify the duration and/or pitch of an audio selection.

To use the Phase Vocoder:

1. Select the audio that you wish to process. If you wish to select the entire document, press ⌘-A.
2. Choose Phase Vocoder from the DSP menu. The

Phase Vocoder dialog appears.




The Phase Vocoder dialog

3. In the Change Duration field, you can enter a new duration for the selection by typing the time in seconds.
4. In the Change Pitch field, you can change the pitch of the selection by entering a new value in cents. (Cents are divisions of a musical octave—one octave is equivalent to 1200 cents.) Common musical intervals are stored in the interval pop-up menu, allowing you to enter a major third, octave, or other intervals. Use the direction pop-up menu to control whether the pitch is shifted upward or downward.
5. In the Analysis Settings field, select the number of bands and FFT (Fast Fourier Transform) size to determine the quality of the output. The Phase Vocoder works by analyzing the frequency content of the audio selection and placing the found frequencies into tracks. These tracks are then used to control an oscillator-based resynthesis that uses the pitch and duration modifications you enter. In general, using a smaller FFT size brings less smearing of the audio output than higher FFT sizes. Using a larger number of bands setting used increases the accuracy while tracking of harmonic content of the source sound. In general, setting the FFT size larger than the number of bands will

give undesirable results. Due to the nature of the Phase Vocoder's algorithm, optimum results are achieved when it is used with solo instruments and steady state sounds (such as a voice or solo flute line) rather than complex tones (such as an orchestra playing).

- Click OK. Peak processes the audio. To hear the results, initiate playback.


 *Phase Vocoder is not available in Peak LE.*

Rappify

The Rappify command applies extreme dynamic filtering to a selection. As one Peak user described it, "Rappify can turn your hi-fi into lo-fi!" If the target material has a pronounced beat, this has the effect of reducing the material to its most essential rhythmic components. Try using this function with a variety of different music material for some surprising and exciting results.


To Rappify a selection:

- Select the audio that you wish to process. If you wish to select the entire document, press ⌘A.
- Choose Rappify from the DSP menu.
- In the dialog that appears, select the amount of "rappification" you wish to mix back into the original, with 100% being entirely rappified and 0% being unchanged.



The Rappify dialog

- Click OK. Peak processes the audio. To hear the results, initiate playback.


 *Rappify is not available in Peak LE.*

Repair Click

The Repair Click command will eliminate a selected click or "spike" in the waveform using the setting designated in the Repair Clicks dialog (explained next).

To repair a single click:

- Place the Insertion Point located the click you wish to repair.
- Choose Zoom To Sample Level from the action menu (Shift-Left arrow).
- Select the click in the waveform. Please be sure that your selection is no more than 100 samples.
- Choose Repair Click from the DSP menu.

 *Repair Click is not available in Peak LE.*

Repair Clicks

The Repair Clicks command allows you to find and repair pops or clicks in an audio document. The Repair Clicks dialog automates the process of finding and removing clicks (usually indicated by a sharp "spike" in a waveform), much like a search and replace dialog in a word processor.

The Repair Clicks operation works by looking for discontinuities from sample to sample. For example, a sample value of -100 followed by a sample value of 10,000 is likely to be a click. Once the area of the click is identified, a smoothing technique is used to maintain the original shape of the area being repaired.

If you are working with mostly digitally induced clicks, the Repair Clicks dialog will become an indispensable tool. Extremely damaged signals such as those of a scratching and popping vinyl record will require more careful repair in addition to using the Repair Clicks dialog, such as Change Gain, Delete, and the Pencil

Tool. Clicks such as those of a scratching and popping vinyl record lose their detectability once they are sampled using Analog to Digital converters.



The Repair Clicks dialog

Smoothing Factor

Smoothing Factor determines how much smoothing is applied to the click. Material with high frequency information may require lower smoothing factors to preserve the high frequencies. In general, a setting of 40-60 percent will repair most clicks.

Detection Setting

The Detection Setting value determines how the clicks are located. Higher values locate only the most severe clicks, while lower values will detect less severe clicks. Note that lower values such as 10% also have a greater chance of misjudging audio for a click. In general, a setting of 40-80% works well.

Repair Size


The Repair Size setting affects how many samples around the click are used in determining the new shape of the repair. Repair size can vary from 5 to 100 samples, with a repair size of 50 samples working well in most circumstances. Peak will then interpolate what the correct waveform should be, and repair the click.

Buttons along the bottom of the Repair Clicks dialog allow you to control repairing, auditioning, and undoing click repairs:

- Click the Repair button when you wish to repair a click found by the Next Click button.
- Use the Next Click button to search for the next potential click in the audio selection.
- Once a click is located, you may listen to the click

using the Audition button. The Audition button plays the click using the Pre-roll and Post-roll settings from the Auditioning... dialog under the Preference menu.

- If you repair a click and are unsatisfied with the results, simply click on the Undo button.
- If you would like to repair all of the clicks in the audio document's selection without having to repair each one individually, click the Repair All button.

 *Be sure not to confuse repair size with the size of the selection containing the audio you want to scan and repair. The repair size refers to the size of each individual repaired click.*

To repair multiple clicks in an audio document


1. Select the entire audio document or the area in the audio document you wish to repair click.
2. Choose Repair Clicks from the DSP menu.
3. Click the Next Click button. Peak will search for any clicks. If none are found, you can try again with a lower detection setting.
4. Audition the click using the Audition button. The click should sound in the middle of the auditioned area.
5. Once the click is found, click the Repair button. Click the Audition button to make sure the click was adequately repaired. If it was not adequately repaired, use the Undo button, modify the smoothing factor or repair size and click the Repair button again.
6. Proceed from step 3 until all clicks are removed, or simply click the Repair All button. If you wish to stop the Repair All process, press ⌘-period.

