Pristine Sounds 2000
Function reference

This function reference covers both Pristine Sounds 2000 Light and Pristine Sounds 2000 Pro. The differences between the two versions are stated functionwise.

February 7, 2000

Pristine Sounds 2000

Function reference

This chapter explain how and why the functions & effects in Pristine Sounds ‘99 work. The following paragraphs are used when describing the functions. Not all paragraphs are used in all functions.

Description of effect

A brief description of the effect.

Main benefit

The most obvious reasons to use this effect. In some cases it might be more useful to use other variations of the same theme, in that case existing in other functions. Such reasons are described here.

Description of parameters

A description of all parameters in the effect. This includes all buttons and sliders, the name of the parameters and the usage of them.

What could go wrong

It is not always possible to succeed the first time. Some families of effects such as noise reduction and 3D-filtering are known to cause problems for many users. Under this paragraph, possible problems and solutions to those are listed.

For you that want to know more...

If you have the habit of research it might be interesting how and why things work. Not all functions have this paragraph, either they are not very complicated or the principle is described somewhere else. It is not *necessary* to read this paragraph! It only exist for you who want to know more...

Preview

Most functions and effects can be “previewed”. Please refer to this section when referred to in respective function.

Press the 'Preview' button in the functions window to test its parameters. This will tell you if it sounded like you expected it to. You can analyze the result in both time-, and frequency space and do true A / B comparisons and real-time previews.

(Real-time previewing is not available in Pristine Sounds ‘99 Light.)

Description of function:

You can test your settings so they appear as you thought. This automatically takes a couple of seconds of your waveform and applies the effect to it.

Main benefit:

This is easier than applying the effect to some part of the waveform and un-doing if you are not satisfied.

Description of parameters:

Under the "Play original"-, and "Play effect" buttons there are some statistics about the waveform: Maximum & minimum sample value, RMS. and the number of channels.

Play buttons:

Press to hear either the original or the processed version of your waveform. The "Play original" and "Play effect" should be obvious, but the "Play combination" deserves some explanation: This enables true A/B listening. Define what will be played in each channel: the left/right original sound or the left/right processed sound. If you choose left original as the output on the left channel you will hear the originals left channel in the left speaker. Choose left effect in the right speaker to hear the effects left channel. For example; if the effect in question is noise suppression, you will hear the left channel in the middle and some noise in the right speaker, and no noise in the left speaker.

Stop buttons:

Press to stop waveform playing.

2 / 4 / 8 seconds:

The amount of time of the waveform that will be previewed. One second of silence will be inserted after this so that effects like reverbs or echoes can be heard fading out.

Position slider:

Set the actual point of testing.

Waveform / Spectral view:

Use the waveform view to see the waveform amplitude, or use the spectral view to examine the signals frequency components.

Intensity:

For use with spectral view, to brighten or dim the view.

Graphical windows:

A quick snapshot of the "un-processed"-, and the processed waveform.

Volume...:

Displays the "real-time parameters" toolbox.

Residual output:

This is an option for testing purposes only. The resulting output is the difference between the original and the effect.

Example:

1) When removing noise, the residual output is the noise itself.

2) Then applying reverb, the output will be the reverb part.

3) When using a filter, the output will be the added components only.

Real-time...:

By pressing this button, the preview dialog box will fold down into a small toolbox beside your current effect window. You can now use the controls of the effect window and hear the result playing at the same time. By pressing ’Stop!’ you exit the previewing mode. The real-time preview toolbox has the following parameters:

Bypass processing:

Use this option to toggle between the processed data and the original. This can be used for exactly determining the performed work.

Left channel only:

Use this function to process only one channel. This will make processing twice as fast when previewing stereo waveforms. Some effects can not use this trick.

Slow playing:

When the CPU's processing power is not enough, the playback is slowed down. Using a really low playback rate to begin with will ensure that the rate is less likely to fluctuate. A varying playback speed can be annoying, and this option makes playback constantly low.

Automatic treatments (batch processing)

Description of effect:

A “treatment” is a list of effects (up to 8) that are executed in turn without user intervention. You can use up to 128 files to be treated with your favorite effects.

To be able to use an effect in a treatment, you must first save the setting in question in the functions dialog box. You can then choose it in the treatment dialog box.

You do not need to have any waveforms opened when you use the Automatic treatment. You can load files from disk directly.

Most effects & functions available in Pristine Sounds except functions that need user intervention (Frequency space edit, pencil edit etc.) are available in a list to choose.

This function is not available under the Light license.

