

# DeNoise 2.7 & DeNoiseLF 2.7

[www.kagi.com/clickrepair](http://www.kagi.com/clickrepair)

(for Java-enabled platforms)

## User's Manual

*Brian Davies*

[clickrepair@kagi.com](mailto:clickrepair@kagi.com)

*Preamble – Please Read.*

*In 2004, I began to write software for removing clicks heard on typical LP recordings, once they have been captured for transfer to CD. The resulting application, ClickRepair, now has a wide circle of users from whom I have received much encouragement. DeNoise and DeNoiseLF are applications designed for the reduction of continuous noise. They have been in experimental use since 2006 and they are now released for general use. DeNoise is designed in particular for the reduction of hiss and other broadband noise, DeNoiseLF for the reduction of rumble and hum.*

*If you find these applications useful, please help by posting information to news and discussion groups, and ratings feedback on the site where you found it. This way the circle of users will increase, and I shall feel encouraged to devote the time and effort required to maintain and improve the software.*

---

DeNoise and DeNoiseLF are written in Java, a system which allows the same application to run on different platforms (Mac OS X, Windows, Linux). The Windows installer is produced using “Inno Setup 5” by Jordan Russell. The manual is produced using Apple’s “Pages” application, typeset in Garamond, and exported as PDF.

I wish to thank those users who have contributed to the design of the Interface. I owe a particular debt to Michel Poulain and Willy Kunz, who spent many hours working on the French and German translations, and also provided much feedback.

# Table of Contents

<b>1</b>	<b>Introduction</b>	<b>5</b>
	<i>Basics</i>	5
	<i>Processing Strategy – DeNoise</i>	6
	<i>Processing Strategy – DeNoiseLF</i>	6
	<i>Important Warnings</i>	7
	<i>Common Feature of ClickRepair, DeNoise, and DeNoiseLF</i>	7
	<i>Organization of Manual</i>	7
	<i>Recommendations</i>	8
	<i>Further Information</i>	8
<b>2</b>	<b>Installation and Registration</b>	<b>9</b>
	<i>Apple Mac OS X</i>	9
	<i>Microsoft Windows</i>	9
	<i>Registration</i>	10
	<i>Access to Registration Information</i>	10
<b>3</b>	<b>Quick Start – DeNoise</b>	<b>12</b>
	<i>Initial Interface</i>	12
	<i>Basic Controls</i>	12
	<i>Choosing Files</i>	13
	<i>Setting the Noise Profile</i>	13
	<i>Using a separate Noise Sample File</i>	15
	<i>Processing</i>	15
	<i>Progress Indicator</i>	15
<b>4</b>	<b>Producing and Organizing Files</b>	<b>17</b>
	<i>File Requirements</i>	17
	<i>Recommended Order of Processing</i>	17
	<i>Notes about File Structure</i>	18
	<i>File-lengths</i>	19
	<i>Sample Sizes</i>	19
	<i>Reading from CDs</i>	20
<b>5</b>	<b>User Interface Basics</b>	<b>22</b>
	<i>Visual Interface</i>	22

	<i>The Sound Bar</i>	22
	<i>The Display Window and Scale Controls</i>	23
	<i>Menus</i>	23
	<i>Keyboard Shortcuts</i>	24
	<i>Audio Monitoring</i>	25
	<i>Sound Latency</i>	25
<b>6</b>	<b>Using DeNoiseLF</b>	<b>27</b>
	<i>Rumble Reduction</i>	27
	<i>Hum Reduction</i>	29
	<i>Removing Harmonics</i>	30
<b>7</b>	<b>Audio Output and the Sound Buffer</b>	<b>31</b>
	<i>Java Sound Output settings</i>	31
	<i>The Sound Bar</i>	31
	<i>Appearance of the Display</i>	32
	<i>Experimenting using Preview</i>	33
	<i>Committing a Preview</i>	33
<b>8</b>	<b>Settings – DeNoise</b>	<b>34</b>
	<i>Spectral Audio Information</i>	34
	<i>Spectral Gain Information</i>	35
	<i>Adjusting the Noise Floor</i>	36
	<i>Achieving a Compromise</i>	36
	<i>Using the Automatic Facility</i>	38
	<i>Adjusting the Reduction Limit</i>	38
	<i>Preset Noise Profiles</i>	38
	<i>Noise Floor Balance</i>	39
<b>9</b>	<b>Batch Processing</b>	<b>41</b>
	<i>Organization</i>	41
	<i>Assembling a Batch</i>	42
	<i>Processing</i>	43
<b>10</b>	<b>The Preferences Dialog</b>	<b>44</b>
	<i>Preferences</i>	44
	<i>Language</i>	44

<i>Sound Buffer Length</i>	44
<i>Log Files</i>	45
<i>Files and Processing</i>	45
<i>Breakpoints and Markers</i>	45
<i>Resetting Defaults</i>	46
<b>11 Other Features</b>	<b>47</b>
<i>Behaviour of Buttons</i>	47
<i>Abort Button</i>	47
<i>File Dialogs</i>	47

# 1 Introduction

## Basics

DeNoise and DeNoiseLF are applications for reducing continuous noise in audio files, typically those produced by digitizing vinyl (LP) and shellac (78) records, or magnetic tape. The most common form of such noise is commonly heard as hiss, rumble, and hum.

Unless you are prepared to understand at least something of the problems, and how the applications can be used to address them, then the best results may be obtained by doing nothing!

Many potential users of these applications will already have been using ClickRepair for some time, and will find this harsh admonition somewhat surprising, given the ease with which the latter can be applied to a wide variety of audio material using default settings. The vital difference is:

- Finding and removing clicks can be likened to looking for objects which have been dropped on a carpet. You can hunt for them and pick them up, and you can teach a robot to do that too. You can also smooth over any damage done to the carpet.
- Hiss, and other continuous noise, is more akin to trying to remove the milk from a cup of tea or coffee. It can't be done, and hiss can't be removed either, except in the unusual circumstance that the noise and the signal lie in non-overlapping frequency bands.

Let me elaborate on the second point:

- The bad news is that there is no way, given only the noisy audio file, to exactly recover the original audio, thereby achieving an exact reconstruction.
- The good news is that, in many cases, the perceived noise may be reduced to a level that is unobjectionable, maybe even barely perceptible, without introducing objectionable distortion and processing artifacts in its place.

The applications work with mono or stereo sound files up to 32 bits at 192 kHz. This includes standard CD and DVD audio formats. They are designed to run automatically once appropriate settings have been chosen. It is also possible to pause the programs at any time so as to perform re-processing with different settings.

The maximum sample rate was increased in version 2.2 of DeNoise and DeNoiseLF. When the sample rate exceeds 96kHz, DeNoise down-samples to half of the incoming rate, then up-samples to produce output at the original sample rate. For files captured at 176.4kHz or 192kHz, there is a consequent ultra-high frequency roll-off, from about 35kHz, which should not be an issue if the audio material is in need of hiss reduction. DeNoiseLF uses wavelets, so there is no impact on the ultra-high frequencies.

## Processing Strategy – DeNoise

The DeNoise application works in the frequency domain, dividing the incoming audio stream into 80 overlapping frames per second, and transforming each such frame into 1024 equally spaced frequency bands (2048 if the sample rate exceeds 48 kHz).

For sound sampled at the CD rate – 44.1 kHz – this means that each band has a width of 22Hz. From about 200Hz this enables fine-grained processing; however it is not satisfactory for the treatment of rumble, hum, and the like. This is the reason for the development of a separate application (DeNoiseLF) for the treatment of low-frequency noise.

Most readers will be familiar with the operation of graphic equalizers. Normally they are restricted to no more than 30 bands, and they are adjusted infrequently. DeNoise operates as a type of “intelligent” graphic equalizer. There are 1024 (or even 2048) frequency bands, and the incoming audio is processed in overlapping frames at the rate of 80 frames per second. The gain in each of the bands is adjusted continuously from frame to frame, using algorithms which attempt to judge what would be perceived as noise, and what would be perceived as desirable audio. The objective is to suppress the former without doing perceptible damage to the latter.

## Processing Strategy – DeNoiseLF

A little thought will show that the above approach runs into severe difficulties for low-frequency noise, particularly hum. The latter will require a frequency resolution in the order of 1Hz, possibly better. For CD quality sound, this implies at least 22,000 frequency bands, requiring FFT transformation of frames of about 1 second in length. At a frame rate of 80 per second this is impractical; lowering the frame rate is detrimental to the reduction of hiss and other high-frequency noise.

The fundamental limitation is a simple matter of mathematics: the more accurately one determines the frequency content of a signal, the less accurately the timing of that content may be determined. In quantum physics, a similar relation is called the Heisenberg uncertainty principle.

Wavelets are a natural signal analysis method incorporating this relationship between frequency and time resolution. DeNoiseLF makes extensive use of wavelets to separate out the low-frequency audio components before treatment. Treatment itself is done using a sophisticated algorithm which attempts to track the noise and the audio content and separate one from the other. Wavelet reconstruction is then applied so that the higher frequency content is reconstructed exactly.

The algorithms are statistical estimation techniques. It is simply impossible to reconstruct the original audio precisely. Statistics deals with probabilities, not certainties.

Since low-frequency noise can affect the operation of ClickRepair, while DeNoiseLF will not affect clicks and crackle in any way, it may be advantageous to apply it first.

## Important Warnings

While excellent results are possible with appropriate use, please note that successful operation depends critically on the settings, which feed into the models used for the analysis. Note, therefore:

- Choosing the settings requires skill, which will only be gained by becoming familiar with what is involved, and by careful listening. The applications are based on science; their use is an art.
- DeNoise has no default settings, only some pre-set profiles for common noise types.
- DeNoiseLF has different algorithms for general noise (rumble, etc.) and hum, and a choice must be made before starting processing.
- If you find settings which work well, you should save them for future reference. Even though they may need adjustment in a different situation, they may be a good starting point.
- Do not delete the original file until you are happy with the quality of the new file.
- The Quick Start chapter is not a substitute for reading the manual – after you have made some initial experiments, please take the time to read it.
- Before purchasing a license key, be sure to appraise the usefulness of the software for your particular needs. That is the purpose of the trial period!

## Common Feature of ClickRepair, DeNoise, and DeNoiseLF

These applications implement completely different scientific methods, but their actual coding as useable software involves about 70% of code in common. This includes:

- The handling of sound files, supported audio formats, input and output, etc.
- The organization of the main interface as presented on the computer screen.
- Implementation of preferences, and custom settings.
- Batch processing.
- Shortcuts, menus, etc.

This implies that some of this manual is a duplication of the ClickRepair manual.

## Organization of Manual

This manual is organized into the following chapters:

- Installation and Registration (Chapter 2).
- Quick Start – DeNoise (Chapter 3).
- Producing and Organizing Files (Chapter 4).
- User Interface Basics (Chapter 5).
- Using DeNoiseLF (Chapter 6).
- Audio Monitoring and the Sound Bar (Chapter 7).
- Settings – DeNoise (Chapter 8).
- Batch Processing (Chapter 9).
- The Preferences Dialog (Chapter 10).

