Sonic NoNOISE® Guide

Version 9.0

Legal Notices

Guide Part Number 9329-65057-00 REV A 9/10

Documentation Feedback

At Avid, we are always looking for ways to improve our documentation. If you have comments, corrections, or suggestions regarding our documentation, email us at ***techpubs@avid.com***.

contents

chapter 1

Introduction

Welcome to the Sonic NoNOISE Plug-In Suite for Pro Tools®.

NoNOISE is a set of TDM and AudioSuite™ plug-ins for Pro Tools systems. Each NoNOISE plug-in provides a specific type of noise reduction and audio restoration processing (with many options available in each processor).

NoNOISE processing, originally developed by Sonic Solutions, is used in the mixing and mastering stages of audio production. The NoNOISE set of plug-ins isolates and eliminates audio artifacts such as hiss, scratches, hum, and mechanical and impulsive noise.

NoNOISE can be used to restore old recordings, remove unwanted noises from field recordings, and repair audio materials that have suffered damage.

The AudioSuite version of NoNOISE supports sample rates up to 192 kHz, and the TDM multi-mono version of NoNOISE supports sample rates of up to 96 kHz.

AudioSuite

Manual Declicking Assists in removing unwanted noises such as clicks, pops and low-frequency artifacts. The algorithms analyze audio on either side of the anomaly and, based on this information, synthesize replacement sound for the original.

Production Declicking Lets you analyze and declick sound files that contain multiple clicks.

Decrackling Corrects dense, impulsive noise in which small impulses crowd against one another (crackling) by performing a type of sliding interpolation that isolates good audio and uses it as the basis for resynthesis.

TDM

Broadband Denoising Analyzes and identifies noise content, then lets you create a filter to remove the noise from audio content. Analysis is performed using the Broadband DeNoise AudioSuite plug-in, and filtering is applied using the Broadband DeNoise TDM plug-in.

The Broadband Analysis AudioSuite plug-in provides analysis functions for the Broadband Denoising TDM process, and is not supported by non-TDM Pro Tools systems.

High-Resolution (High-Res) Filters Useful for removing artifacts such as hums and buzzes, and for applying filter types commonly used in the mastering stages of audio production.

TDM and AudioSuite Support

Due to their types of processing, some NoNOISE plug-ins are TDM while others are AudioSuite.

NoNOISE TDM plug-ins are multi-mono only.

NoNOISE AudioSuite plug-ins do not support multi-mono processing.

NoNOISE plug-in formats

Process	Plug-In Format	Pro Tools Software
High-Res Filters	TDM, Multi-Shell II	Pro Tools HD
Broadband Denoising Processor	TDM	Pro Tools HD
Broadband DeNoise Analysis	AudioSuite	Pro Tools HD
Manual Declick	AudioSuite	Pro Tools HD Pro Tools
Production Declick	AudioSuite	Pro Tools HD Pro Tools
Decrackle	AudioSuite	Pro Tools HD Pro Tools

Sample Rate Support

TDM The TDM Broadband DeNoiser and all of the High-Res Filters support sample rates up to 96 kHz.

AudioSuite The Manual Declicker, Production Declicker, and Decrackle AudioSuite plug-ins support sample rates up to 192 kHz. The Broadband DeNoising Analysis AudioSuite plug-in supports sample rates up to 96 kHz.

Contents of the Boxed Version of Your Plug-In

Your Sonic NoNOISE plug-in package contains the following components:

- Installation disc
- Activation Card with an Activation Code

System Requirements and Compatibility

To use NoNoise TDM plug-ins, you need:

- An iLok USB Smart Key
- An iLok.com account for managing iLok licenses
- A qualified Pro Tools®|HD system.

To use NoNoise AudioSuite plug-ins, you need:

- An iLok USB Smart Key
- An iLok.com account for managing iLok licenses
- One of the following:
  - A qualified Pro Tools system

  – or –

  - A qualified Avid® Xpress®, Avid Xpress DV or Avid DNA™ system

Avid can only assure compatibility and provide support for hardware and software it has tested and approved.

For complete system requirements and a list of qualified computers, operating systems, hard drives, and third-party devices, visit:

www.avid.com/compatibility

Registering Plug-Ins

Your plug-in purchase is automatically registered when you activate your iLok license (see "Authorizing Plug-Ins" on page 7).

Registered users are eligible to receive software update and upgrade notices.

For information on technical support, visit www.avid.com.

Working with Plug-Ins

Besides the information provided in this guide, refer to the *Pro Tools Reference Guide* for general information on working with plug-ins, including:

- Inserting plug-ins on tracks
- Using clip indicators
- Navigating the Plug-In window
- Adjusting plug-in controls
- Automating plug-ins
- Using plug-in presets

Important Differences in NoNOISE for Pro Tools and NoNOISE for Sonic Solutions

If you are experienced with NoNOISE on Sonic Solutions SonicStudio™, you will notice that NoNOISE for Pro Tools employs most, but not all features found in the Sonic Solutions product. In addition, some features are implemented differently due to differences between Pro Tools and SonicStudio operation.

Features *not* available in the Pro Tools version include:

- Noise Estimate Files (use Pro Tools plug-in Settings files instead)
- Complex Filtering
- AudioClient
- AutoSonic Server
- Audio Workers
- QueueManager
- Sonic Solutions-Proprietary Sample Rate Conversion Algorithms

Conventions Used in This Guide

All Pro Tools guides use the following conventions to indicate menu choices and key commands:

Convention	Action
File > Save	Choose Save from the File menu
Control+N	Hold down the Control key and press the N key
Control-click	Hold down the Control key and click the mouse button
Right-click	Click with the right mouse button

The names of Commands, Options, and Settings that appear on-screen are in a different font.

The following symbols are used to highlight important information:

User Tips are helpful hints for getting the most from your Pro Tools system.

Important Notices include information that could affect your Pro Tools session data or the performance of your Pro Tools system.

Shortcuts show you useful keyboard or mouse shortcuts.

Cross References point to related sections in this guide and other Pro Tools and VENUE guides.

About www.avid.com

The Avid website (www.avid.com) is your best online source for information to help you get the most out of your Pro Tools system. The following are just a few of the services and features available.

Product Registration Register your purchase online.

Support and Downloads Contact Avid Customer Success (technical support); download software updates and the latest online manuals; browse the Compatibility documents for system requirements; search the online Knowledge Base or join the worldwide Pro Tools community on the User Conference.

Training and Education Study on your own using courses available online or find out how you can learn in a classroom setting at a certified Pro Tools training center.

Products and Developers Learn about Avid products; download demo software or learn about our Development Partners and their plug-ins, applications, and hardware.

News and Events Get the latest news from Avid or sign up for a Pro Tools demo.

chapter 2

Installation and Authorization

Installing Plug-Ins for Pro Tools

Installers for your plug-ins can be downloaded from the Avid Store (http://shop.avid.com) or can be found on the plug-in installer disc (included with boxed versions of plug-ins).

An installer may also be available on the Pro Tools installer disc or on a software bundle installer disc.

To install a plug-in:

1 Do one of the following:

- Download the installer for your computer platform from www.avid.com. After downloading, make sure the installer is uncompressed (.dmg on Mac or .ZIP on Windows).

– or –

- Insert the Installer disc into your computer.

2 Double-click the plug-in installer application.

3 Follow the on-screen instructions to complete the installation.

4 When installation is complete, click Finish (Windows) or Quit (Mac).

When you launch Pro Tools, you are prompted to authorize your new plug-in.

Authorizing Plug-Ins

Software is authorized using the iLok USB Smart Key (iLok), manufactured by PACE Anti-Piracy.

iLok USB Smart Key

An iLok can hold hundreds of licenses for all of your iLok-enabled software. Once a license for a given piece of software is placed on an iLok, you can use the iLok to authorize that software on any computer.

An iLok USB Smart Key is not supplied with plug-ins or software options. You can use the iLok included with certain Pro Tools systems (such as Pro Tools|HD-series systems), or purchase one separately.

Authorizing Downloaded Software

If you downloaded software from the Avid Store (http://shop.avid.com), you authorize it by downloading a license from iLok.com to an iLok.

For more information, visit the iLok website (www.iLok.com).

Authorizing Boxed Versions of Software

If you purchased a boxed version of software, it comes with an Activation Code (on the included Activation Card).

To authorize software using an Activation Code:

1 If you do not have an iLok.com account, visit www.iLok.com and sign up for an account.

2 Transfer the license for your software to your iLok.com account by doing the following:

- Visit www.avid.com/activation.

– and –

- Input your Activation Code (listed on your Activation Card) and your iLok.com User ID. Your iLok.com User ID is the name you create for your iLok.com account.