<u>Description of parameters:</u>

Add file:

This will enable you to add a file from disk.

Add opened:

This will enable you to add a waveform that is already opened in Pristine Sounds.

Remove:

This will removes a file from the list.

Clear:

This will remove all files from the list.

Play after processing:

The waveform will be played when the processing is done. When playing is done, the next waveform will be processed.

Save in output directory:

This will define where to save the processed files.

Do not close after processing:

This will let you access the files after processing, for example if you wish to add something extra to each file individually.

The way to work with automatic treatments is to define up to 8 effects, and at the same time defining which setting to use. There are 8 rows of identical parameters:

The approach is to "*Apply [**choose effect**] using the [**choose setting**]*". First you choose an effect from the **drop down list**. Most effects available in Pristine Sounds (except functions that need user intervention) are available here. Then choose which setting to use. (Use Browse to find the setting).

The effects will be executed in the same order they are displayed. You may have empty rows with no effects in between actual effects. The empty row will be ignored.

Press GO! to execute the treatment.

Load / Save:

You can save or load a treatment. A treatment is the list of effects, - not the list of waveforms.

3D-Filter

Description of effect:

The 3D-filter is a filter/compressor hybrid. This function can compress/expand certain frequency bands.

Main benefit:

This filter takes in respect that different sounds have different volumes. Soft sounds, can be filtered independently of loud sounds. The method can also be described as multi-band compressor.

The overall impression of a sound can be altered without too much conventional filtering.

A mix can be made "softer" or "harder".

Description of parameters:

This effect contain two identical parameter sets. The following description applies to both of them.

Reset:

Resets the slider-set

dB scale:

For the upper slider-set this means: Filter applies when a sound is this loud, or louder. For the lower slider-set this means: Filter applies when a sound is this weak, or weaker.

Filter sliders:

For every given frequency band, this amount of amplification is applied.

Reset buttons:

Flatten the filter.

Pristine mode option.

Filter will be applied on a “sound” by “sound” basis instead of a frequency by frequency basis. This will utilize the CPU more.

Load / Save:

Save or load a filter.

What could go wrong:

Due to the nature of the function, it is easy to make mistakes because the function has no counterpart in the “real world”. This is a new kind of function, so the chance for mistakes is obvious. Due to the complexity of the filter, a large set of examples have been included for guidance.

Has this happened when trying to create a filter?

- You can hardly hear any effect: Try to make the two dB-thresholds closer.
- You get to much effect: Try to rise the two dB-thresholds.
- You get a pumping sound / less attack: Try to make the two filter-sets more alike. Don't over-do it!

Graphical compressor

Description of effect:

This is a 'point-and-draw' graphical compressor.

Main benefit:

The compressor can be defined exactly by the user by drawing the curvature with the mouse.

Description of parameters:

The graphical area shows the current curvature of the compression. It can be altered with the mouse.

Compression slider:

This slider can draw "for" you. Set this slider to a appropriate value, and the diagram will follow. When you start drawing with the mouse again, the properties of this slider will be ignored. The "compression" is represented by the slope of the curve.

Threshold slider:

This slider can draw "for" you. Set this slider to a appropriate value, and the diagram will follow. When you start drawing with the mouse again, the properties of this slider will be ignored. The "threshold" will be represented by the point in the diagram where the slope begins.

Attack:

This slider defines the attack time of the compression. This can be regarded as a “reaction” time of the compression. Spans from 0.1ms to 400ms.

Release:

Ranges from 1 ms to 9000ms

Load / Save:

You can save and load your favorite compression.

Smooth:

Not everybody can draw a straight line...;-) With this option you can ease things out.

Reset:

Sets the compression to 1:1 (No compression)

Join stereo:

If you have a stereo sound, you probably want the same compression on both channels, or the sound will tend to jump between the Left & Right channels.

What could go wrong:

If the release is set to quick, for example under 10ms, the compression will appear to distort the sound, which will not sound good. A proper value should be 100ms and more.

Remove DC-offset

Description of effect:

Centers the wave around the baseline.

Main benefit:

This effect will not *audibly* alter the sound. It will however adjust any baseline shift added to the waveform during recording or bad effect processing.

Description of parameters:

The effect does not take parameters.

For you that want to know more...:

A baseline shift can occur if the recording equipment does not filter out a DC voltage on its inputs. If that is the case, and you feel that it is of considerable strength, maybe you should look into the analog electronics you use. If you have never experienced any unwanted features, that is probably not the case! The offset of a recording can not be heard by the human ear, as the frequency of a stable baseline happens to be the interpretation of exactly 0 (zero) Hz. That also, is considered a frequency, and can therefor be operated as such. It is for example possible to filter out a baseline offset with any high pass filter. In Pristine Sounds it is done by filtering out any frequencies below 5 Hz.