- Other Features (Chapter 11).

## **Recommendations**

Once you have downloaded and installed the software, I recommend that you start the program, and use Chapter 3 (Quick Start – DeNoise) and/or Chapter 6 (Using DeNoiseLF) to process the sample files off the website. This will serve:

- As an initial introduction to the application(s).
- To ensure that you are properly set up, in relation to producing the required files, using whatever software you choose for that purpose.

Subsequent exploration of the features of the software, and the corresponding chapters of this manual, may then be made as necessary. For DeNoise, a careful reading of Chapter 8 is mandatory. For DeNoiseLF, please read Chapter 6 carefully.

## **Further Information**

The design of DeNoise and DeNoiseLF is directed toward users who want, primarily, reliable automatic treatment of sound files, using a stand-alone tool. There are no immediate plans for plug-in versions.

A detailed history of all versions of the applications is on the web-site, [www.kagi.com/clickrepair](http://www.kagi.com/clickrepair).



## 2 Installation and Registration

### Apple Mac OS X

Users should download the file DeNoise.dmg. Upon double-clicking the file, it will mount as a volume. The contents should be dragged to the desired location on your hard disk – most likely the applications folder.

DeNoise and DeNoiseLF for Mac OS X require a Java run-time, and have been tested on Mac OS X 10.3 to Mac OS X 10.7. The latter version (Lion) does not have Java pre-installed but a Java run-time is available from Apple – you will be prompted to install it the first time you attempt to run a Java application. The situation with the upcoming Mac OS X 10.8 is unclear - visit my website for information which will be posted when it is available.

There is no installer for Mac OS X. I recommend that you create a folder “DeNoise” in your Application folder, and keep the applications and the manual in it. For easy access, drag the Application icons to the dock as well.

### Microsoft Windows

Microsoft discontinued Java deployment for Windows some time ago. DeNoise users should download and run the installer DNinstall.exe, which does not come with a Java run-time. You will be prompted to install one, if necessary. The applications have been tested with Java 1.6 on both Windows XP and Windows 7.

Shortcuts to the Application and to the Manual are installed in the Windows Start>Programs menu at Start>Programs>DeNoise.

### Alternative Deployment (All Platforms)

Linux, Mac OS X and Windows users can all run the Java “jar” files directly. To do this, download and unpack the Linux version (a zip file), which will deliver two jar files and this manual (pdf). These jar files are identical to the jar files included in the Mac or Windows versions, but they lack the normal packaging of applications specific to those platforms. Double-clicking DeNoise.jar or DeNoiseLF.jar should launch the applications using the built-in Java launcher, provided the system is equipped with a modern Java run-time. Linux users may need to invoke them from the command line, or they may respond to double-clicking. This method of deployment can be useful, on Mac or Windows systems, if the launch mechanisms that are part of the standard downloads fail to work for some reason. There is no package or installer provided for Linux users.

## Registration

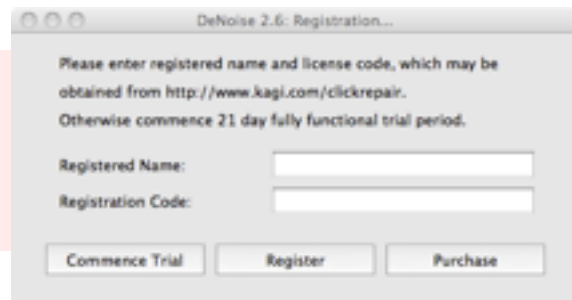
Each time you start DeNoise, a registration dialog will appear. Click “Commence trial” or “Continue trial” in order to evaluate the program. There is a limit of 21 days trial use. Evaluation mode imposes no limitations on the features of DeNoise. However, you will be reminded of the number of days remaining each time you start the program.

DeNoiseLF also has a 21 day trial period, but has no registration mechanism. It will be registered automatically once DeNoise is registered, that is, a single license suffices for the pair of applications.

License costs depend on the method of purchase. Click the “Purchase” button, or follow the link from the home page: [www.kagi.com/clickrepair/license](http://www.kagi.com/clickrepair/license) (under the “Software” tab), for further details.

Some users, for example residents of many European countries, will be charged applicable tax (VAT, TVA, etc.) at the relevant rate for their state/country of residence.

The programs are licensed to an individual user. If you give it to a friend (for evaluation – please do!) they should buy their own license key to continue using it.



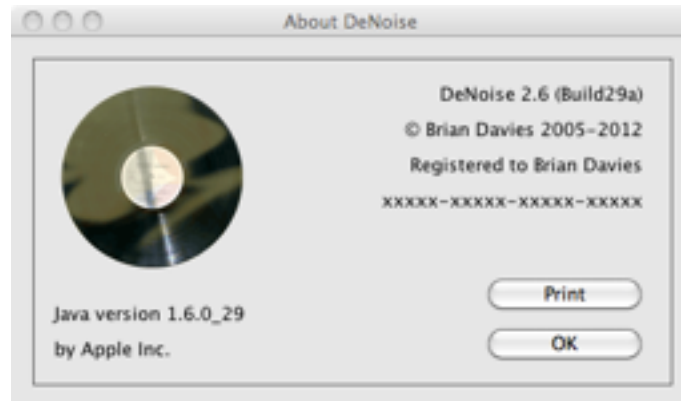
When purchasing the key, you will be asked for your name – minimum six characters. This information is encoded into the license key, and you will need it again, in precisely the same form, to register the program. (Note that it is not case sensitive – John Brown, JOHN BROWN, and john brown are equivalent.) The information should be emailed automatically by Kagi, to the email address you gave, once the transaction is approved.

## Access to Registration Information

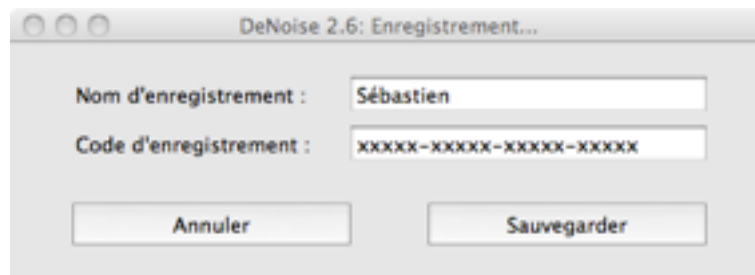
Immediately after registering DeNoise, the “About” box will appear, as shown below. Notice the print button – if you have a printer, I recommend printing the information immediately.<sup>1</sup> Otherwise, write it down and keep it in a safe place. The information will also remain available by choosing “About DeNoise” in the menus.

---

<sup>1</sup> In fact, on the first couple of launches after registering, the About Box can only be dismissed either by using the “Print” button or by waiting for 10 seconds.



Users may also access the information via the “Registration” item on the “Help” Menu. The resulting dialog is shown below. The example is for a French user, whose name had an accent that was mangled by the Kagi website (Sébastien was SA(c)bastien). Although the user name may not be changed in a way which disagrees with the license code, accented characters are recognized and used by DeNoise when making this check. Note in the example below that the user has entered the correct characters in the DeNoise registration despite the incorrect Kagi website translation. Also, upper- and lower-case letters may be interchanged.

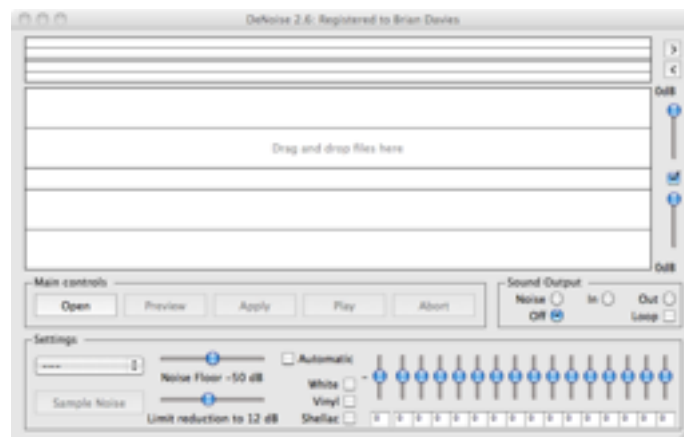


The information in this dialog may be saved into a text document using cut and paste.

## 3 Quick Start – DeNoise

### Initial Interface

On starting DeNoise, the default window has the appearance shown below. There are controls that affect the way that noise is estimated, and a large window in the middle for reviewing the settings – initially blank. Do not be put off by the multitude of sliders; the 16 vertical sliders on the bottom right, together with the “Noise Floor” slider, set the noise profile, and this can often be done automatically. In the picture here the profile is flat and the “White” checkbox is automatically selected to indicate this.



### Basic Controls

DeNoise does not rely on menus for the most important operations, just buttons, check-boxes, and sliders. The more important functions have keyboard shortcuts, which match the menu shortcut keys. Shortcut keys may be displayed in a separate window, accessible from the “Help” menu. For example, the first step in using the application requires a file to be chosen for processing. The alternatives are:

- Drag and drop drop one or more files.
- Select “Open” from the File menu.
- Use the menu shortcut key + “O” to make this menu selection.
- Click the “Open” button using the mouse.
- Press the “Return” key – see Chapter 5 for details of how to find these shortcuts.<sup>2</sup>

Drag and drop is new since DeNoise 2.4. If only one file is dropped, it becomes the input file, for immediate processing and the output file – a new file with a different name – is automatically created. The convention for naming this file, and the choice of folder where it goes, is determined on the basis of the preference settings. A dialog will appear showing the full path name of the output

<sup>2</sup> Unlike menus, this last action is not coupled with using a menu shortcut key. A complete list of keyboard shortcuts is given in a later chapter; menu shortcuts may be read off directly from the menus.

file. If more than one file is dropped, the “Batch Setup” dialog is opened and the files are added using the normal rules for batch processing. See Chapter 9 for details of how this works.

If you want complete control over the naming and destination of the output file, do not use drag and drop – use the Open command. After a dialog allowing the input file to be selected, there will be a second dialog allowing the output file to be named and placed.

Setting of the sliders and check-boxes may be done individually, or collectively by using the “Settings” menu or the on-screen popup menu to access the pre-defined and custom settings. This is all described in detail in later chapters.

## Choosing Files

Before you can go on, it is necessary to decide how to organize input and output files. DeNoise operates by opening the (input) file to be processed, as read only, together with a second (output) file to which the repaired audio is written. The files must be uncompressed, mono or stereo, in 16, 24, or 32-bit “AIFF” or “WAVE” format. The output file will be in the same format as the input file. More details of file formats are given in the next chapter.

To choose the files, click the “Open” button, which will bring up a dialog for choosing the input file. Once this is done, a second dialog will appear for naming the output file.

The two files may not be the same file. This means that either they must reside in different folders (directories), or they must have different names.