3 Transfer the licenses from your iLok.com account to your iLok USB Smart Key by doing the following:

- Insert the iLok into an available USB port on your computer.
- Go to www.iLok.com and log in.
- Follow the on-screen instructions for transferring your licences to your iLok.

For more information, visit the iLok website (www.iLok.com).

4 Launch Pro Tools.

5 If you have any unauthorized software installed, you are prompted to authorize it. Follow the on-screen instructions to complete the authorization process.

Removing Plug-Ins

If you need to remove a plug-in from your Pro Tools system, follow the instructions below for your computer platform.

Mac OS X

To remove a plug-in:

1 Locate and open the Plug-Ins folder on your Startup drive (Library/Application Support /Digidesign/Plug-Ins).

2 Do one of the following:

- Drag the plug-in to the Trash and empty the Trash.

– or –

- Drag the plug-in to the Plug-Ins (Unused) folder.

Windows

To remove a plug-in:

1 Choose Start > Control Panel.

2 Under Programs, click Uninstall a program.

3 Select the plug-in from the list of installed applications.

4 Click Uninstall.

5 Follow the on-screen instructions to remove the plug-in.

chapter 3

Using NoNOISE AudioSuite Plug-Ins

This chapter explains how to use three of the NoNOISE AudioSuite plug-ins. AudioSuite plug-ins are applied as a file-based process to selected audio.

Manual Declicking See "Manual Declicking" on page 9.

Production Declicking See "Production Declicking" on page 12.

Decrackling See "Decrackle" on page 19.

The Broadband Analysis plug-in is an AudioSuite process. This AudioSuite plug-in provides analysis for the TDM Broadband Denoising plug-in. See "Broadband Denoising" on page 23 for more information.

NoNOISE AudioSuite plug-ins do not support multi-mono processing.

Combining Processes For Best Results

Throughout this chapter, tips are provided that suggest combinations of NoNOISE processes for maximum results. For example, applying a preliminary Manual Declicking pass can improve the results yielded with the Decrackle plug-in (see "Decrackle" on page 19 for more information).

Manual Declicking

Manual Declicking is designed to remove or reduce unwanted noises such as individual or isolated clicks, pops and low-frequency artifacts. It offers five different interpolation algorithms that are capable of correcting even very difficult audio anomalies. (For details on each available algorithm, see "Interpolation Algorithms" on page 10.)

The algorithms analyze audio on either side of the anomaly and, based on this information, synthesize replacement sound. NoNOISE substitutes the replacement sound for the original sound. If you don't like the result of the interpolation, you can Undo the operation, and try again with adjusted settings.

To use Manual Declicking:

1 Locate and zoom in on a click in an audio track.

2 Use the Selector to make a selection that fully contains the click.

3 Choose AudioSuite > NoNOISE Manual Declicking.

4 In the Manual Declicking window, choose a process type from the pop-up menu (See “Interpolation Algorithms” on page 10 for details on these options.)

5 Configure AudioSuite settings and modes as needed. See the *Pro Tools Reference Guide* for details on these standard AudioSuite controls.

6 Click Process. The selected audio is replaced with continuous, interpolated audio based on the selected process type.

To undo:

- Choose Edit > Undo.

Manual Declicking operates on a single channel, so when working with multichannel material, you must declick each Track separately.

Interpolation Algorithms

There are several interpolation algorithms or *interpolators* available. Each is suited to a particular type of audio problem and context.

The B type interpolator is the most general. The majority of declicking situations can be handled by simply choosing this selection.

Type A and Type C Pitched Interpolators

The Type A and Type C interpolators are pitched waveform interpolators. A waveform interpolator is most useful in dealing with periodic waveforms, such as brass instruments or the human voice. The Type A and C interpolators take context information from multiple periods (a period is the distance between successive peaks in the waveform) to the left and right of the area identified by the selection.

The difference between these two types is that the Type C has protections built into it for certain cases for which waveform interpolation algorithm produces bad results. Type A lacks these protections. The Type A interpolator will often produce an interpolation, but the results may not always be pleasing. The Type C interpolator will sometimes fail to interpolate the signal designated. (If this occurs, try a different interpolation algorithm.)

Type B General Interpolators

This is the most general-purpose of the declicking algorithms, and works well on complex musical waveforms (for example, where several signals are combined with instruments that produce non-periodic waveforms, such as sax or strings).

The basic Type B interpolator examines the audio on either side of the click to determine the context for resynthesizing audio to fill the gap. In most situations this basic interpolator will produce the best results.

Type B L –>R and B R–>L General Interpolators

There are two variations of the command that load the context information in a particular way. If, for example, a click occurs just prior to the beginning of the attack of a piano note, the basic Type B interpolator would include part of the piano note in its resynthesis, producing the impression that the piano starts a bit early.

In this case, the B Type, L–>R (left-to-right) variation would avoid getting the piano note into the interpolation. Likewise, the R–>L (right-to-left) variation might be used in an instance where a click follows immediately after a sudden change in the audio waveform.

Type D Declicker and Type E Decrackler Interpolation

These are very high-order interpolations that may be used to correct problems that elude other interpolation algorithms. Both interpolators use 64-bit precision arithmetic to produce very high quality interpolation.

The Type D interpolation is only capable of replacing up to about 2 milliseconds (0.002 seconds) of sound before it runs out of memory. The Type E Interpolator provides a very similar algorithm that can be used on large sections of audio.

If there is a particularly problematic area, then Type E manual interpolation can help clean up small regions. As with the other interpolation algorithms, Type E operates directly on the sound file.

Production Declicking

Production Declicking lets you analyze and de-click audio files, regions, and audio selections. (To remove individual or isolated clicks, see “Manual Declicking” on page 9.)

Production Declicking is an AudioSuite plug-in, and provides many tools for detecting and minimizing clicks, as explained in the following sections.

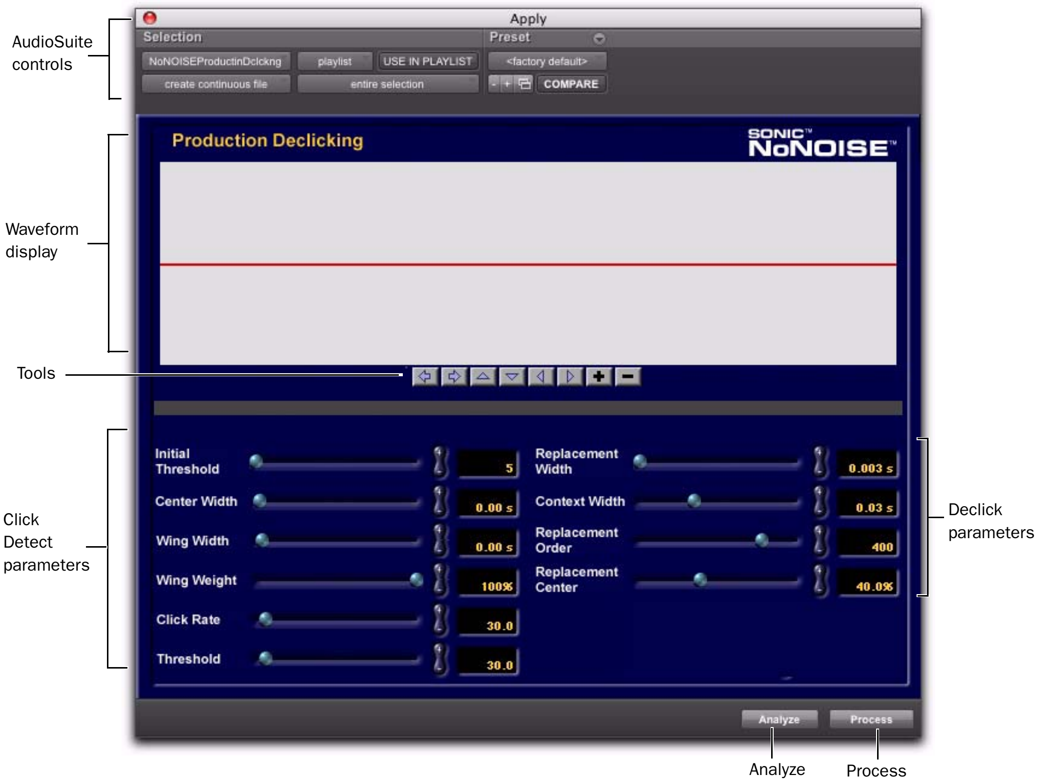

Figure 1. Production Declicking

To declick a soundfile using Production Declicking:

1 Select an audio file, region, or range of audio in a track, to declick.

2 Choose AudioSuite > NoNOISE Production Declicking. The Production Declicking window opens, with an empty Waveform display.