Deep bass

Description of effect:

With this effect you can add undertones to your sounds, making them appear more powerful and fatter than before.

This effect may best be suited for sound-effects and not for musical purposes.

Description of parameters:

You can add two sets of undertones to the sound. The "First undertone" and the "Second undertone". They have identical parameters as follows:

Enable:

You can enable either the first or/and the second undertone by clicking the appropriate check boxes.

Use sound below...:

This slider will select, which source-tones should create undertones. If set to (for example) 140Hz, only the frequencies below 140Hz will create undertones. Above that threshold, noting will be used. A tone of frequency 105

Hz will create an undertone of 57.5 Hz (first undertone) or 26.25 Hz (second undertone), while a tone of 150Hz won't do anything at all.

Amplify added deep bass:

By setting the slider to 0dB, the new undertone will be as loud as its “parent” tone. A tone of 100Hz at -40dB will therefor create an undertone of 50Hz at -40dB and if the slider is set to -5dB, the new tone will be at -45dB

What could go wrong:

Actually. It is quite easy to make the undertones appear very unnatural. As a common practice, a proper setting should be: setting the "Use sound below..." to high, and "Amplify added deep bass" to low, not the other way around. But of course every situation demand its very own solution.

Envelope editing

Description of effect:

With this function, you can draw the volume directly on the waveform with a pencil tool.

Description of parameters:

Main graphical window:

There are one graphical window for every channel of audio. This is a "paintable" area that represent a segment of the current waveform. If you have made a selection, this segment will be placed in the beginning of the selection. It is not possible to zoom in or out this picture, and the segment length is fixed (and happens to be 262144 samples).

Use the mouse to draw the volume with the currently selected function.

Smooth:

The curve you have drawn with the mouse may possibly be a little bit jagged. Press ’Smooth’ to make the curve more smooth.

Lock L+R:

Check this option if you intend to apply the same curve on both left and right channel.

Play Preview, Stop playing:

You can preview the segment before you apply the changes.

Apply:

This will apply the current changes.

Apply and scroll buttons:

Press the < button to apply, *and* scroll a little bit to the left. Press << to apply *and* scroll much to the left. Press > to apply *and* scroll little to the right. Press >> to apply *and* scroll much to the right.

Undo:

This will undo the changes you have made. You cannot undo more than to the state from where you entered the function.

Undo and scroll buttons:

Press the < << > >> buttons to ignore the current changes, and then scroll.

Fade in / Fade out

Description of effect:

This effect is used for fading the volume in or out.

Main benefit:

You can set the characteristics of the fade by simply altering the curve. (Fade hard, or soft.)

Description of parameters:

Behavior:

An image displaying the current fade curvature. You can not draw in this image.

Offset:

Maybe you don't want to fade out completely, but only to 10% or 80% If you are fading out, the Magnitude value is referring to the final volume, and if you are fading in, it is the initial value.

Curvature:

This parameter states how fast the apparent fading will be. Note that whatever selection you have made will still be the length of the fade. "Fast" in this case is far more subjective, as it is only the "envelope" of the fade that is considered. A value of 0 (zero) means that the fade is linear (in the behavior image you see a straight curve) and a value of 100 means a soft curve, slow in the beginning and end, and fast in the middle. This will be heard as a "fast" fade as the volume change is noticed more in the middle of the selection.

<u>What could go wrong:</u>

The behavior image, could sometimes fool the eye of how fast the fade will appear. The fast fade will tend to change the volume only in the middle of the selection. The normal procedure will therefor be to make a longer selection (even outside the boundary's where you don't want fading) if the fade is coming too quick!

<u>For you that wants to know more...:</u>

Looking at the curvature, it is possible to see the behavior of an "arctan" function (math-stuff...). Actually the mathematical function is (simplified):

Volume=(pi/2+arctan(speed*time)/pi

where the middle of the selection is "time=0". Thus changing the speed-criteria will zoom in and out of an arctan function.

Force to silence

Description of effect:

Sets the selection to zero volume.

Main benefit:

The Force silence sets the selection to zero. To avoid clicks and pops, the silence is faded in and faded out within a thousand of a second.

Description of parameters:

If the material is stereo, there is a choice to apply to any channel.