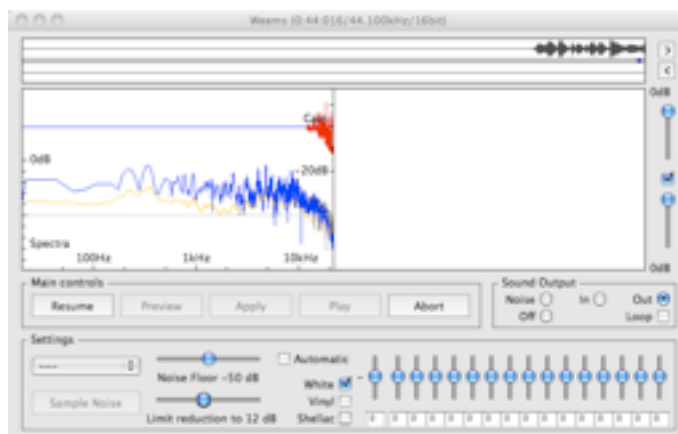
DeNoise has two modes, controlled by the “Append -nr” check-box in the Preferences. If this option is selected (the default), the save dialog will assume that you want to use this when naming the output file in relation to the input file, i.e., “beatles.aiff” becomes “beatles-nr.aiff”. If the option is de-selected, the save dialog will suggest that the output file have the same name as the input file, and DeNoise will refuse to save it in the same folder (with the suggested name).

Each time you choose an input file, the dialog will go back to the folder previously selected. Each time you choose a location for an output file, it will go back to the previously chosen folder either for input or output, depending again on a choice which may be set in the Preferences. Of course, you are free to make changes using the dialogs themselves.

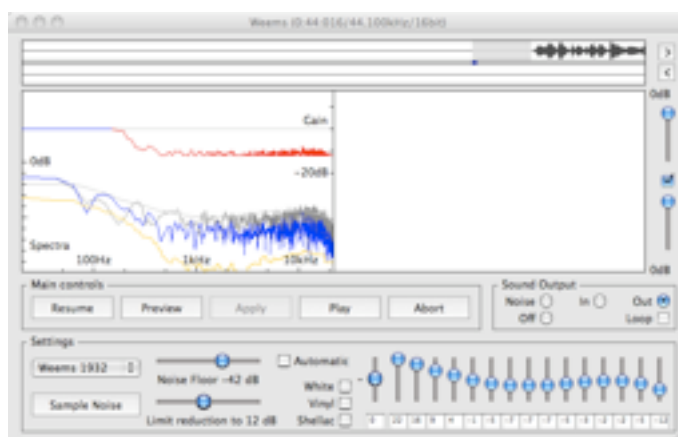
## Setting the Noise Profile

After the files are open, the window title will change to indicate the file name, total length, the sampling rate, and the bit size. The “Open” button will change to a “Start” button. Now comes the important part – setting all those sliders which determine the noise characteristics, and the amount of noise reduction that will be attempted.

For the purpose of this quick-start, I assume that your file has some of the lead-in track left intact.



The file which is used here can be down-loaded from the DeNoise web pages. It is part of a 1932 recording of 'Ted Weems' orchestra, and has already been de-clicked, de-crackled, and mixed down to mono. Having put on some good headphones, or organized some other way to monitor events, click the “Start” button. Once the music begins to come through, click the “Pause” button. The result, for the Ted Weems file, is shown above.



The next step is to sample the noise. In the second picture above you will see that I have selected a portion of the lead in track in the upper narrow window, where it is highlighted in light grey. Compare this picture with the previous one to see that three of the buttons, “Play”, “Preview” and “Sample Noise”, previously greyed out, are enabled as soon as some audio is selected in this way. Click the “Play” button, or press the space-bar, to play the selected sound, and listen carefully to check that it is noise, and that you have not accidentally selected the start of the music. Then click the “Sample Noise” button, and all of the sliders, with the exception of the “Limit Reduction” slider, will be automatically reset.

Do not select the “Automatic” option at this stage – read and digest the contents of Chapter 8 first.

## Using a separate Noise Sample File

If you have saved small snippets of noise samples as separate short files, then I recommend the following procedure:


- Open the sample file.
- Using the “Preferences” dialog, choose to pause at the end of processing.
- Process the file without any regard to the settings.
- Once DeNoise pauses, select the desired noise sample and click the “Sample Noise” button.
- Either name and save the settings, or click the “Resume” button and choose to keep the new settings. If you process the actual audio file next, those settings will remain.

## Processing

For a quick start, and to avoid possible over-processing, leave the reduction limit at the default value of 12dB. In any event, a 1932 recording with no audible hiss will sound quite unnatural, and almost certainly not as good as a restoration which employs a more conservative approach.

Now everything is ready to go, except that the beginning of the file was already processed with totally inappropriate settings! No problem – there are two possibilities:

- (i) Click the “Abort” button to get the following dialog:



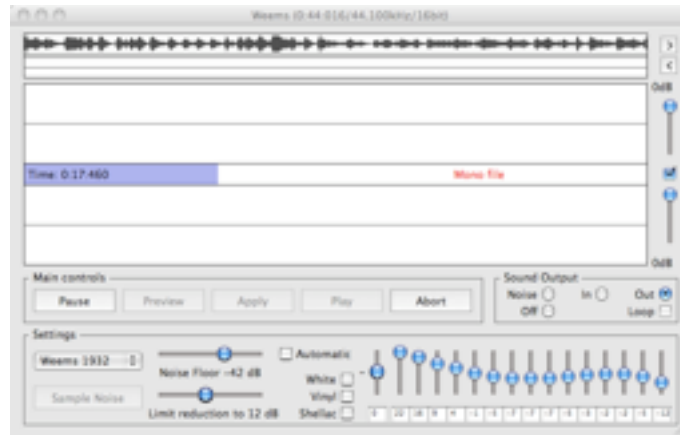
Now click “Start again”, which will take you back to a new start with the file already chosen.

You must click the “Start” button to actually commence processing again.

- (ii) You can re-process the audio that is still in the Sound Buffer, with these new settings, and continue on with those new settings. This is described later, in Chapter 7.


## Progress Indicator

Once everything is set up, click “Start” or “Resume”, and processing will continue unabated. A short segment of the file (the default is 524,288 samples – about 12 seconds) is resident in computer memory (RAM) at any given time; as new sound samples are read in, repaired samples are written out to make room for them. In this manual, this data cache is called the “sound buffer”. You can observe the state of this buffer in the upper display; there is also a progress indicator bar in the lower display. For more information on the sound buffer and its use, consult Chapter 7.



Apart from listening during processing, you can always pause, then select some or all of the audio in the sound buffer and press the space bar to play it. You can compare the incoming and outgoing audio by clicking the “In” and “Out” radio-buttons; you can also loop the playback using the “Loop” check-box. Should you also listen to the removed noise, do not be alarmed that it is far from uniform. DeNoise has a noise-reduction strategy which results in very little or even no noise removal where it is judged it will not be perceived by the human ear.

Notice also that in the picture here I have the sound turned on. Turning it off will result in quicker processing of the rest of the file. I have also saved the settings with the name “Weems 1932”.



At the end, but only after the repaired file is completely written out, a dialog box will appear. The program is then ready to process the next file. As with many of the dialogs, there is a “Don’t show again” check-box. I recommend that you do not turn off these dialogs until you are proficient in use of the program.



## 4 Producing and Organizing Files

ClickRepair users need only read the section about the recommended order of processing.

### File Requirements

Before you can repair audio files, they must first be captured. This is done by playing the record into a suitable sound editing program via a sound card or other analog to digital converter, to capture the data as a digital stream. There are a number of useful resources for learning about this, for example [www.delback.co.uk/lp-cdr.htm](http://www.delback.co.uk/lp-cdr.htm) and [www.a-reny.com/iexplorer/restoration.html](http://www.a-reny.com/iexplorer/restoration.html). Both give thorough coverage of the process, from connecting the sound equipment to the computer and capturing the recording to preparing and burning the CD. The files must be uncompressed mono or stereo, in either “AIFF” or “WAVE” format. Acceptable sample sizes are 16, 24, or 32-bit integer, (aiff, wav type 1) and normalized 32-bit float (aifc type fl32, wav type 3). Be aware, however, that internal processing is 25 bit integer in all cases. The output file will automatically be in the same format as the input file.

The signal from the turntable equipment itself is totally unsuitable for conversion to a normal digital CD. This is because records are produced with a heavy emphasis on the high frequencies, and de-emphasis on the low frequencies, and must be played back through a suitable equalizer. This equalizer de-emphasizes the high frequencies and re-emphasizes the low frequencies. This is the reason why a good quality vinyl LP has relatively little “hiss”; equally, it is the reason why “hum” and “rumble” can be difficult problems to deal with.

For shellac (78) records, I recommend some careful research; one place to start is Roger Wilmut’s page [www.rfwilmut.clara.net/](http://www.rfwilmut.clara.net/)

The most likely choice is to connect the turntable equipment to the phono input of your music system, and capture the output from auxiliary out sockets into the computer. In addition, there are other solutions for sound input, such as purchasing a dedicated preamplifier, or a USB sound card system.

If you already have some files which have been captured and burned to CD, see below.

### Recommended Order of Processing

After capturing the sound to file, you may want to perform some or all of the following:

- Remove subsonic and low-frequency noise with DeNoiseLF.
- Remove the clicks and crackle with ClickRepair.
- Reduce continuous noise, such as hiss, with DeNoise.
- Separate the tracks of the record into distinct sound files, so they can be navigated on the CD.
- Fade the beginning and end of each track before burning the CD.

General editing functions, such as audio capture, separating and fading tracks, etc., will require an audio editing program.


Some users have questioned the order recommended here, given that there are other commercially available packages which recommend general noise reduction as the first step. The reasons for the recommended order, when using the nominated applications, are simple:

- DeNoiseLF isolates and processes just the signal components below about 250Hz. This may be beneficial for the operation of ClickRepair.
- Clicks and crackle may be partially removed by DeNoise, leaving them audible, but harder to find.
- Clicks and crackle can significantly upset continuous noise reduction processes which operates across the entire frequency spectrum, including DeNoise.
- These last two points indicate why ClickRepair should be used before DeNoise.

One other word of warning. Before subjecting files to noise reduction with applications like DeNoise, listen to some historical recordings that have been professionally restored. You will find that some of the hiss and other noise is still there! This can be quite deliberate – so tread carefully!

## Notes about File Structure

There are many types of sound files. They are generally organized in “chunks”, with small “header chunks” indicating the actual structure and type. If you attempt to open a file which DeNoise is not designed to accept, a dialog will appear, giving relevant, if somewhat technical, information. An example of such a dialog is shown here. The most likely reasons for rejection of a file are:



- The file uses something other than 16, 24 or 32 bit (integer) sample type, or 32 bit normalized float. Sample types not supported include 8 bit integer, and float formats other than wav type 3.
- The file uses some form of compression. Note that WAVE files can use various compression formats, whereas AIFF files are uncompressed. Other file types (MP3, AAC, ...) are also unacceptable.