3 Click Analyze. NoNOISE analyzes the audio for the location and a description of each click.

4 The results of the analysis pass are shown in the Waveform display.

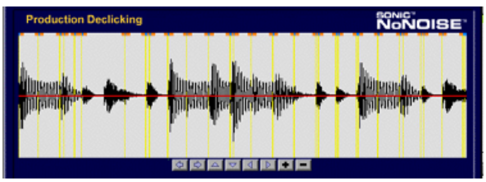

Waveform display after analysis

Yellow vertical lines show the location of each detected click, based on the current Click Detect Parameters. (See “Click Detect Parameters” on page 14 for details.)

The Waveform display also shows the following Declicking parameters:

◆ Orange horizontal bars reflect the Context Width.

◆ Blue horizontal bars reflect the current Replacement Width.

◆ Replacement Center position can be adjusted using the slider or increment arrows.

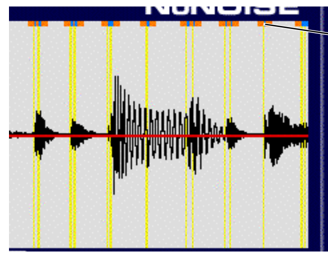

Adjusting Declicking parameters in the waveform display

5 If the current settings did not detect clicks to your satisfaction:

◆ Adjust the Click Rate, Threshold and other Click Detection parameters using their sliders, then click Analyze to perform a new analysis pass using the new settings. (For details, see “Click Detect Parameters” on page 14.)

6 You can fine tune the declicking process using the Replacement Width, Context Width, Replacement Order, and Replacement Center sliders. (For details on these settings, see “Declicking Parameters” on page 16.)

7 Click Process.

The Declick pass removes the clicks from the audio, based on the current Declick settings.

Click Detect Parameters

The following sections provide specifications and examples of each of the Click Detect parameters.

Click Detection Parameters

Parameter	Range	Default Value	Recommended	Extreme
Initial Threshold	1 to 5000	5	5 to 10	>100
Center Width (ms)	.06 to 46.44	1.00	0.5 to 2.5	>5.00
Wing Width (ms)	.101 to 44.64	2.00	2.00 to 10.00	>10.0
Wing Weight (%)	0 to 100	100	100	0

Initial Threshold

When the detection algorithm identifies a candidate click site, it measures and assigns a value for total energy. This is compared against the Initial Threshold, and against the energy in the Wings (see "Wing Width" on page 15) to determine if the site is an actual click.

Initial Threshold is the lowest value that is recognized as a click. Click Detect accepts values from 1 to 5000 for Initial Threshold, but it takes a very large click to produce a total energy value greater than 200. Higher values have the effect of excluding virtually all clicks.

Initial Threshold has the effect of limiting the number of clicks detected. This helps to avoid detection of spurious click sites and limits the size of the click list so that it remains manageable within the system's memory.

For 78 RPM phonograph recordings, an Initial Threshold value between 5 and 10 is recommended. To detect only loud clicks, use a range of 100–200.

To find all clicks, use the minimum setting of 1, but be aware that it is possible for the click list to become very large and processing will be slow.

To avoid overly large click lists, you can run multiple detect/declick passes with progressively lower thresholds, allowing even the longest pieces to be declicked. This also helps to maintain your control, by allowing you to evaluate and modify the results at each stage.

Initial Threshold is used to fine tune the analysis and click detection process before you analyze a passage. The two controls for Click Rate and Threshold are used after an analysis pass, to edit the results before applying declicking. See "Click Rate and Threshold" on page 18 for details.

Center Width

Center Width is the length of the frame that the detector analyzes. It represents the duration of a typical click. If a click is much shorter than the Center Width, the click detect bar may not be properly centered over the click.

If a click is longer than the Center Width, the system will most likely still detect it, but the site will be listed shorter than the actual click. You may replace the entire length of the clicks during the click removal process by setting a wide Replacement Width setting.

Typical clicks range from about 0.5 to 2.5 milliseconds in length. We recommend keeping the Center Width within this range. The initial default value is 1 millisecond. For special cases, the value might be set as high as 5 milliseconds.

If a recording contains several distinct kinds of clicks, for example if a broken record contains loud and long clicks near a break, and also has intermittent normal, short clicks, it is best to process the recording in two passes. Perform the first pass using a large value of Center Width; after declicking and removing those clicks, perform another pass using a shorter value of Center Width.

Wing Width

Wings are the sections that precede and follow a candidate click site. They provide the context information for recognizing a valid click. The click detector calculates total spectral energy in the wings and in the site candidate, then subtracts energy in the wings from that in the site candidate. Click amplitude is the difference between the center and the wings.

Wing Width is the time in seconds on each side of the site candidate that is used for analysis. The minimum value is 0.0001 seconds (0.1 milliseconds).

The maximum value of Wing Width is 2048 samples, corresponding to approximately 42.7 milliseconds at 48kHz sample rate, or 46.4 ms at 44.1kHz sample rate.

The default value of Wing Width is 0.002 seconds (2 milliseconds).

When using the recommended click detect algorithm, this type of narrow Wing Width is preferred.

However, if the Wing Width is too small, it may interact with signal, causing normal transients to be detected and eliminated as clicks. Normal music waveforms (such as male voice, trumpet, and trombone sounds) exhibit impulsive behavior that can be mistaken for clicks.

If the Wing Width is much less than one waveform period, the detector will sometimes list a spurious click at the beginning of each period. If the source material includes strong low frequencies, setting Wing Width to 10, 15, or even 20 milliseconds will eliminate this problem, but at the risk of not detecting some clicks separated by less time.

Wing Weight

The NoNOISE click detector subtracts total spectral energy in the wings surrounding a possible click site from the energy within the site area as defined from the Center Width. This difference value is used to determine valid clicks to be eliminated.

Wing Weight is a value between 0 and 100 percent, applied to the wings energy value before it is subtracted from that of the candidate click site. In essence, it tells the system how much to value the wings in determining the validity of a click.

The amplitude of a detected click is lowered when the wing amplitude is subtracted from the click amplitude. Lowering the Wing Weight can be thought of as decreasing the contrast between the site candidate and its wings.

This is useful in determining the validity of a click, but can sometimes result in valid clicks being rejected. Lowering Wing Weight compensates for this, but increases the risk of falsely detecting sites that are not real clicks.

Optimum settings for Wing Weight varies with program material. The initial default of 100% (Wing Weight has full effect) is suitable for the greatest range of program material.

If the program material is more band-limited than the clicks, as with clicks on 78 RPM discs, Wing Weight may be set to a lower value.

At a setting of 0, the criterion for click site identification is based strictly on the magnitude of the filtered signal. This is useful when the audio material is severely band-limited and the click's principal energy is out of the program's frequency range.

Declicking Parameters

The following sections provide specifications and examples for each declicking parameter.

Ranges and Defaults of Declick Parameters

Declick Parameters

Parameter	Range	Default Value	Recommended
Replacement Width (ms)	.091 to 2.88	2.500	2.500 to 3.00
Context Width (ms)	1.0 to 81.27	30	30
Replacement Order	3 to 1024	400	400
Replacement Center (ms)	0 to 100	40	40
Click Rate/Threshold	0 to 500	30	30

Replacement Width

The declick algorithm is capable of filling a gap ranging from 4 to 4096 audio samples. Replacement Width determines the number of samples that will be replaced when removing the click.

Replacement Width is expressed in seconds, from 0.00009 to 0.092 seconds (0.09 to 92.0 milliseconds) at 48 kHz sample rate. At 44.1 kHz, the maximum value is slightly higher. The default value is 0.0025 seconds, or 2.5 milliseconds.

The default value of 2.5 milliseconds is reasonable for normal disk recordings. You should always try to replace the minimum amount of sound necessary to eliminate clicks. It's generally best to replace just a little bit more than the actual length of the click.

A good rule is to examine the source material to determine the length of the average click, then add approximately 20 percent. The typical click from a 78 RPM phonograph record is 0.5 to 2.0 milliseconds, making the setting of 0.0025 to 0.003 a good default.

The length of sound replaced is not related to Center Width in the click detect pass, nor to the width of the bar that appears over the click site.

If the Replacement Width is set too low, then partial clicks may be left in the audio. If Replacement Width is set too high, the chances of generating low frequency artifacts increase.

In the Waveform display, NoNOISE shows two types of bars above the audio. Blue bars shows the audio that will be replaced. Orange bars show the context that will be used for the interpolation.

Context Width

Context Width is the length of audio to either side of the click that is used in resynthesis. Its value is expressed in seconds.

Larger values of Context Width require longer time to process, but are less likely to produce artifacts. The default value of 0.030000 seconds (30 milliseconds) provides a reasonable compromise between speed and accuracy.