To repair a single click from an audio document:

1. Select the area around the click, centering the click in the selection.

2. Choose Repair Clicks from the DSP menu or use Repair Click and skip step 3.
3. Click the Repair button. Then click the Audition button to make sure the click was adequately repaired. If it was not adequately repaired, use the Undo button, modify the smoothing factor or repair size and click the Repair button again.

You may need to lower the detection setting in the Repair Clicks dialog to find some clicks, depending upon their severity. Be careful not to lower the detection setting dramatically — lower it gradually for the best results.


 *Repair Clicks is not available in Peak LE.*

Remove DC Offset

This function allows you to remove any DC Offset in your audio file. Peak scans the audio for DC offset and then removes it. Peak will scan the left and right channels of a stereo file independently. DC Offset is usually caused by problems in the analog to digital conversion process. The result is that the waveform is not centered on the base line—it is offset either higher or lower than the center line. The Remove DC Offset function is particularly useful for preparing audio for processing with the RealAudio Encoder.

To use Remove DC Offset:

1. Select the audio that you wish to process. If you wish to select the entire document, choose Select All from the Edit menu (⌘-A).
2. Choose Remove DC Offset from the DSP menu. Peak will scan the audio, and automatically remove any DC offset that might be present.

 *Remove DC Offset is not available in Peak LE.*


Reverse Boomerang

The Reverse Boomerang command mixes a reversed copy of the selected audio with the original. This creates a variety of interesting and useful results. Try

using Boomerang on drum loops, voice, and sound effects.

To use Reverse Boomerang:

1. Select the audio that you wish to process. If you wish to select the entire document, choose Select All from the Edit menu (⌘-A).
2. Choose Reverse Boomerang from the DSP menu.
3. In the dialog that appears, select the amount of reversed sound you wish to mix back into the original, with 100% being entirely reversed, and 0% being unchanged.



The Reverse Boomerang dialog

4. Click OK. Peak processes the audio. To hear the results, press the Spacebar to initiate playback.

Reverse

The Reverse command reverses the current selection. In a reversed selection, the last sample becomes the first sample, the second-to-last sample becomes the second sample, and so-forth. The effect is similar to playing a record or cassette tape backwards.

To reverse a selection:

1. Select the audio that you wish to reverse. If you wish to select the entire document, choose Select All from the Edit menu (⌘-A).
2. Choose Reverse from the DSP menu. Peak reverses the selected audio. To hear the results,

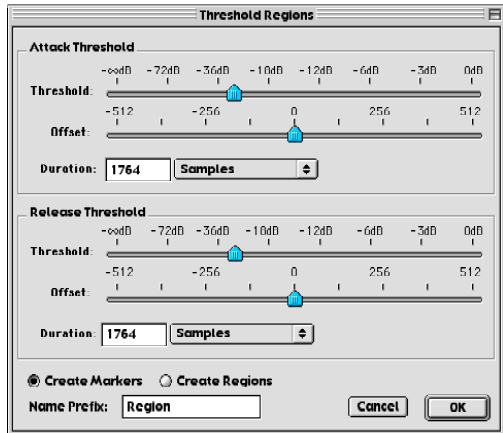
start playback.

Threshold

The Threshold command allows you to split up an audio document into its component parts by analyzing the amplitude levels in the audio document and setting a cutoff or threshold amplitude. For instance, you might use the Threshold command on an audio document that contains successive notes from a musical instrument to split them up, or on a drum loop to break it up into its component parts. You can save the segments with Markers, or as Regions.

To use the Threshold command:

1. Select the audio you wish to process and choose Threshold from the DSP menu. After Peak analyzes the amplitudes in the selection, the Threshold dialog will appear, allowing you to select a threshold amplitude for both attack and release values.



The Threshold Regions dialog

2. Drag the threshold indicator left or right to set the threshold amplitude. As you drag the indicator, new markers will appear in the audio document forming markers or regions, depending on your settings. The Offset sliders allow you to “nudge”

the onsets of markers or regions by plus or minus 0 to 512 samples.

3. Select Create Regions to create regions instead of markers. The separate Release Threshold, attack and sludge settings affect the region end points, allowing you to eliminate silence from the region end points.
4. Adjust the Attack value. This setting sets the amount of time that audio must stay above the given threshold to qualify as a new marker or region.
5. When you have finished, click OK.
6. After the audio document has been “thresholded” to your satisfaction, you can use the Export Regions command in the File menu to export the separated regions into new windows or files.
7. To select and play regions in order from left to right, press the Page Up key on your computer keyboard. To select and play regions in order from right to left, press the Page Down key.



User Tip: Use the Threshold command to create several looping points. To convert a marker to a Loop Start or Loop End point, double-click on the marker and change it to “Loop Start” or “Loop End” in the Edit Marker dialog. Also, try rearranging the regions generated by the Threshold function in the Playlist or by using Cut and Paste to create new interesting compositional and rhythmic ideas!



Threshold is not available in Peak LE.

Conclusion

You have now learned how to manipulate and process audio using Peak's native DSP capabilities. In the next chapter, you will learn how to use third-party Premiere, AudioSuite, TDM, and VST plug-ins with Peak.