Uncompressed sound files are huge – about 10MB for each minute of CD quality data. It follows that most, but not all, of the file is the sound data itself. This is the only part of the file to which DeNoise will make any changes. In particular:

- All chunks which occur before the sound data are copied to the new file verbatim.
- All data which occurs after the sound data is copied to the new file verbatim.
- File name extensions are copied from the input file, so as not to change file associations.
- On Mac OS X, the file type and creator, which is used by the Finder in determining file associations, are also copied.

## File-lengths

DeNoise can only process AIFF files for which the sound data may be indexed using 32-bit integers (the largest 32-bit integer is 2,147,483,647). Because they have a different format, DeNoise can only process WAV files whose actual length is less than about 2GB (2,147,483,647 bytes). Attempts to open longer files will result in an error dialog. For stereo files at a resolution of 16/44.1kHz, a 2GB file is a little over 3 hours in either format.

DeNoise has one other dialog related to file problems, shown here. It is a warning that the file length given in the header disagrees with the actual (physical) file length on disk. This alert is provided for information only – in most cases it will not affect the use of the file, either by DeNoise, or by other audio editing programs.



In more detail:

- Provided that the file is at least as long as the sum of its required chunks, this error will not affect processing, and there will be no problem in proceeding.
- If the file is incomplete, this warning will be followed by the larger error alert above, with the number of missing bytes stated. Processing will not be possible in this case.
- DeNoise does not change these (or any other) headers, even if they are inconsistent.

## Sample Sizes

The sound data used for standard CDs is 16-bit, meaning that each sample is a whole number (integer) in the range -32768 to +32767. For 24-bit data, the range is -8388608 to +8388607. It is important, when processing CD sound data, to adjust the levels so that the loudest passages do not exceed the maximum level – saturation renders the audio virtually unusable. In DeNoise, audio level is measured relative to this critical saturation level, which is labeled 0dB. A signal which leaves

one data bit of headroom has a level of -6dB, a signal leaving two bits of headroom has a level of -12dB, etc.

Audio levels are not measured in numbers and fractions, rather in decibels (dB). This is a scale for measuring multiplicative increase and decrease in an additive way – a logarithmic scale. An audio level of -6dB corresponds to one-half the saturation level, -12dB corresponds to one-quarter, -18dB to one-eighth, etc.

Note that a signal whose amplitude is reduced by one-half has its power reduced by one-quarter – power is proportional to the square of the amplitude. This is why half power is -3dB while half amplitude is -6dB. In the DeNoise display, we are viewing the amplitude.

In considering what 16-bit audio quality means, the following comments may help. In a 16-bit number, the 16th bit is used to indicate whether the data is positive or negative. The lowest bit contains as much quantization error as signal, since 100.45 might be recorded as 100, 100.55 as 101, magnifying the difference. Therefore the maximum signal-to-noise ratio attainable on a standard CD is about 87 dB (14 1/2 bits), quite adequate for capturing a record, provided that the transfer is not made at too low a level. But, if the transfer is made at a level which never peaks (for example) above -12dB, 2 bits of the available digital data is unused, and 12dB of the available signal to noise capacity is discarded.

Internal processing uses 25-bit integer representation,<sup>3</sup> regardless of whether the data files are 16, 24, or 32-bit, integer or floating. Capturing and processing in 24-bit format give greater flexibility for editing the audio, even if the data is finally truncated to 16-bit for burning CDs; indeed this is the format widely used by professions for audio editing and production (usually at 96 kHz as well). Note, moreover, that the resolution and dynamic range of audio captured as 16-bit data is not improved by subsequent conversion to 24 or 32-bit formats – for this it is necessary to perform the analogue to digital conversion using hardware which provides 24-bit.

## Reading from CDs

Some users may have previously captured records and burnt them to CD, and now wish to re-process them. Obviously it would be convenient to read such files directly from the CD, without first extracting the tracks.

The sound data on a standard CD is 16 bit stereo, sampled at the rate of 44.1 kHz. The physical format of the data is such that it has to be “ripped” to recover a valid computer sound file. On Mac OS X, each track on a CD is presented, by the Finder, as a file of type “AIFC”, whose encoding is of type “sowt”. This is a standard file type, and since the operating system is doing the “ripping”, the data can be read directly by DeNoise. Of course, data cannot be written back to the CD, but since DeNoise will not edit the original file, it is possible to process directly from it.

---

<sup>3</sup> Effectively 24-bit audio – the extra bit is required by ClickRepair, which uses the same file handling code.

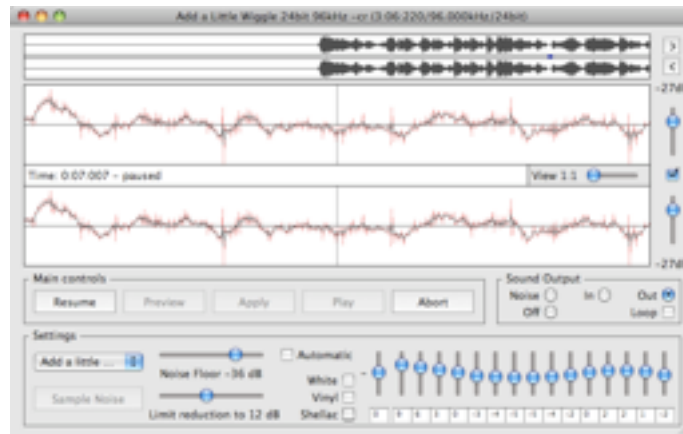
It has been explained already that DeNoise makes no changes to any part of an audio file, except for the sound data itself. Therefore, files processed directly from CD will be saved with the same headers and file extension as were passed in by operating system. You will be able to open these new files with any program which is able to deal with the original file directly.

If you want to test this, drag one of the tracks from the CD, to the desktop, using the Finder. This will copy it. Now attempt to open it with your Audio Editor or other application. If it works for the original, it should work for the file produced by DeNoise.

## 5 User Interface Basics

### Visual Interface

The main features of the user interface were already seen in Chapter 3. However, it is possible to view either the spectral information or the waveform, and I show a picture of the latter here. To get it I pressed the Alt/Option key. Every time processing stops, there is a dialog to indicate that the Alt/Option key functions this way; obviously you will wish to disable this dialog after a while.



The red trace is the incoming audio, the black trace the treated audio. In this example it is seen that some significant processing is being done, because this is a quiet patch in an old shellac recording, and it is clearly seen that the audio output is a smoothed version of the more jagged audio input.

The various elements of the display will now be described in some detail.

### The Sound Bar

The upper (smaller) view displays the contents of the sound buffer – by default about 12 seconds in length. It may be used to navigate around and also to select sound in the sound buffer for playback, reprocessing, and sampling. There is also a narrow central strip, in which a small blue indicator shows what segment of the sound buffer is being displayed in the lower view. The two buttons to the right of the sound bar scroll the samples displayed in the lower view, and this is also indicated by movement of the indicator.

I shall refer to this upper view as the “sound bar”, since it has some similarity with the horizontal scroll bar of a word or image processing program. The indicator shows how the samples displayed in the lower view relate to the totality of samples resident in the sound buffer. Moreover, clicking in the sound bar, or in the adjacent right and left shift buttons, causes the lower display to scroll. In the default situation where spectral information is being displayed, the arrow buttons and arrow keys always scroll by a single processing frame (1/80 second).

## The Display Window and Scale Controls

The lower view is used for examining either the spectral information (the default) or the waveform (an option). I call it the “display window”, although technically, it is not a window. The waveform display is linked with three scaling controls (two vertical sliders and one horizontal slider). The incoming noisy signal is shown in red, the outgoing treated signal is superposed in black. Unless the degree of noise reduction is significant, the red will hardly be visible.

If the program stops at a marker, or is stopped using the “Pause” button, the central strip displays that fact, together with the time, and the name of the marker, if appropriate. When the program is running automatically, this central strip serves as a progress indicator. Finally, when processing is complete, it displays some statistics.

To the right of the display window, there are two slider controls, with a check-box in between. The sliders control the vertical scaling employed when displaying the waveform. When the check-box is selected, vertical scaling is automatic. Moving the vertical sliders automatically de-selects the check-box. To magnify the display, move the slider down, which reduces the display peak level. This might seem counter-intuitive, but is a consequence of showing the peak level.

The waveform display is scaled automatically, usually the most convenient scheme. Peak levels are displayed in decibels (dB), just above and below the sliders. The sliders also move down as the peak level falls. Information about sample sizes and peak levels was given in the “Sample Sizes” section in Chapter 4. Generally, there is no need to worry about these controls.

The horizontal scale used in the display window is controlled by the “View” menu, and also by a small horizontal slider, visible in the central strip of the display window whenever the waveform is being displayed. You can use either for changing the scale.

The effect of increasing the ratio is that more samples are displayed across the display window, at the expense of having a less detailed view. The scales progress in powers of 2, starting with the most detailed “View at 1:1”, then “View at 2:1”, up to “View at 256:1”. The viewing scale also determines how many sound samples are displayed per pixel across the display window.

As the scale increases, the fraction of the sound buffer which is displayed increases. The indicator of the sound bar changes to show this, becoming progressively wider.

## Menus

On all platforms, there are six menus, with menu items as given here:

- File – Batch Setup, Breakpoints, File Dialogs, Open, Start, Pause, Abort.
- Edit – Preview, Commit, Sample Noise.
- Settings – [saved settings], Custom.

- Sound – Play/Stop, Off (faster), Removed Noise, In (Original), Out (Repaired), Loop, Mixer Settings.
- View – from 1:1 to 265:1.
- Help – Registration, Keyboard Shortcuts.

Note that there are no default settings on the Settings menu. Initially it is empty.

On Mac OS X, the Application menu has the items:

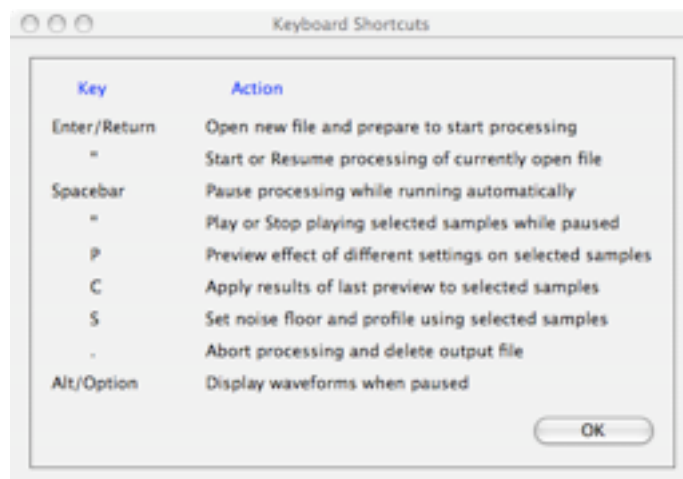
- Quit DeNoise.
- About DeNoise.
- Preferences.

On other platforms, these three appear as extra items on the File and Help menus.

Even if you are using the in-window controls (buttons, etc.) and/or keyboard shortcuts most of the time, the menus serve to provide extra information where there is insufficient room for this on screen. For example, the Sound menu items contain more explanation of the four possible choices.