The minimum and maximum Context Width depends on the settings of Replacement Order and sample rate:

◆ Minimum Context Width = ([Replacement Order+1], sample rate (44100 or 48000))

◆ Maximum Context Width = f([4096 – 256 – Replacement Order], sample rate (44100 or 48000))

For best results, set the Context Width to at least twice the minimum value. Using the maximum Replacement Order of 512 at 44.1 kHz sample rate, the minimum context is about 0.012 seconds (12 milliseconds), and the Context Width should be set to at least 24 milliseconds. Setting Context Width too low will produce artifacts sounding like anything from low frequency to bursts of noise, depending on the context. However, setting the Context Width too high substantially increases processing time.

Replacement Order

Replacement Order sets the precision of the resynthesis calculation, and concerns the power of interpolation. The default value for Replacement Order is 400, with a minimum value of 3 and a maximum of 512. In general, the larger the Replacement Order the better the interpolation but the slower the processing. If either Replacement Order or Context Width are increased, the processing time increases in proportion to the sum of the increase in these two numbers. For example, if Replacement Order is increased by 10 percent and Context Width by 10 percent, then total processing time increases by about 20 percent.

The default of 400 is a good value to use in most situations. If you are experiencing low-frequency artifacts, then the value of Replacement Order should be increased. The trade-off is that it will take longer to replace each click.

The default Context Width and Replacement Order have been found suitable for about 80 percent of Declicking projects. For extremely large clicks, Context Width may be raised to about 35 milliseconds (0.035 seconds) and Replacement Order to 512.

Replacement Width and Order

To use a large value Replacement Width and avoid artifacts, increase Replacement Order. Keep in mind that using a maximum Replacement Order of 512 limits Context Width at 44.1 kHz sample rate to:

[4096 – 256 – 512] / 44100 = 75.5 milliseconds

Replacement Center

Replacement Center specifies an offset (as a percentage) that shifts the area of interpolation in relation to the marked click site.

Clicks are often followed by some amount of ringing, making it necessary to continue replacement for some time after the click itself. In transcribing a phonograph record, for example, the stylus and tonearm resonate in response to the impulse of a scratch on the record, producing a damped oscillation that extends for some time after the click. This ringing is not detected by the click detect functions.

Replacement Center shifts the replacement area to the left or right in relation to the click center according to a percentage from the left edge. A setting of 0 percent positions the click at the left edge of the replacement area, while a setting of 50 percent places the click directly in the middle. A setting of 100 percent positions the click at the very right edge of the replacement area.

If Replacement Center is set so that the ringing of the click extends beyond the Replacement Width, the ringing is not interpolated and will be heard in the program. Not only that, the ringing will be used as context information, causing erroneous interpolation.

Adjust this parameter to ensure interpolation of any ringing. Large clicks (more than 0.010 seconds in length) can exhibit a lot of ringing. Setting the Replacement Center between 5 and 25 percent will ensure that the ringing will fall within the replacement area.

For 78 RPM record clicks, 40 percent is a good value to use. Replacement Center is almost never set higher than 50 percent, as this would shift the interpolation forward, ahead of any ringing.

The Replacement Center parameter can be of use in extraordinary situations. For example, a record that had been broken and then glued back together resulted in clicks of about 2.5 milliseconds in duration, followed by over 10 milliseconds of ringing. A Replacement Width of 15 milliseconds and a centering of 20 percent turned out to be the best setting in this case.

Click Rate and Threshold

These settings are used after analysis to fine tune how many interpolations are actually performed out of the click sites marked in the click detect pass.

A click is a short section of sound whose spectrum differs markedly from that of surrounding audio. Generally, a click contains far more high-frequency energy than the adjacent signal. Each candidate anomaly is assigned a value for total spectral energy. Sites that read below the specified Threshold are disabled.

The Initial Threshold parameter in click detect determines which candidate sites are marked. After that, the Threshold (Rate) parameter can be used to enable and disable clicks from the list.

Threshold

The Threshold parameter may be changed by using the Threshold controls. This value controls the action of the declick pass, regardless of the value specified for the List. Rate and Threshold limit the number of clicks actually interpolated.

Threshold lets you set a specific value that defines the amount of click energy required to trigger interpolation. A Threshold of 100 means that only clicks that have a total energy value higher than 100 (arbitrary units) are interpolated.

Click Rate

A Click Rate of 3.0 clicks per second means that the system will calculate an average of 3 interpolations per second. If the source file is 100 seconds long and the Rate is set to 3, the declicker will interpolate the 300 loudest clicks.

If Threshold is set too high (or Rate is too low) then the declicker will not remove as many clicks as expected. If Threshold is too low (and Rate too high) the number of interpolations increases, along with processing time and the possibility of a bad interpolation.

Decrackle

The Decrackle process is effective for correcting dense, impulsive noise by performing a type of sliding interpolation that isolates good audio between impulses and uses it as the basis for resynthesis.

Impulsive noises in recordings come in two general varieties. Clicks, pops, ticks, and spikes are sizable impulses that break the flow of audio in a way that is comparatively easy to recognize and isolate. Usually, such glitches are spaced far enough from one another so that the audio on either side can be used to reconstitute the area of the click. The Manual Declicking process is effective in attacking this type of problem. (See “Manual Declicking” on page 9.)

The other common type of impulse noise is crackle, in which small impulses crowd against one another, producing a nearly continuous noise, like bacon frying in a pan. This type of artifact requires a different processing approach.

Usually, the Decrackler is used together with Declicking. The general procedure is to perform a light Declicking pass and take out the large, conspicuous clicks. It is acceptable to leave a click here and there. Then run a Decrackle pass, as described in this section.

Decrackle does not distinguish between pure crackle and isolated clicks: It tries to eliminate both. The point of running a Declicking pass first is that Decrackle has only a certain amount of processing for each frame of data. If it has to spend too much getting the larger clicks out, then it will not have enough left for the crackle.

To use NoNOISE Decrackle:

1 Select an audio file, region, or range of audio in an audio track to decrackle.

2 Choose AudioSuite > NoNOISE Decrackle.

Decrackle

3 Adjust AudioSuite settings as appropriate.

See the Pro Tools Reference Guide for tips on using AudioSuite controls.

4 Click Process to perform a Decrackle pass. Be patient, as this process may take awhile (especially with longer selections).

5 When the decrackle pass completes, listen to the new version of your audio. If satisfied, close the NoNOISE Decrackle plug-in. If you are not satisfied, continue with the following instructions.

6 Choose Edit > Undo, to undo the Decrackle process.

7 Adjust the Decrackle parameters and click Process again. See “Decrackle Parameters” on page 20 for details.

8 Repeat until you have removed crackle and other impulsive noise to your satisfaction.

Decrackle Parameters

The following parameters control Decrackling:

Decrackle Parameters

Parameter	Range	Default Values	Suggested	Extreme
Clip Fraction (%)	50 to 100	90	75 to 90	< 75
Synthesis Order	10 to 256	75	75 to 128	65, > 150
Damping Factor	0 to 1	.001	.001 to .015	> .02
Amplitude Weighting	–1 to 1	0	–0.5 to 0.5	–1, 1

The four processing parameters provide considerable control over the results of the decrackle process. Best results are obtained through experimentation.

We recommend that you experiment with the different parameters by decrackling short portions of a file. This provides a good sense of each parameter's effect in a reasonable amount of time.

Clip Fraction

This number affects the detection part of the Decrackle algorithm, in which audio samples are separated into good and bad categories. The synthesis process then replaces the samples that are deemed bad with synthesized material that matches the surrounding sound.

The Clip Fraction is the percentage, or fraction of the samples that will be left in the good category. The range of the parameter is from 50% to 100%.

These samples will pass through the Decrackle process unchanged. The higher the Clip Fraction, the less aggressive the decrackling. If the Clip Fraction is set to 100%, the output file will be a copy of the input file.

For heavier decrackling, make the Clip Fraction smaller.The lowest possible setting is 50%, because any lower allows too little of the original signal to provide a basis for resynthesis. Lowering the Clip Fraction very slightly is enough to make a noticeable difference.

If the Clip Fraction is set too high then there will be some amount of crackle remaining in the program. If the Clip Fraction is set too low you may begin to decrackle good audio, which could result in audible low frequency artifacts.

Synthesis Order

The Synthesis Order determines the precision of the Decrackler's resynthesis. In general, the larger this number, the cleaner and more artifact-free the output. The default value of 75 is suitable for the majority of source materials.

If low-frequency occur in the processed output, try raising the value of the Synthesis Order to 100 or even 128. However, raising the Synthesis Order markedly increases the amount of time required to process the soundfile. Synthesis Order should generally be left at 75 unless the process is producing unacceptable numbers of artifacts.

Paradoxically, processing artifacts tend to occur more with clean recordings that have high signal-to-noise ratios. For (relatively) recent material, such as tape recordings from the early 1950s, it is common to set the Clip Fraction to 0.98 and the Synthesis Order to 128. For 78s from the 1930s, however, it is common to set the Clip Fraction to 75% and the Synthesis Order to 75.