## Keyboard Shortcuts

A complete list of keyboard shortcuts (not to be confused with the usual menu-shortcut keys associated with the command/control key) may be viewed by selecting this item on the “Help” menu.



The resulting window is shown here. It is not a modal dialog and may be left open during normal operations, for ready reference, although it can be rather obtrusive.



## Audio Monitoring

Radio buttons in the “Sound Output” control group allow you to turn sound output on or off. Monitoring, particularly with a pair of headphones, or reasonable speakers, can provide invaluable feedback. These buttons may also be controlled from the menus.

Audio monitoring causes processing to be real time, which will be much slower for computers with even moderate processing power. Processing speed does not depend on the amount of noise.

- Selecting “Out”, the default, the repaired file will be played as it is processed, which is some time after it is read into the sound buffer from file.
- Selecting “In” will send the unprocessed data to the audio system, which can be useful for “A–B” comparison, particularly when used with the “Play” button (detailed instructions are given later).
- Selecting “Noise” will enable you to monitor what is being removed by the program.

It is inevitable that the removed noise will have some correlation with the audio itself. While you do not want to be dulling the audio, remember that the strategy of avoiding the removal of noise which is masked by the audio means that some of this correlation is actually desirable.

- The sound can be turned off, removing the restriction to real-time processing. This is the fastest possible mode, and the most convenient provided you are confident of the result.

## Sound Latency

Given that the computational delays will depend critically on available computing power, the processed data is not sent to the sound system directly, rather to a second buffer, which in turn feeds the sound system. If this second buffer empties, either because other programs/services consume a large fraction of available resources, or because the CPU is not sufficiently powerful, then the sound will be broken up, as if the computer were stuttering.

The algorithms use techniques from advanced digital signal processing theory, and are quite taxing on processing power. I have attempted to make the coding as efficient as possible, but the algorithms have been chosen for their effectiveness more than their speed. Turning the sound off may be the only viable setting with slower processors, in which case it is still possible to play the contents of the buffer from time to time, because the repairs have already been made.

The upshot of all this latency is that by the time you hear something, it will already be about 1 second in the past before you can click “Pause”, and the program can respond. Since the (main) sound buffer – the buffer displayed in the sound bar – is longer than this, the audio is still accessible for review and for adjusting the repair settings. In order to review it, you will therefore have to navigate back, as described later. Reprocessing audio in the sound buffer involves using the “Preview” and “Commit” buttons, and will be described in Chapter 7.

It would be more convenient to go automatically to the exact point at which the problem was actually heard, in response to clicking the “Pause” button. This would involve much tighter synchronization of the processing loop with the user interface, resulting in a significant diversion of processor power to this task.


## 6 Using DeNoiseLF

DeNoiseLF is somewhat easier to use than DeNoise; there are less controls and they are set by careful experimentation. File organization and user interface are the same as for DeNoise.

There are two modes, governed by the selection of the pair of radio buttons marked “Hum” and “General” in the “Settings” group. In addition, the Hum mode has four associated check-boxes for simultaneously removing the first four harmonics of the hum – see below for details.

I have provided three presets as starting points for experiment.

- “Default Rumble”, which applies noise reduction up to about 40Hz.
- “Default Vinyl Noise” is a variation with a higher frequency limit of 150Hz.
- “Default Hum” – it switches the radio buttons to “Hum” at 50Hz, without harmonics.



Frequencies in the range 20-220Hz may be selected using the slider, the label will read “Limit” (for general noise) or “Peak” (for hum).

The “Noise Floor” settings, one for each channel, are in units related to the model used.<sup>4</sup> The check box indicates whether or not the left- and right-channel sliders are linked together. Moving the left slider will move both; moving the right slider will not affect the left, rather the check-box will be de-selected. Select the check-box again to re-link the controls, which will automatically reset the right-channel slider to match the left one.

### Rumble Reduction

When the application is operating for the removal of rumble or other broadband low frequency noise, processing is in two consecutive steps:


- A fixed, sharp, cut-off filter is applied to the incoming signal, with a cut-off frequency of 20Hz. Fixed means that all subsonic audio is removed, regardless of signal level.

---

<sup>4</sup> Technically, to an excitation; this is not a direct measure of the actual noise amplitude although it is related.

- The audio stream emerging from this filter is then processed by a dynamic algorithm which attempts to estimate the dominant low-frequency noise at a level set by the “Noise Floor” sliders.
- At higher signal levels, the effect of the dynamic process on the audio will not be perceptible; at low signal levels the objective is to reduce the perception of undesirable noise.

As with Chapter 3, I have provided on the DeNoiseLF web pages the sample file, “LF-sample.aif”, that I used in writing this chapter. The main problem with this sample is 120Hz hum at a quite intrusive level, but there is also some low-frequency and subsonic noise with might be described as rumble. Because the hum is so prominent, it is best treated after this subsonic noise has been reduced, although that is not completely necessary. Nevertheless, that is the approach I took, choosing to process the file twice, first for the reduction of general low-frequency noise, then subsequently for the removal of the main problem, the hum.




After a couple of trial runs, in which I auditioned different noise floor settings, I decided that the first step was adequately achieved with the noise floors at -50dB but with the limiting frequency increased to 60Hz. The picture above shows the result of this processing on the lead-in track: I clicked the “Pause” button at about 3 seconds into the file, before the audio begins. To the right of the vertical line – marking the pause position – is the untreated incoming audio. To the left is its decomposition into outgoing audio (black) and removed noise (red). Notice that the scale is set to 64:1 – the removed component, although at a relatively high level because there is no significant audio content, is at a very low-frequency, there being well over 1000 samples between peaks.

I mentioned in Chapter 1 that DeNoiseLF does not use a spectrally-based processing method. Nevertheless, the spectrum, in the range 20-250Hz, may be viewed by depressing the Alt/Option key, and this is a valuable diagnostic tool when seeking appropriate settings. The picture below shows the result of depressing this key.

- The light grey trace is the incoming, untreated signal. It is mainly obscured by the blue trace, indicating that the incoming and outgoing signals are the same.
- The blue trace is the outgoing, treated, signal. It is only lower than the blue trace in a frequency band extending to about 60Hz.
- The hum at 120Hz is clearly seen, and is not being reduced by the present process.


- In fact, if the waveform is viewed at a scale of 16:1 rather than 64:1, it is clear that most of the outgoing signal in the lead-in track is just the hum.
- Remember that this spectrum is not localized, about 1 second of audio is used centered on the selected place.



## Hum Reduction

Having allowed processing of this file to complete, thus reducing rumble and sub-sonic noise, I processed the new file (LFSample-dr) to remove the hum. The first step is to determine its frequency; from the spectrum it is clearly at 120Hz. This is quite difficult to deal with, since there is likely to be desirable audio near to this frequency from time to time. To treat the problem, select the “Hum” radio button and move the “Peak” slider to 120Hz.

It is important to set this frequency accurately when attempting hum reduction.




It remains to find noise floor levels which are high enough to reduce the hum sufficiently, but no higher than necessary. This is best done by processing the first few seconds of the file, pausing the process, then selecting a few seconds of the audio in the sound bar for audition. At this juncture the best way to proceed is to select the “Loop” check-box and use the “Preview” button. While the loop is playing, the noise floor sliders can then be adjusted a little at a time, listening to both the output and the removed noise. The final settings which are selected will be the result of personal

preference, a balancing act between effective noise suppression and minimal degradation of the audio. With the settings I finally chose there was no remaining audible hum or visible peak, although I tried to minimize the damage being done to the audio itself. Notice that I set the right channel noise level 5dB higher than the left channel. Remember also that DeNoiseLF is not a notch filter, as is evidenced by the spectra shown in the picture above – not all the content has been removed at the selected frequency of 120Hz. How much audio is removed is determined by the Noise Floor controls, and it can be very little indeed. Setting the level too high, however, can result in loss of original audio.

## Removing Harmonics

Harmonics up to the fifth may be removed simultaneously, up to a maximum frequency of about 320Hz. For the sample file above, the second harmonic would be at 240Hz, and since there was no trace of that in the spectra, there was no need to worry about it.



A more common situation is that the hum is related to an AC power supply, at 50 or 60Hz (depending on the country). The harmonics of 60Hz, up to the 5th, are at 120Hz, 180Hz, 240Hz, and 300Hz. If any or all of these are evident in the spectrum, either in a quiet passage or in the lead-in track, then they may be suppressed by selecting the appropriate check-boxes. As the frequency slider is increased, some of these check-boxes may be grayed out to indicate that these harmonics are at too high a frequency for treatment.

The picture above shows the treatment of some very annoying “buzz”, at a fundamental of 50Hz, but including odd harmonics at 150Hz and 250Hz. In fact, even higher harmonics were present; these were reduced to an acceptable level by subsequent processing with the DeNoise application. Processing in the other order (DeNoise before DeNoiseLF) is not recommended.

## 7 Audio Output and the Sound Buffer

### Java Sound Output settings

Audio monitoring involves sending sound data to a sound data line which has been made available to Java applications. The number and functionality of such lines is determined jointly by the operating system and the Java run-time, and is outside the control of Java applications. Sometimes there is a choice of output lines; also these lines often have an associated “Master Gain” control to set the overall sound output level.

These features may be accessed from the “Mixer Settings” item on the “Sound” menu. There is a drop-down list, which will list all available output mixers, together with a “Master Gain” slider control, and a button marked “Play Test Tone”. Clicking on this button causes a short pure tone, at A440Hz and maximum allowable amplitude, to be sent to the selected line. If the gain is set too high, the tone will sound harsh, due to clipping. If it is set too low, the sound output from DeNoise will be at too low a level.

I have chosen the default setting of +5dB, for the default Java Sound Audio Engine, to match what might be expected from typical applications such as Apple QuickTime or Windows Media player.

To change the settings, simply choose something that is suitable, then click the “OK” button.

### The Sound Bar

The sound bar shows all the audio data which is still resident in the sound buffer. The default length of this buffer is 524288 samples, for files up to 48kHz – about 12 seconds in practice. This can be increased by a factor of 2x or 4x, to 24 or 48 seconds, using the “Preferences” dialog.

For files above 48kHz, the longest buffer setting may result in “Out of Memory” messages, since all workspaces get doubled or quadrupled in length – not just the sound buffer. Sound output, moreover, always remains at the lower sample rate because of limitations with Java sound output.

To review any part of the sound buffer, simply click on the desired spot, and the display window will move to that place to within the closest frame (1/80 second). If the point clicked is not exactly where you want to be, then use the right or left arrow keys to scroll the selection showing in the display.

First click “Pause” – the sound bar does not respond in automatic mode.

To make a selection of audio samples in the sound buffer, click and drag in the sound bar from one end to the other of the selection; or double click to select all. The selected samples will be highlighted in light gray. Once a selection is made, the “Play” and “Preview” buttons will be enabled.

A single click in the soundbar makes the display window scroll, for which reason a single click, followed by a shift-click, is not used as a selection mechanism for sound playback. This differs from a common behaviour pattern of many programs, for which a click does not cause a scroll, whereas a shift-click does extend a selection, but it avoids ambiguity in the meaning of a single click in my applications. You will need to click, hold, and drag to make a selection. Note also that if a double-click is too slow, it will be treated as two single clicks.

## Appearance of the Display

Whenever the sound bar is clicked to navigate to a given position, the display changes accordingly.

- For DeNoise, the standard form of display is the spectral information, since this is of most use in experimenting with settings. Scrolling, for example by using the right- and left-arrow keys, moves forward or backward by one processing frame (1/80 second).
- Depressing the Alt/Option key will switch to a view of the waveform. Scrolling still works while the key remains depressed.
- For DeNoiseLF, the standard form of display is the the waveform, with the outgoing signal painted in black, and the removed component in red. When the “Pause” button is first clicked, this is seen quite clearly since the vertical line shows the point at which this decomposition is occurring. This was seen quite clearly in pictures shown in the previous chapter.
- Depressing the Alt/Option key will switch to a view of the spectral information. Scrolling still works while the key remains depressed, again making this a useful view for judging the effect of the current settings.

Note that these viewing modes remain the same when using the Preview functions described below.

## Listening to the Sound Buffer

Once processing has paused and a selection made, as indicated by the grey section on the soundbar, click the Play button to play that selection. There is a small blue indicator, in the central strip of the sound bar, that will track the sound as it is played. During this time, the button will display “Stop”. Notice that, if sound output is off, the processed sound is played as the default option. If you want to play back the original sound or the removed noise, select the appropriate radio buttons.

The “Loop” check-box results in an endless loop, until either you choose to stop, or to resume processing. You may change the radio buttons while listening – but don’t forget that there is a lag before it affects the sound being heard.

The shortcut for both “Play” and “Stop” is the spacebar. There are also menu shortcuts for switching between the three audio-out options (In, Out, Noise), and for switching the “Loop” option.

Having found a desired selection, it may be used as a noise sample (if appropriate) or reprocessed using the “Preview” and “Commit” buttons, as described next.



## Experimenting using Preview

The effect of using different settings on the selected samples may be auditioned using the “Preview” button. Simply change one or more settings (for example, the “Limit Reduction” slider) then click “Preview”. While the preview is in progress, the button will change its text to “Stop”.

- If “Loop” is not selected, the selection will be played just once, with the new settings applied. When this finishes, the “Apply” button will be enabled.
- If “Loop” is selected, different settings may be previewed in an endless loop. Since processing never finishes, the “Apply” button will never be enabled.
- The only way to enable the “Apply” button in this case is to turn off “Loop” and wait for the preview to complete.
- Similarly, if the preview is stopped using the “Stop” button, the “Apply” button will not be enabled, because the process was not allowed to complete.
- The only way to enable the “Apply” button in this case is to click “Preview” again and wait for the process to complete.

## Committing a Preview

Once a preview is complete it may be applied or discarded.

- To discard it, simply de-select the samples in the sound bar, or just click the “Resume” button.
- If the selection is only part of the the sound buffer, the “Apply” button will commit that selection, using a fade-in and fade-out of the removed noise.<sup>5</sup>
- If the selection extends to the end of the sound buffer, the “Apply” button will commit those settings without a cross-fade at the end, and the new settings will continue to apply if processing is resumed.
- There is no “Undo” facility once “Apply” has been clicked, but you can preview and reapply with the previous settings while the affected samples remain in the sound buffer.

## The Last Buffer

When the last samples have been processed, there is still a full sound buffer that may not have been auditioned, and which has not been committed to file. An opportunity for review may be set using the preferences: there can be a one second warning beep and/or an indefinite pause before finalizing the file.

---

<sup>5</sup> The underlying audio is, of course, not subject to this cross-fade; only the removed noise.

## 8 Settings – DeNoise

The “Settings” group contains controls with the following functions:

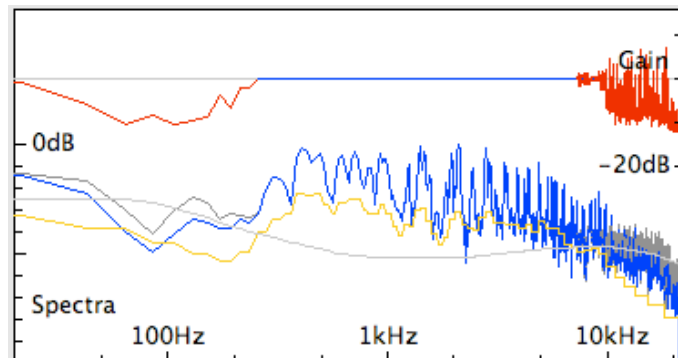
- Two horizontal sliders which set the overall noise level and the reduction limit.
- 15 vertical sliders which set the profile of the noise across the frequency spectrum.
- 15 text-entry fields linked to the noise profile sliders.
- A vertical slider and text-entry field for setting the L-R balance of the noise floor.
- Three check-boxes which may be used to set the noise profile to one of three presets.
- A check-box next to the “Noise Floor” slider for automatic noise tracking.
- The “Sample Noise” button which is used for automatically setting all of the above controls except for the “Limit reduction” slider.
- A drop-down list which provides the convenience of saving and recalling useful settings.

Although this is a very large number of controls, normally a noise profile will be set automatically, either via automatic analysis of a noise sample, or by selecting one of the three check-boxes to set a default profile – see later for a description. The results depend critically on a good choice of settings, and achieving this requires some understanding of what DeNoise is doing and what the available information means.

The centre frequency of each of the 15 bands controlled by the vertical sliders will be displayed if you hover over the slider with the mouse. There is also a tool-tip for the L-R balance control, as a reminder of its purpose.

### Spectral Audio Information

As mentioned in the introduction, DeNoise works in the frequency domain, dividing the incoming audio stream into 80 overlapping frames per second, and transforming each frame into a large number of equally spaced frequency bands. When the application is paused, it is possible to view this information for a selected single frame. Here is such a display from the file used as an example in the Chapter 3 (Ted Weems Orchestra). I selected a frame which is in the middle of some high-pitched brass, and the many sharp harmonic peaks are evident.



The bottom four traces contain information about the audio spectral analysis and the spectral distribution of the noise, with the frequencies running to the theoretical maximum frequency, which is half of the sampling rate (half of 44.1 kHz in this example). The horizontal axis marks off the frequency on a logarithmic scale, a common choice for this type of display. The vertical scale of these traces is marked on the left-hand side of the box, starting from 0dB at the top, which is the maximum signal. Each tick down the vertical represents a 10dB drop. The four traces are:

- The grey trace is the frequency content of the incoming, untreated, audio.
- The blue trace is the frequency content of the outgoing, noise-reduced, audio. Where the two are identical – from about 250-8000Hz for this example, the blue obscures the grey.
- There is a light grey trace, this time a smooth curve, which lies mainly below the audio from about 200-10,000 Hz. This is the noise level which is being used in the noise-reduction algorithm.
- There is an orange curve, sometimes above the noise level, but below it at low and high frequencies. This is the estimate of the amount of noise whose perception is masked by the audio.

Notice that at frequencies below about 250Hz, the discrete frequency bands are so far separated that the incoming and outgoing audio spectra are clearly interpolated by straight-line segments. It is the coarseness of the equally spaced frequency bands, on the more natural logarithmic frequency scale, which limits the utility of this method of noise reduction at lower frequencies.

## Spectral Gain Information

It is vital to realise that there is no way to precisely separate signal from noise. That would be akin to trying to remove the milk from your coffee. What is actually done is to pass the audio through a 1024-band equalizer (2048 bands if the sample rate exceeds 48 kHz), adjusting the gain of each of the equalizer bands 80 times per second, in response to analysis of the noisy audio data. The gain – mostly but not wholly negative – is displayed by the single curve at the top of the display, with the decibel labeling on the right. The gain curve is superimposed over a light grey reference to indicate 0dB.


- That part of the curve for which the gain is within  $\pm 0.5$ dB is displayed in blue. This is the component of the signal which is not adjusted, whether or not there is judged to be any noise.
- That part of the curve for which the gain is adjusted by more than  $\pm 0.5$ dB limit is shown in red.

Notice, in the pictures being displayed here, that not all of the equalization is attenuation. Sometimes a positive gain is applied, corresponding to a statistical estimate which indicates that the noise has acted to cancel the audio in that band.


In the picture shown above, it is interesting to note that the blue line correlates quite well with that part of the spectrum for which the orange curve is above the noise estimate. This accords well with the maxim “If it’s not broke, don’t fix it”. DeNoise applies this principle by not attempting to treat noise if that seems unnecessary. This uses some of the science of psycho-acoustics, which has been so successful in permitting lossy audio codecs such as mp3 and aac (mp4).

## Adjusting the Noise Floor

The noise floor can be adjusted, as a whole, using the “Noise Floor” slider. As a simple example, I increased the noise floor by 12dB, from -45dB, which was set automatically from the noise sample, to -33dB. If the picture shown below is compared with the one above, it is quite clear that a lot more of the incoming audio is now condemned as being at or below the noise level. At this point no new processing has been performed; all that has happened is that the smooth grey noise curve has moved up 12dB, and that is the only difference between the two pictures.



If applied, this change to the noise floor has a considerable, and deleterious, effect on the analysis and processing. To show this, I reprocessed this audio with the new setting, and then went back to view exactly the same frame. Here is the result:



The effect is that the band within which no processing is applied has narrowed, and a lot more of the audio content from 5 kHz to 10 kHz has been attenuated. This will dull the audio without serving any useful purpose, so one has to be careful not to over-estimate the noise.<sup>6</sup>

## Achieving a Compromise

The previous section has warned about the effects of over-estimating the noise level. The obvious answer might seem to be to err on the side of under-estimating it, using the “Noise Floor” control

---

<sup>6</sup> The only reason the result is not far worse is that the application of psycho-acoustical principles has saved the day to a considerable extent. The display does not indicate what the result might otherwise have been.

in the other direction. In this section I shall try to explain that this can be just as bad, possibly even worse, than over-estimating the noise level. What is needed is a good compromise.

The bane of the type of methods being used here is the introduction of distortion that gets heard as various kinds of musical noise. Typically this can take the form of random bursts of very short tones at random frequencies, sometimes called “bird song”. Other forms of distortion can also be perceived in some cases. Basically, what is happening is that most of the noise is being suppressed most of the time, but because the noise floor has been under-estimated, random noise bursts remain, but now localized both in time and in frequency. The effect can be even more objectionable than the original noise.

The way to judge this is by careful listening while applying the maximum possible reduction. If you have downloaded it, I suggest you open the “Ted Weems” file once again, start processing, then pause after about 10-15 seconds. Otherwise use one of your own files. Now take the following actions:

- Select the “Loop” check-box.
- Select the “Out” sound radio button.
- Select several seconds of sound in the sound bar.
- Click the “Preview” button.