Damping Factor

The Damping Factor affects the way that the Decrackle algorithm tracks high-level transient information. The higher the Damping Factor, the more the process will tend to smooth transients in the source material. Large transients sometimes produce low-frequency artifacts in the output. A small amount of damping, such as the default value of 0.001 will (in most cases) smooth the material just enough to prevent artifacts without adversely affecting transient response. Although the range of the Damping Factor extends as high as 1.0, the highest value recommended for normal work is about 0.015.

If Damping Factor is set too high, there may be a loss of transient response and, in extreme cases, loss of overall dynamic range.

Amplitude Weighting

While Clip Fraction determines the overall percentage of samples retained unaltered in the processed output, Amplitude Weighting determines how these are distributed between high and low amplitude sections of the source. At the default value of 0.0, all portions of the source file are processed equally.

As the value of Amplitude Weighting is increased (more positive), the processing becomes concentrated in higher amplitude sections. One can think of it as decreasing the Clip Fraction in proportion to signal amplitude.

Negative values may also be entered for Amplitude Weighting, in which case processing is concentrated in the sections of lower amplitude. The maximum range of the Amplitude Weighting parameter is from +1.0 to –1.0. In practice, values less than plus or minus 0.5 are used, with possibly higher values for special purposes such as distortion removal.

As Amplitude Weighting diverges (positively or negatively) from zero, it is recommended that the Clip Fraction be reduced by some percentage as well. Otherwise, there may be no processing at all in some regions of the signal.

Using Decrackle to Remove Peak Distortion

The Decrackle function can often ameliorate or remove breakup and distortion associated with high signal levels. By using the Amplitude Weighting factor, decrackling becomes concentrated entirely in the highest peaks of the waveform, using the good portions of the wave to reconstruct the portions that are flat-topped or otherwise distorted. In many cases, this approach is able to restore the distorted portions enough to reduce the audible distortion significantly.

To use Decrackle to remove clipping and other high-level distortions:

1 Select a (clipped) audio file or region in a track.

2 Use the AudioSuite Gain plug-in (or equivalent) to lower the overall gain of the audio by 6 to 10 dB. This ensures that headroom exists for correction. If the source audio is clipped, then it is to be expected that the reconstruction will extend beyond the original top of the waveform. Headroom must be above the clipping level for this reconstruction to take place.

3 Run a Decrackle pass on the resulting audio file with the Amplitude Weighting parameter set high (0.7 or 0.8), and the Clip Fraction somewhat reduced (80% to 85%).

The range of variation in source materials and possible distortion types is huge. Experimentation, using a short section of the source file, is recommended to determine the optimal settings for distortion removal.

chapter 4

Using NoNOISE TDM Plug-Ins

This chapter describes and explains how to use Broadband Denoising and the High Resolution (High-Res) Filters. These NoNOISE plug-ins are TDM plug-ins, and can be used for real-time noise reduction and removal on audio tracks, Auxiliary Inputs, and Master Faders.

Broadband Denoising

Broadband noise, or hiss, is one of the most common forms of audio degradation. Noise can be introduced from any of a number of sources, including the noise floor inherent in analog tape recording and thermal noise from microphones, preamps, and other equipment.

Overview of Broadband Denoising Procedure

To eliminate hiss and other noises, it is necessary to first *analyze* noise content and then adapt or *apply* the denoising process to the material. The basic steps include the following:

1. Analyze the sound file Use the Broadband Analysis AudioSuite plug-in to analyze the noise content, and save the Noise Estimate as a plug-in settings file. The Noise Estimate is shown as an editable green line in the plug-in display. (See "Analyzing the Noise" on page 24.)

2. Apply the Noise Estimate for real-time Denoising Insert the Broadband Denoising TDM plug-in on the audio track you analyzed and import the settings file. (See "Broadband Denoising in Real-Time" on page 27.)

About Broadband Denoising

Broadband Denoising operates by means of analysis and resynthesis. A Fast Fourier Transform (FFT) frequency analysis is performed on a sample of noise from the material to be processed. The level of noise in each of 2048 individual frequency bands is determined. The output of this analysis is a Noise Estimate.

In actual denoising, the source material is also subjected to a 2048-point FFT analysis. The level of signal in each frequency band is compared against a threshold level determined by the Noise Estimate. Based on this comparison, the processing algorithm determines whether a given band at that particular instant contains audio signal or only noise.

If a frequency band is found to include elements of the desired signal, it is left untouched. If it is determined that the signal in that band is only noise, the level of that band is reduced by an amount determined by the Attenuation processing parameter. The results of this comparison and adjustment for all bands is a modified version of the original FFT frequency analysis.

A reverse FFT is then performed using the new, adjusted version of the signal analysis, reconstituting the audio signal with noise attenuated by the specified amount. Because the Denoiser operates in more than 2,000 individual bands, the removal of noise is precise and can leave the original audio signal unaffected.

Broadband Denoise can process only one soundfile at a time, but supports all Pro Tools multi-mono processing features.

Analyzing the Noise

The first step in denoising is to derive a Noise Estimate from the recording to be denoised. The Noise Estimate is the fingerprint of the noise as analyzed from the sound waveform display. It determines local threshold values for each frequency band, or bin.

The Noise Estimate determines the result of the entire denoising process, so it is important to ensure that the estimate taken is valid and represents the true noise floor of the source sound file. Consider the following suggestions for optimum results.

Suggestions for Better Noise Estimates

Work in Sections With many projects, it is usually necessary to derive a separate Noise Estimate for each song, cut, or take. If these are contained in a single sound file, then that file should be denoised in sections, so that the optimal set of estimates and parameter settings can be used for each cut.

The denoising algorithm depends on a constant level and spectrum in the noise floor. Noise floors are seldom constant except within single pieces of recorded music (and sometimes not even then).

Unless there is strong reason to believe that each cut in a compilation was recorded in the same session, with the same equipment, at precisely the same levels, onto the same media, stored in the same way, transferred to the same intermediate media in precisely the same way, and converted to digital samples in the same way, then it is advisable to take separate estimates for each cut or take.

Likewise, if the character or level of the noise floor can be heard to change at all during the recording, then the best results are be obtained by dividing that piece into sections to be denoised individually.

After denoising, the individual sections can be joined together to create a seamless whole. The Noise Estimate should be taken from a relatively short segment in a quiet part of the recording to be denoised.

Denoising Stereo Tracks When denoising stereo material, it is also recommended to derive a separate Noise Estimate for each channel. Or, load a mono-derived settings file into a linked, multi-mono Broadband Denoising TDM plug-in.

For Best Results Optimum results are obtained when the Noise Estimate is taken from a section of pure noise between about 0.3 and 0.5 seconds in length, with a minimum of around 100 milliseconds (0.1 seconds).

The first concern in selecting a segment for analysis is that it represents the noise floor throughout the recording. In many cases, there is an apparently clean segment of noise prior to the start of program. Beware of such segments, as they may not represent the noise in the remainder of the recording.

It is not uncommon for recording engineers to fade up or punch in prior to the start of music, leaving an early clean segment containing only a portion of the real noise floor, leading to an inaccurate estimate.

Taking an Estimate when Noise is not Isolated

In some instances, there may be no section where the noise can be measured without signal. Under such circumstances, select a section that is relatively quiet and free of non-harmonic sounds such as cymbals or bells.

When forced to take a Noise Estimate in the presence of signal, you must correct the situation by editing the Estimate (see "Using Fit" on page 26). If the signal present is harmonic in nature, such as a sustained note or chord, or a vowel sound in the case of spoken word, it is often much easier to identify the frequency components that represent the source signal among the noise.

Taking a Noise Estimate

Once a suitable portion of the sound file is identified, it is time to take the Noise Estimate.

To take an initial Noise Estimate:

1 In any mono audio track, select a short section of audio that is pure noise. (See "Analyzing the Noise" on page 24.)

2 Choose AudioSuite > NoNOISE Broadband Analysis.

3 Click Analyze.

After analysis is completed, the panel changes to show an Estimate display. A green line indicates the editable Noise Estimate.

4 Choose Save Settings As from the plug-in Settings menu, and save the initial estimate. This lets you load your settings into the TDM Broadband Denoising plug-in. After editing and optimizing the Noise Estimate, you can save your final settings as well.

The Noise Estimate divides the area framed by the time selection into windows of 2048 samples each. An FFT analysis is performed on each of these windows, and the results of these FFTs are averaged to produce a smoothed, composite FFT.

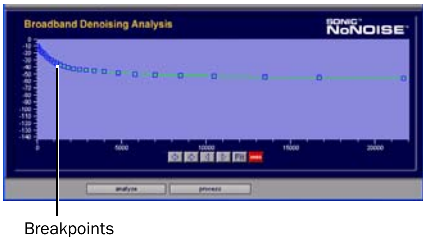

Broadband Denoise Analysis AudioSuite plug-in

You can edit the green curve and breakpoints of the Noise Estimate. See “Editing a Noise Estimate” on page 26 for more information.

After fitting and optimizing the Noise Estimate, you save it as a settings file for use in real-time Broadband Denoising.

Editing a Noise Estimate

For most broadband noise reduction of noise other than typical tape hiss, it is often necessary to edit the estimate manually, differentiating between the program signal and the noise floor.

To edit the Noise Estimate:

1 In the NoNOISE waveform display, select a range of audio and click the Zoom icons to zoom in. The components of the audio signal usually are visible as prominent, more or less evenly spaced spikes in the analysis display. (This is why, if an Estimate must be taken in presence of signal, it is advisable to avoid sections of overly-dense or non-harmonic material.)

2 Edit the Noise Estimate Line by dragging the boxes superimposed over the line. The boxes can be moved only on the vertical axis. If, for example, a number of piano harmonics are infecting the estimate, the boxes can be moved so that the Noise Estimate Line follows the valleys that represent the underlying noise floor.

3 To edit multiple boxes at once, drag in the Noise Estimate display to create a time selection that contains the nodes. The numeric display in the Track header area shows the selection start frequency, selection end frequency, and selection span in Hertz.

4 To edit the boxes within the selection, click and drag a box. All boxes within the selection move together.

Noise Estimate editing requires you to exercise skill and judgment. The ability to extrapolate a useful Noise Estimate Line from a contaminated noise sample will grow with your experience.

Using Fit

Fit produces a smooth Estimate Line from any individual section of the Estimate display. Fit is useful when noise-only content is unavailable. In such a case, you can select the largest noise-only section in the NoNOISE display and use Fit to map the selected range of Noise Estimate across the entire region or file being analyzed.

To fit the Noise Estimate:

- After analyzing and loading the Noise Estimate, click Fit.

NoNOISE displays a Noise Estimate using the selected region of the composite FFT.

Saving a Noise Estimate

To save a Noise Estimate file:

1 Choose Save Settings from the Settings menu in the NoNOISE Broadband Denoise Analysis AudioSuite plug-in.

2 Choose a location and name for the saved file, then click OK. Pro Tools saves the file to disk.

Broadband Denoising in Real-Time

By inserting the TDM NoNOISE Broadband Denoising plug-in on an audio track, Auxiliary Input, or Master Fader, you can apply your Noise Estimate in real time during playback.

To use Real-Time Denoising:

1 Insert the NoNOISE Broadband Denoise TDM plug-in on a Pro Tools audio track, Auxiliary Input, or Master Fader. Often, this will be the same audio track you analyzed to create the Noise Estimate.

2 Import a previously saved Noise Estimate settings file (see “Saving a Noise Estimate” on page 27 for more information).

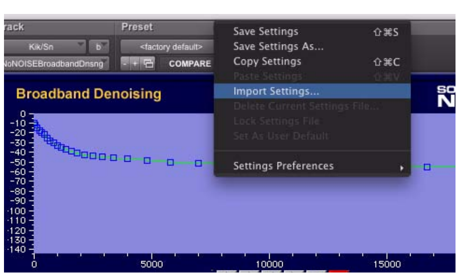

Broadband Denoising, TDM, importing settings

During playback, the Noise Estimate saved in the settings file will be used to silence the described noise from the audio track in real time. You can further adjust the Denoise parameters from within the TDM plug-in window for NoNOISE. (See “Broadband Denoising Parameters” on page 27 for details on the available controls.)

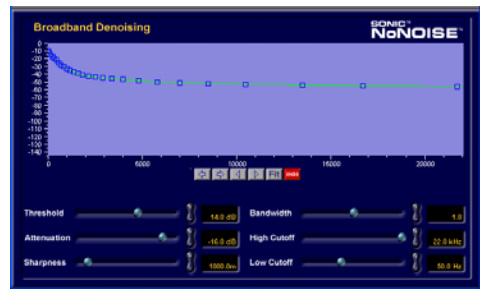

Broadband Denoising, TDM

Broadband Denoising Parameters

The following parameters control Broadband Denoising:

Broadband Denoising TDM Parameters

Parameter	Range	Default Value	Suggested	Extreme
Attenuation	–120 to 0	–16	–10 to –30	
Threshold	–60 to 60	14	8 to 25	30
Sharpness	0.5 to 5	1	1 to 1.5	> 1.75
Bandwidth	0.5 to 3	1.8	.8 to 2.4	
Low Cutoff	0 to 22,050	50	50 to 500	25, 12,000
High Cutoff	0 to 22,050	22,050	22,050	12,000

By carefully adjusting the Real-Time Denoise parameters (especially the Threshold and Attenuation controls) and listening to the results, you can optimize the broadband denoising to remove the greatest amount of objectionable noise while avoiding undesirable artifacts. The Bypass button can also be used to compare processed with unprocessed audio. Sets of processing parameters may be saved.

Attenuation

This value in decibels sets the maximum attenuation to be applied in any frequency band. (A setting of 0 dB applies no noise reduction.) The higher (more negative) this value, the greater the reduction in noise, but with increasing danger of producing audible artifacts in the audio signal.

The amount of noise reduction perceived is normally about half the maximum attenuation ±3 dB depending on the material and the other denoise parameters. A good starting point for the maximum attenuation is to take the amount of perceived noise reduction you wish to obtain and then double it. If, for example, you wish to obtain a perceived noise reduction of around –8 dB, then start with a maximum reduction setting of –16 dB.

Typical values for this parameter range between –10 (mild), –20 (moderately aggressive), and –30 (extreme). If the maximum attenuation setting is too extreme, ambience or high frequency response may be lost.

Threshold

The Noise Estimate defines the curve of the thresholds that apply to each of the over 2,000 individual frequency bands used by the denoise process.

The Threshold parameter allows the curve as a whole to be moved up or down. Together with the Maximum Reduction parameter, this provides the basic throttle that determines how aggressively denoising is applied.

As the Threshold value is raised, more and more of the signal is processed.

At extremely high settings, a distinctive watery sound may be heard on the audio signal. If the Threshold is set too low, little or no noise reduction is obtained.

The Threshold can be thought of as the fine line between noise and music, globally raising or lowering the entire Noise Estimate curve relative to its original position. Local adjustment of the threshold according to frequency bands is effected by adjusting the Noise Estimate curve.

The Threshold point is set (somewhat arbitrarily) at –6 dB. Threshold settings and maximum reduction settings should generally be adjusted together for best results. Typical values for this parameter range between 8 (mild), 16 (moderately aggressive), and 25 (extreme).

Sharpness

The denoising process works much like a multiband downward expander.

As signal level in a particular band drops, the process reduces the gain in that band even further, using an internal attenuation curve.

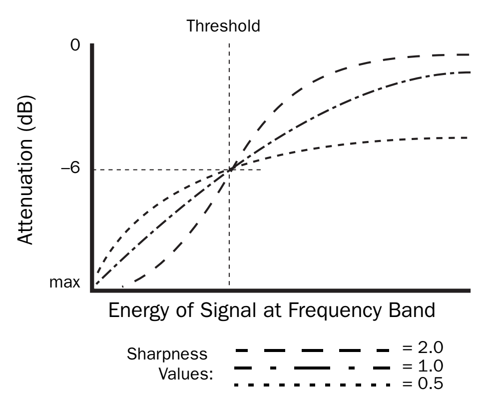

Sharpness Curves

The Sharpness parameter sets the slope of this curve. Higher values cause quicker attenuation as instantaneous energy falls off from the Threshold value, resulting in a response similar to that of a noise gate. If the Sharpness is too low, there may be no reasonable amount of noise reduction possible, despite the settings of other parameters.

Generally this parameter should be set as high as possible without audible ill effects on the program. A value of 1.0 is recommended for common tape hiss problems, while a value of about 1.2 is useful for standard 78 RPM type recordings.

If the sharpness is too high, you may hear a phasing problem in the music, sometimes described as an underwater effect. The noise that remains may also become more unstable, producing a rapid fluttering of the noise floor.

Bandwidth

The Denoiser has been described as a multiband downward expander with many individual bands. Actually, there is a bit more involved. Denoising with each band adjusted separately produces an unnatural-sounding result.

In NoNOISE, individual bands share information for more natural sound.

Bandwidth governs this process. Higher values produce more sharing. A higher value (in most cases) creates a more natural sounding result, but with the risk of audible pumping of the residual noise floor. A low value for Bandwidth eliminates the possibility of noise pumping, but may sound less realistic and more muffled. Typical values to use here might be 0.8 (little sharing), 1.8 (good standard setting), 2.4 (a lot of sharing).