The “Preview” button differs from the “Play” button in the fact that it processes the selected audio anew and sends the result to the audio output.<sup>7</sup> Of course, if the settings are not changed, this makes no difference. Now, however, while the loop is playing, take the following actions:

- Move the “Limit reduction” slider to maximum, so as to ensure that musical noise, distortion, and untreated noise in general, is easy to hear above the music.
- Decrease the “Noise Floor” by 6dB from the level which was found from sampling the lead-in track, so that the noise is under-estimated.

You will notice immediately that most of the hiss returns at this setting. (If it doesn’t, you probably increased the noise floor. Remember that it is negative; from -39dB to -45dB is a decrease.)

- Now, on each successive loop, increase the noise floor by just 1dB, and listen to the effect on the entire segment.

Each increase will decrease the perceived hiss, until it virtually disappears when you get back to the original setting. But, this improvement in noise reduction may be accompanied by some less than pleasant distortion if the noise floor is just a little too low. Once the noise floor has been increased to the point where no further noise reduction is achieved, further increase simply dulls the audio. The automatic noise sampling algorithm attempts to find this compromise automatically.

---

<sup>7</sup> This is fully described in the previous chapter.

## Using the Automatic Facility

This is a major new feature in version 2 of DeNoise. It is designed to track the noise floor insofar as the noise extends into the high frequencies, so it is most suited for the reduction of hiss.

When this check-box is selected, the noise profile remains fixed, except for the balance slider. As each frame of audio is treated, the 15 noise bands are also analyzed to assess how “noise-like” they are. The noise floor and balance sliders are then adjusted, automatically and dynamically, which affects the absolute noise profiles (the smooth grey curve in the spectral display). In deciding to use this facility, the following points should be taken into account.

- Automatic noise floor tracking does not attempt to adjust the noise profile.
- Automatic noise floor tracking does depend on setting a reasonable noise profile before starting.
- There are no known algorithms for deciding, with absolute certainty, between noise and music.
- If high-frequency musical content is erroneously classified, the noise floor can be over-estimated.

My recommendations are:

- Do not use automatic noise-floor tracking if the noise levels are substantially constant for the audio file being treated. The built-in psycho-acoustic algorithms are already quite adaptive.
- Conversely, try using this feature when the noise increases (or decreases) significantly throughout the file.<sup>8</sup>

## Adjusting the Reduction Limit

The more noise reduction is attempted, the more likelihood there is of degrading what you are trying to restore. So, once the noise floor is set to the best compromise – most easily achieved in the case that you have a good noise sample to use – continue to experiment using the “Preview” button while looping the audio, only now adjust the “Limit reduction” slider by 2dB at a time. The setting of 12dB which is set as a default, and which I suggested for the Weems track, seemed to me to be a good compromise between reducing the hiss to a reasonable level while not degrading the audio content nor introducing objectionable processing artifacts. For LP (vinyl) records something less (~6dB) will be more appropriate in most cases. Ultimately, the choice is always a personal one.

## Preset Noise Profiles

The simplest assumption for the distribution of noise is that it is “white”. This means that it is a purely random addition to the original audio: to each audio sample a random noise sample has been added, drawn from the same statistical distribution and with no correlation from one noise sample to the next. The resulting frequency distribution of the noise is uniform.

---

<sup>8</sup> This is common with 78rpm records as the linear tracking speed changes greatly across the disc.

Much of the noise on a vinyl record is also modeled quite well by this. But this noise, having been picked up by the cartridge, is then subject to the same RIAA equalization as the audio itself, and will have a distribution looking like the RIAA equalization curve – the bass is accentuated and the treble attenuated. Indeed, this is the purpose of RIAA equalization.

Similar considerations apply to shellac recordings, although the connection is not quite so strong.

For the convenience of users, I have therefore supplied three preset noise profiles which may be selected via three check-boxes:

- White – most likely used for treatment of tape-based material if no noise sample is available.
- Vinyl – for use with LP records if noise sampling is not possible.
- Shellac – I have added a bias to this profile so as to favour preservation of the mid-range. It really is much better, with 78 records, to get profiles from individual noise sampling, or by observing the spectral information in very quiet passages.

### **Noise Floor Balance**

The noise profile is always the same for both channels. The balance control, however, permits the noise floor settings to differ by up to 24dB between the left and right channels. A positive setting means that the noise floor is higher for the left channel compared to the right channel; conversely with a negative setting the noise floor is lower for the left channel than for the right channel. This difference is always applied as a split: for example if the noise floor is set to -45dB and the balance is set to +9dB, then the left-channel noise floor will be -40.5dB (4.5dB higher) and the right-channel noise floor will be -49.5dB. The difference is, of course,  $(-40.5) - (-49.5) = +9\text{dB}$ .

When the noise profile is set using the “Sample Noise” button, on a stereo file, the balance control is also set automatically. When making manual adjustment, however, the same principles apply as were already discussed under the heading “achieving a compromise”.

With mono files this control has no effect.

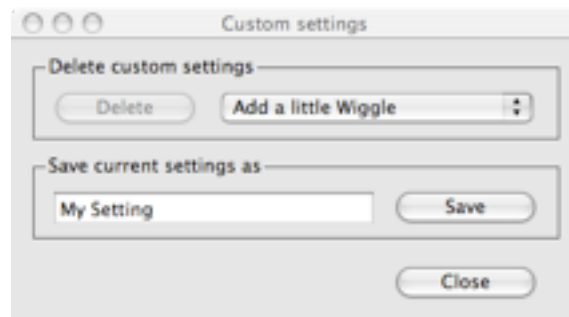
### **Saving and Recalling Settings**

Settings may be saved with any chosen name (up to 32 characters in length), and recalled at a later time, using the drop-down list in the “Settings” control group or in the “Settings” menu. I call either the “Settings Menu” in the following. All settings are remembered from one session to the next.

Note that the settings menu changes from the currently named settings to “---” as soon as manual adjustment are made to any of the settings. This is an indication that some settings have changed.

The settings menu comes with no presets whatsoever. If you have settings which you would like to add, select “Custom...” which is initially the only menu item. A dialog will appear, shown below.

Simply type in the name for your settings, then click “Save”. The current settings will then be saved with the chosen name. The Settings dialog may be used to manage the list, as follows:



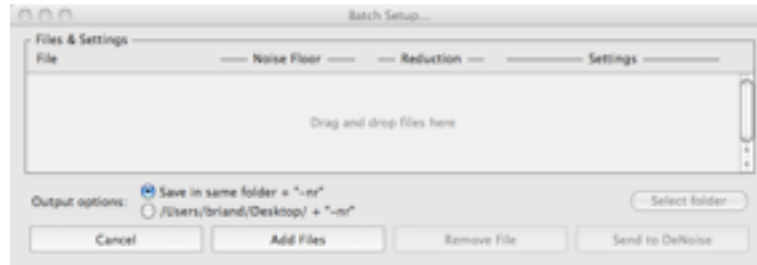
- To add the current set, type in the name and click “Save”.
- To change a set, type in a name which already exists, click “Save”, and confirm in the ensuing dialog that you want to update the settings of that name.
- To remove a set, select the unwanted settings in the drop-down list, click the “Delete” button, and confirm in the ensuing dialog.

It is important to note that the Settings dialog does not offer any possibility for adjusting the settings – its sole purpose is to save settings which are actually in use, and to manage that list. The only way to change an existing named setting is to save the new values with the same name. After the confirmation dialog, the old values will be replaced by the new ones.



## 9 Batch Processing

To prepare to process a batch of files, choose “Batch Setup” from the “File” menu. This will bring up the dialog shown here.



Note that this dialog grows in height as more than 4 files are added, however, the drag and drop area will always be sufficient to display 4 files. The maximum number of files which can be added is 256, of which a maximum of 16 will be visible at any one time. This is controlled by the vertical scrollbar, which will change appearance once there are more than 16 files.

The display will be slow to update if a lot of files are selected at one time, since each file has to be checked, and a new skeleton file created.

### Organization

The first thing to decide is how to organize the output files in relation to the input files. This is controlled by two radio-buttons and a “Select folder” button. The first option is for each output file to be saved in the same folder as the corresponding input file, and named by adding “-nr” to the name of the input file (or another string if this has been changed in the “Preferences” dialog). The other radio-button allows the output file to be saved to a selected folder, the pathway to which is displayed alongside the radio-button.

In this mode, the name of the new file will be extended only if this is the selected preference in the “Preferences” dialog, whereas the file name is always extended when saving to the same folder.

When the second option is selected, the “Select folder” button will be enabled. It may be used to change the output folder if the currently displayed choice is not the desired destination. Only folders will be displayed in this dialog, since we are not yet to the point of choosing files. Navigate to, and select, the desired output folder, and click “Select”. All these settings are saved from one session to the next.

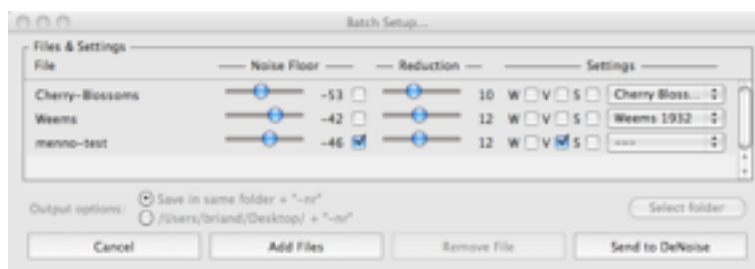
Note that the file dialogs used for batch setup are not the standard file dialogs provided by the operating system, rather they are Java “Swing” components, and have somewhat different behaviour.

The choice of destination may not be changed once files have been added to the batch.

## Assembling a Batch

Once you are ready to add files, click the “Add Files” button in the “Batch Setup” dialog, or drag and drop files into the dialog window in the highlighted area.

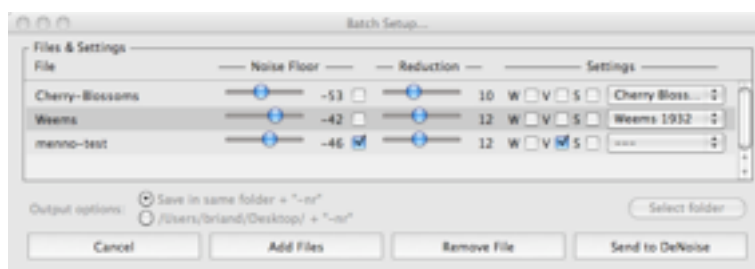
Clicking the “Add Files” button will bring up another file dialog, this time permitting only files to be selected. Select one or more files, and then click the “Select” button. The files will be added to the dialog, as shown here. To extend the selection in the file dialog, use the shift key. Random multiple selection may be made, using the “Apple” key on Macintosh, the “Control” key on Windows.



If you choose files which cannot be processed by DeNoise, then an error dialog will appear for each such file, in turn. Also, if you try to organize the files so that an output file will over-write its input file, this will generate an error dialog, and the file will not be added.

Note that there will be no confirmation dialogs asking permission to overwrite existing files. It is possible to lose existing files if you choose to replace them in this way when batch processing. Note also that the new files are created as soon as the files are added to the “Batch Setup” dialog. Moving these files before processing is complete will cause unpredictable errors.