The optimum Bandwidth setting depends entirely on the program material. It should be high enough to retain high frequency response but low enough to avoid pumping or distortion caused by noise being modulated by the high harmonics of a signal.

High Frequency Cutoff

The denoise function lets signal in frequency bins above the high frequency cutoff point pass through the process untouched (that is, they are not processed). This limits processing to frequencies below the high cutoff point. This can be used in situations where noise is not objectionable above a certain frequency, but in most cases, this parameter is left set at 22,050 Hz.

In extreme cases, good results may be obtained by processing the upper and lower frequencies separately, using different parameter values. Run the audio through the denoiser twice, using the low and high cutoff frequencies to define the area to be denoised.

Low Frequency Cutoff

This parameter allows signal below the specified frequency to pass through unchanged. This can be useful if noise is not objectionable below a certain frequency and you wish to leave it alone. Low Frequency Cutoff is usually left set around 50–100 Hz. If cutoff is below 25 Hz, there may be artifacts because wavelength exceeds the analysis window.

High-Res Filters

The NoNOISE High Resolution (High-Res) Filters provide many types of filters (see “High-Res Filter Types and Parameters” on page 31).

To use a High-Res Filter:

1 Insert a NoNOISE High-Res Filter plug-in on any audio track.

2 Select a process from the Type selector.

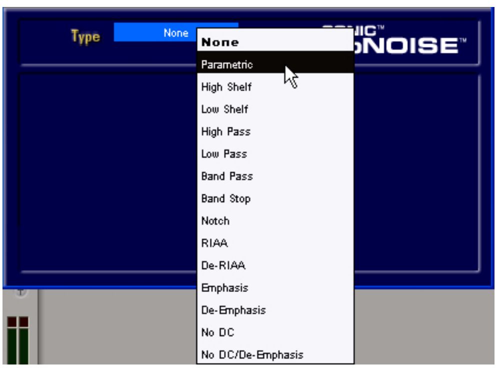

Filter Type selector in the High-Res Filter plug-in

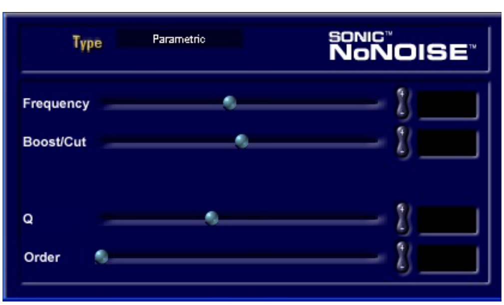

Parametric filter

General Filter Parameters

Frequency

The frequency is the reference point of the filter. There are two different interpretations of the frequency in NoNOISE, depending on the filter type used:

- In filters that have a bandwidth or Q parameter, the frequency references the midpoint of the affected region. Usually the frequency is the most affected frequency of these types of filters.
- Filter types that affect frequencies above or below a particular frequency reference the frequency (cutoff frequency) as the –3 dB point from the boost or cut specified. A High Shelf filter designed to give a 6 dB boost above a crossover frequency (frequency) of 10kHz would have a boost of 3 dB at 10kHz.

Q

Bandwidth and Q are two different ways of specifying the width of the filter. The width of the filter is measured from the –3 dB down points on either side of the frequency. Bandwidth represents this width in absolute Hz. A bandwidth of 1000 means that the filter is 1000 Hz wide between the –3 dB points. Q represents the width of the filter relative to the way that we hear.

A Q setting of 2 always has a half-octave bandwidth regardless of the frequency. In mathematical terms, the Q is equal to the frequency divided by the bandwidth.

- Q = Frequency / Bandwidth
  - Also: Bandwidth = Frequency / Q
  - A Q of 1 = a one-octave filter width. A Q of 2 = a half-octave filter width. A Q of 4 = a quarter-octave filter width. A Q of 0.5 equals a 2-octave filter width.

Boost/Cut

The Boost/Cut represents the maximum effect of the filter on the program material. Boost/Cut is expressed in positive or negative decibels (dB).

Order

The order of the filter sets the slope of the filter's transition area. A first order filter usually means that the filter has a transition slope of 6 dB/octave. Each increase in the order adds another 6 dB/octave to the transition. Thus, a 3rd order filter would have a transition slope of 18 dB/octave, and so on. Band Pass and Band Stop filters require an even-numbered order. A second order Band Pass/Band Stop has transition slopes of 6 dB/octave and a fourth order Band Pass/Band Stop has transition slopes of 12 dB/octave.

Stop

Stop expresses the amount of stopripple in the stopband of the filter. The interpretation of Stopband and Stopripple depends on the filter type. (See "High Pass and Low Pass Filters" on page 33.)

Stopband The portion of audio eliminated or attenuated.

Stopripple The minimum attenuation in the stopband. A setting of –40 dB means the signal will be at least 40 dB down in the stopband.

High-Res Filter Types and Parameters

Parametric Filter

The Parametric filter boosts or attenuates a particular region of the audio spectrum.

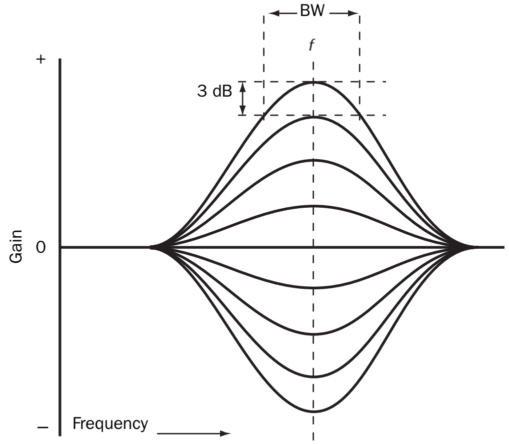

Parametric Filter

There are three parameters that define the response of the Parametric filter type:

Frequency The Frequency is the midpoint of the band affected. The Parametric filter's frequency may be selected over a range of 1.0 Hz to 22,050 kHz.

Q The Parametric filter is a resonance-type filter. Its bandwidth in Hertz may be translated to filter Q by the formula: BW = CF/Q.

Boost/Cut Boost indicates the gain applied at the Frequency. The Parametric filter can supply boost (cut) of ±24 dB.

Notch Filter

The Notch filter is a special case of the Parametric filter, in which the gain at the frequency is set to minus infinity, effectively eliminating all signal at that frequency.

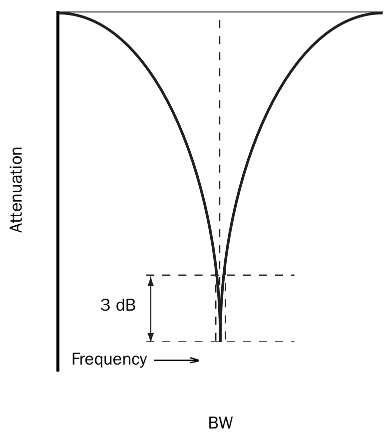

Notch Filter

Specifying a notch filter requires only two parameters: frequency and bandwidth. The gain at the frequency of a notch filter is fixed (at minus infinity dB), so that the frequency is eliminated completely.

Frequency The range of the Notch filter's frequency is from 0.1 Hz to 1/2 of the Nyquist frequency (the Nyquist frequency is half the sample rate).

Q The Notch bandwidth is specified in Q values, with a range from 0.1 to 100.

High Shelf and Low Shelf Filters

Shelving filters apply a fixed boost or cut to all frequencies beyond the cutoff frequency. Shelving filters have two variable parameters. These are different from those of Parametric filters.

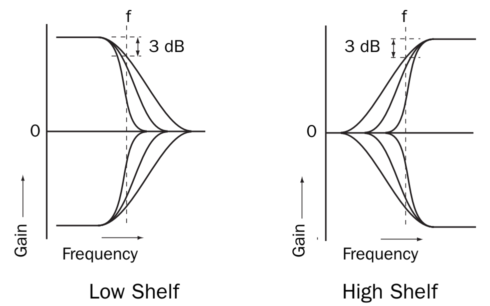

Low Shelf and High Shelf Filters

Frequency Frequency is the point where the signal is boosted or cut by 3 dB, or by 1/2 the specific boost/cut (if less than 3 dB).

Boost/Cut Boost (or Cut) applies to signal above (in the case of High Shelf) or below (in the case of Low Shelf) the cutoff frequency. The range of boost or cut is ±24 dB.

High Pass and Low Pass Filters

The High Pass and Low Pass filters operate in three regions: A passband where signal is minimally altered; stopband where signal is attenuated; and the transition band that separates the two.

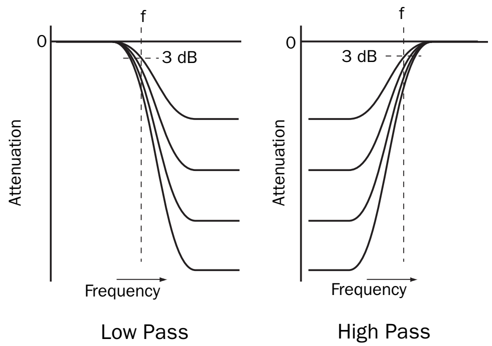

Low Pass and High Pass Filters

In the passband, there is an amplitude fluctuation called passband ripple. Generally, anything less than about 0.5 dB of ripple is inaudible. Passband ripple defaults to 0.1 dB.

There is also some variation in the stopband, called stopband ripple. It represents the maximum value that the filter gain attains in the stopband. For a high pass filter, the frequency parameter refers to the highest frequency at which the gain attains the minimum value in the passband. A low pass filter has a gain of 1.0 (less the pass ripple) below f, and falls off smoothly to stop ripple at some point above f.

Frequency

At the Cutoff Frequency, the signal level is reduced by 3 decibels. The range of cutoffs for both High Pass and Low Pass filters is from 1 Hz to 22.05kHz

Order

Filter order, an integer from 2 to 16, controls the steepness of the transition from stopband to passband. The transition band drops off by roughly 6 dB per octave for each unit increase in order. The higher the order, the greater the chance of audible ringing at the cutoff frequency.

Stop For Low and Highpass filters, Stop is the amount of stopband ripple, the maximum gain reached in the area after the cutoff frequency.

RIAA, Emphasis, and DC Filters

The High-Res Filters includes several types of utility filters for common functions of DC removal, pre- and de-emphasis, and application or removal of the RIAA curve for vinyl records.

These filters have no variable parameters. They are simply on or off.

No DC

The No DC filter is a simple DC reject filter. The No DC filter provides 1 dB of cut at 34 Hz and 3 dB of cut at 18 Hz.

Emphasis and De-Emphasis

The Emphasis filter is a 15/50 microsecond curve, as defined as an option for Compact Disc masters. The De-Emphasis filter provides for removal of this high-frequency boost from material that is previously emphasized.

No DC/De-Emphasis

No DC/De-Emphasis combines this with a filter to remove the Sony F1 (EIAJ digital audio adapter) 15/50 microsecond emphasis curve.

RIAA and De-RIAA

NoNOISE supports RIAA and De-RIAA filters. The RIAA filter imposes the standard RIAA characteristic normally applied, in LP mastering, at the input to a disk cutting lathe. The De-RIAA filter removes the effect of a RIAA filter.

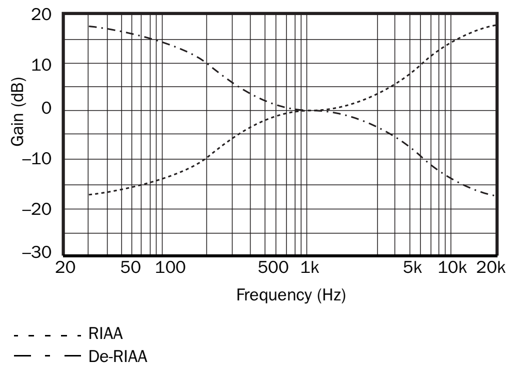

RIAA and De-RIAA

Band Pass and Band Stop Filters

Band Pass and Band Stop are like putting together a high pass and a low pass filter.

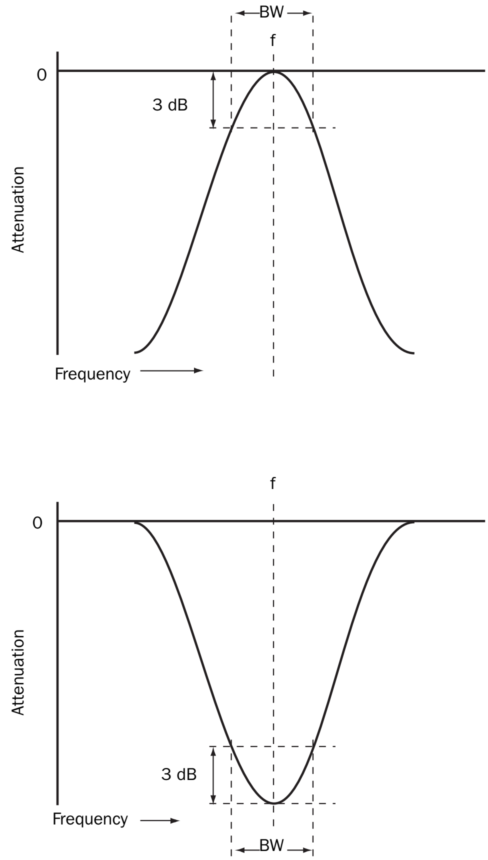

Band Pass (upper) and Band Stop (lower)

Band Pass allows only certain frequencies to be admitted and rejects all others that are out of the range.

Band Stop eliminates a certain range of frequencies and passes all the rest.

Frequency The range of the Band Pass and Band-Stop filter's frequency is from 1 Hz to 22.050kHz.

Q Bandwidth is specified in Q, with a value range from 0.1 to 100.

Order (n) The order of a filter, an integer value from 2 to 16, determines the slope of the filter's response curve. A first-order filter has a slope of 6 dB per octave.

Slope increases by 6 dB per octave for each increment of 1 to the filter's order.

appendix a

DSP Requirements

The number of TDM plug-ins you can use at one time depends on how much DSP power is available in your system. Since the TDM hardware on Pro Tools cards provide dedicated DSP for plug-ins, plug-in performance isn't limited by CPU processing power.

The DSP tables in this appendix show the theoretical number of instances of each plug-in that can be powered by a single DSP chip on Pro Tools|HD cards. DSP usage differs according to card type.

DSP tables show the theoretical maximum performance when no other plug-ins or system tasks (such as I/O) are sharing available DSP resources. You will typically use more than one type of plug-in simultaneously. The data in these tables are provided as guidelines to help you gauge the relative efficiency of different plug-ins on your system. They are not guaranteed performance counts that you should expect to see in typical real-world sessions and usage.

There are a total of nine DSP chips on a Pro Tools|HD card (HD Core™, HD Process™, and HD Accel). HD Core and HD Process cards provide identical chip sets. HD Accel cards provide newer, more powerful DSP chips (making the HD Accel card ideal for DSP-intensive plug-ins, and for high sample rate sessions).

Not all plug-ins are supported on all types of chips. The following tables indicate the number of compatible chips per card.

Using Multi-Mono Plug-Ins on Greater-Than-Stereo Tracks

Plug-Ins used in multi-mono format on greater-than-stereo tracks require one mono instance per channel of the multi-channel audio format. For example, a multi-mono plug-in used on a 5.1 format track, requires six mono instances since there are six audio channels in the 5.1 format.

Monitoring DSP Usage

The System Usage window (Window > System Usage) shows how much DSP is available in your system and how it is being used in the current Pro Tools session.

For more information about DSP usage and allocation, see the Pro Tools Reference Guide.

Sonic NoNoise DSP Requirements

The NoNoise plug-ins have the following DSP requirements:

HD Accel Card

Table 2. Maximum instances of plug-ins per DSP chip for an HD Accel card at different sample rates.

Plug-In	44.1/48 kHz	88.2/96 kHz	Compatible DSP Chips per HD Accel Card
Broadband Denoising TDM	1	1	6
High-Res Filters	9	4	6

HD Core and HD Process Cards

Table 3. Maximum instances of plug-ins per DSP chip for an HD Core or Process card at different sample rates.

Plug-In	44.1/48 kHz	88.2/96 kHz	Compatible DSP Chips per HD Accel Card
Broadband Denoising TDM	1	1	9
High-Res Filters	4	1	9

appendix b

DSP Delays Incurred by TDM Plug-Ins

Virtually all TDM plug-ins incur some amount of signal delay.

This is significant only if you use a plug-in on one channel of a stereo or multichannel signal but not the others, since this can cause the channels to be slightly out of phase.

If you are working with mono tracks, or are processing all channels with the same plug-in, the signal delays are not long enough to be significant and should not be a concern.

Pro Tools systems provide automatic Delay Compensation (and other methods) to compensate for signal processing delays. For detailed information, see the Pro Tools Reference Guide.

Table 3 shows the delays inherent in each plug-in.

NoNoise DSP Delay

Table 3. Samples of delay incurred by each TDM plug-in on Pro Tools|HD cards

Plug-In	Samples of Delay on HD Cards
Broadband Denoising TDM	14
Big High-Res Filters	3

index

T

W

Avid
2001 Junipero Serra Boulevard
Daly City, CA 94014-3886 USA

Technical Support (USA)
Visit the Online Support Center at
www.avid.com/support

Product Information
For company and product information,
visit us on the web at www.avid.com