Individual noise reduction settings are provided for each file in the batch. Controls for these settings appear next to each file name. There is insufficient room to duplicate all the controls in the DeNoise window. Note that the drop-down list has access to, and displays, any custom settings which you have saved, and the check-boxes give access to “Automatic” mode, in addition to the default noise profiles “White”, “Vinyl”, and “Shellac”.



To remove files from the batch, simply click on the name of a file in the “Batch Setup” dialog. The corresponding line will be highlighted, as shown here, and the “Remove File” button will be en-

abled, to allow removal. For removal, the shift key may be used to extend a selection; multiple random selection is not supported.

## **Processing**

When you are completely satisfied with the batch settings, click “Send to DeNoise”, which will dismiss the dialog and cause the “Open” button to change to “Start”. The previous setting will be saved. If you do not alter the setting in the main window, during processing of the batch, they will be restored on completion.

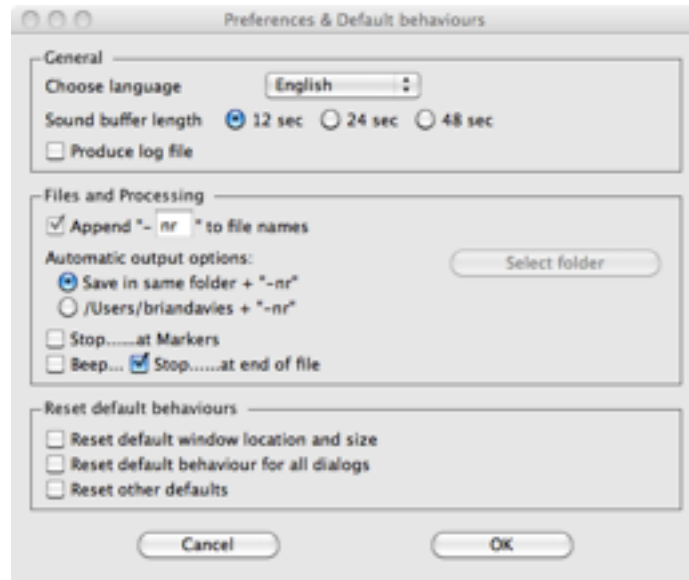
Processing will not start automatically, you must click the “Start” button to commence. Before you do, you might like to review the “Sound output” settings.

If you click the “Abort” button during batch processing, the action will apply only to the file currently being processed. If the action is confirmed, the window title will change to show the next file in the batch, the “Pause” button will change back to “Start”, and processing will be suspended. You must click “Start” to continue. You may also choose to abort files before processing begins, until the list of files is exhausted, and in this way abort all or part of a batch.

## 10 The Preferences Dialog

### Preferences

Using the “Preferences” menu item (Application menu on Mac OS X, Info menu on Windows) brings up the dialog shown here.



### Language

DeNoise has a choice of three languages in the user interface – English, French, and German. If your computer is set to French or German as language of choice, the program should automatically use that language from the first time it is used. To change the language in use, use the drop-down list in the preferences dialog.

The French translations were kindly provided by Michel Poulain; the German translations by Willy Kunz. Both are long-time users of the application.

### Sound Buffer Length

Versions of DeNoise and DeNoiseLF previous to 2.6 had a set buffer length of 524,288 samples – approximately 12 seconds at 44.1kHz. In the new versions, the choice of a longer buffer is intended to make it more convenient to experiments with, and re-process, selected portions of the file while they are still resident in computer memory (RAM). The extra memory requirements are significant, and choosing the longest buffer length (2097152 samples) may cause “Out of Memory” dialogs on some systems. Because the necessary memory allocations are made as the file is opened, it is only possible to change these settings while no file is open.

## Log Files

Logging may be turned on at any time, using the preferences dialog. At the start of each logged session, a file dialog will appear to allow the log file name and location to be chosen. Clicking the “Cancel” button in this dialog will turn logging off. It may also be turned off at any time using the preferences dialog. The log files are plain text for easy use. They record the file names, settings, and times taken.

## Files and Processing

The “Append -nr to File Names” check-box determines the default behaviour for how the output file is named, and where it is saved, in relation to the input file. The default appended string “nr” may be changed by typing a new choice in its place. Remember:

- If the same name is used for both input and output files, they cannot be in the same folder.
- The default locations for each are the folders last used for input and output.
- When batch processing, the save options are controlled directly from the “Batch Setup” dialog.

The “Save in same folder + nr” checkbox determines the folder to which the new file is saved by default, that is, the way the save dialog is initially displayed. If this check-box is selected, the dialog will be set to save the output file in the same folder as the input file, with “-nr” appended. If this check-box is not selected, then the dialog will be set to the same folder as the last output file, with or without an appendage.

The “Beep” and “Pause on completion of processing” check-boxes allows you to enable or disable the warning beep at the completion of processing, and to pause before writing out the final buffer to file, if required. If you are listening to the audio, in real time, through a reasonable sound system, I recommend turning off the beep.

## Breakpoints and Markers

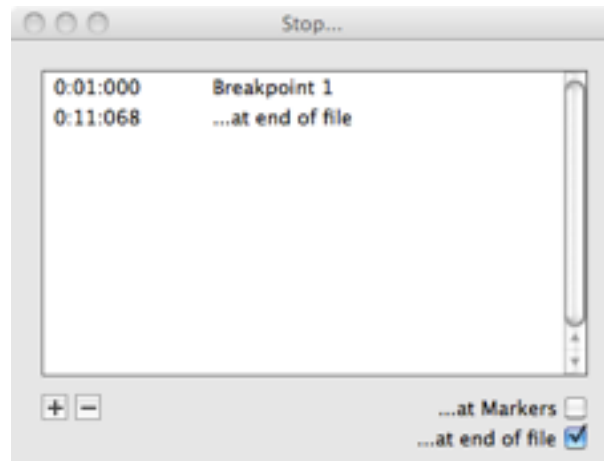
Breakpoints (including Markers) can be extremely useful, allowing the user to choose places in the file where processing is paused so that manual settings adjustments and/or manual repairs can be made. ClickRepair 3.6 greatly expands this feature over previous versions, which only recognized markers placed in the file by selected Audio Editing applications.<sup>9</sup>

Starting with DeNoise 2.4, when a file is opened, a “Breakpoints” dialog is available. This dialog will pick up any existing markers; more importantly, it may be used to set arbitrary breakpoints. The latter will not be added to the file itself; they are remembered only so long as the file remains open. The picture below shows three breakpoints that have been added by clicking on the “plus” button (which brings up yet another dialog for that purpose). Any of the breakpoints may be removed by

---

<sup>9</sup> DeNoise currently supports the MARK chunk, originally defined for AIFF files, although they can also be used in WAV files. A future version might also support the CUE chunk, which is defined for WAV files.

simply clicking on them, followed by clicking on the “minus” button. In the picture there is a fourth breakpoint set at the end of the file. This was done automatically because “Pause on completion of processing” is selected in the preferences. The two check-boxes at the lower right of the “Breakpoints” dialog are linked to the corresponding check-boxes in the “Preferences” dialog; selecting or de-selecting these check-boxes in the one dialog will carry over to the other, and also add or remove this final breakpoint. The breakpoint dialog will open automatically if “Stop at Markers” is selected in the preferences; otherwise it may be invoked from the “File” menu item “Breakpoints...”, once the file is open.



In using these breakpoints, remember that the “Abort” button can be used to restart processing from the beginning of the file. A consequence of having “Pause on completion of processing” as a marker is that processing can be restarted before the file is closed. Breakpoints may be inserted during or after processing. This can be useful if it is already too late to go back – simply insert a marker and then restart processing from the beginning. Remember, too, that with sound output turned off, processing is very fast when using the latest dual and multi-core computers.

## Resetting Defaults

Normally, DeNoise re-opens with the window in the same location, and at the same size, as when the program was last shut down. “Reset Default Window Location and Size” erases this persistent information, and resets the window immediately, but only after the “OK” button is clicked. The information will not be reset again until the Preferences dialog is again invoked.

“Reset Default Behaviour for all Dialogs” refers to the fact some of the dialogs have a “Don’t Show Again” button. This is to prevent such dialogs becoming annoying once you understand the flow of program use. If you wish to see any of these dialogs again, check this box and select “OK”.

“Reset other defaults” resets all check-boxes and radio-buttons not associated with the “Detection & Review” control group.

## 11 Other Features

### Behaviour of Buttons

The “Main Controls” group consists of five buttons, but some buttons have more than one function. This has already been mentioned indirectly, in relation to the “Open” button, which became the “Start” button, then alternatively the “Pause” and “Resume” button. Other buttons with multiple use are the “Preview” and “Play” buttons, which can alternate with “Stop”. These buttons can be operated using fixed keyboard shortcuts, as described earlier, but they never receive the keyboard input focus, and cannot be clicked using the spacebar, or any other normal method for the particular platform. By contrast, the buttons in dialogs<sup>10</sup> may be operated in the default manner by first using the tab key to bring the desired button into focus, then clicking the space bar.

### Abort Button

The “Abort” button results in a dialog giving three choices.

- Abort does just that: it results in the output file being deleted, since it would not be readable.
- Reset initializes processing back to the state it is in immediately after a file has been selected for processing. It is therefore necessary to click the “Start” button again.
- Continue cancels the request to abort, processing restarts immediately.

### File Dialogs

Java has two types of File Dialogs; ‘native’ dialogs as provided by the underlying operating system, and ‘Swing’ dialogs that are provided by the Java Virtual Machine (Java run-time system). Normally there is no need to worry about this, and ClickRepair will choose native dialogs whenever possible.<sup>11</sup> A small number of users have been plagued by random bugs in the native dialogs, as provided by the Java run-time, both on Apple Mac and Windows system. The “File Dialogs” item on the “File” menu provides a work-around in such situations, since the problem is not in the ClickRepair code itself.<sup>12</sup> Another work-around is to use drag and drop for file selection, although this places some limitations on the destination of the new file.

---

<sup>10</sup> Note that the “Shortcuts” window is not a dialog. Its “OK” button cannot be focussed, and the window can be closed using the close button/box in the title-bar, unlike dialogs.

<sup>11</sup> The exception is that Swing dialogs are used for assembling batch jobs, because of some limitations imposed by native dialogs.

<sup>12</sup> The symptom is a message saying that the selected file cannot be found, the name of the selected file having been corrupted as it was communicated to ClickRepair.

## **Keyboard Operation**

Nearly all the menu items have menu shortcut keys, and may be selected in this way. In addition, the buttons in the “Main Controls” control group may be operated by various shortcut keys (without the command or control key); these may be viewed by choosing “Keyboard Shortcuts” – see Chapter 5 for details.

In the main window, the “tab” key cycles the keyboard focus around the objects in the “Settings” group; once an object has the focus it may be operated in the normal manner – the space bar is equivalent to clicking a check-box or radio-button, while the left and right arrow keys operate the sliders.