What could go wrong:

The Force silence is not possible to abuse. (Assuming you didn't silence out your *favorite* part...) Please choose this effect instead of muting the selection with the amplification effect.

For you that want to know more...:

When the characteristics of a sound is changed too fast, (for example during one single sample), noisy “clicks” and “pops” can be heard. These high pitch clicks consists of the high frequency components needed to force the "old sound" into the "new sound" (silence), and this is not a normal property of real life, and can not be heard in nature. A common medicine for these clicks is to zero out when the wave samples are as close to zero as possible. (This is called zero clipping.) But there is still a "force" needed to change the sound to silence in the endpoints, but the click is now much more quite and many are content with that. Pristine Sounds tries to reduce even these clicks by eliminating existing components in its corresponding differential equation.

Frequency space editor

Description of effect:

The frequency space editor is a “sonographic” editor. This tool is for the perfectionist who wishes to have total control of the editing of the waveform. Currently, noise reduction, de-clicking and filtering are editable.

The functions of this tool are applied with a brush and are *painted* onto the waveform. The viewpoint is a frequency view, similar to the 'View | Frequency View' menu.

Description of parameters:

Main graphical window:

This is a "paintable" area that represent a segment of the current waveform. If you have made a selection, this segment will be placed in the beginning of the selection. It is not possible to zoom in or out this picture, and the segment length is fixed (and happens to be 131072 samples).

Use the mouse to paint the effect with the currently selected brush.

Brush graphical window:

This is a preview of the brush. It shows the size and the center-of-strength.

Use the Two brush sliders to define the size of the brush.

Edit brush strength to define the strength of the current effect. Different effects have different units. The Filter has the range of 0-100 dB, and the Noise reduction tool 0-100%.

Filter brush (Amplify & Mute), "De-noise" brush, "De-click" brush, Undo brush are the currently available effects in the Frequency space editor. More effects may be implemented in the future.

The **Filter brush** behaves like an ordinary filter. You can use the Amplify option to rise the gain of a frequency band at a certain point, or the Mute option to lower the gain.

The **De-noise brush** behaves like the noise reduction tools in Pristine Sounds. With this brush you can remove noise extremely locally and precisely.

The **De-click brush** behaves like the VRT in Pristine Sounds. With this brush you can remove clicks extremely locally and precisely.

With the **Undo brush** you can “paint back” the current state.

View channel Left / Right:

You can only look at one channel at a time if you are editing a stereo waveform. Viewing a channel has nothing to do with where you paint.

Paint on channel Left / Right:

If you are editing a stereo waveform, you may not want to paint on both channels with identical moves. Select which channel to paint on. (If you are not viewing that channel, -you won't see anything.)

Apply:

This will apply the current changes.

Apply and scroll buttons:

Press the < button to apply, and scroll a little bit to the left. Press << to apply and scroll much to the left. Press > to apply and scroll little to the right. Press >> to apply and scroll much to the right.

Undo:

This will undo the current state.

Scroll buttons:

Press the < << > >> buttons to scroll without applying your current changes.

Play Preview / Original, Stop playing:

You can preview the changes before you apply them.

See extremely weak signals:

Use this option to enhance the weakest signals of the picture. With this option you can see the smallest components existing in the segment.

Generic noise suppression

Description of effect:

An FFT noise suppression method that do not need a prepared noise image scan.

Description of parameters:

Noise guide:

The arrow pointing at the approximate noise level of the waveform. The noise guide will scan through the file, but if the noise is about the same all over, it will stabilize immediately.

Strength radio buttons:

Please select the strength of your liking. The noise guide will probably point out an appropriate choice.

Dual mono / stereo check boxes:

When the waveform is in stereo, it may be appropriate to use the dual mono mode at some times, if the noise is significantly louder in either of the channels. Often however this will not be the best choice as the result tend to flutter somewhat.

What could go wrong:

This kind of noise suppression may leave traces of the noise when you have removed "most of it, but not everything": -very irritating squeaking. Either go back a step or two, OR, suppress even more.

Tip. Try to find a setting that meet your standard. Then go back one step. As a technician you probably think more about the noise, *than* about the music. -The listeners, listen to the music.

Graphical equalizer

Description of effect:

A graphical ’point-and-draw’ equalizer.

Main benefit:

You can easily draw the filter with the mouse, and with the help of the ’Frequency guide’ you can hear the effect of the filter. This is an FFT filter and does therefor not introduce any phase distortion.

Description of parameters:

Filter behavior:

An image describing the current filter to be used. With the mouse and the left mouse button, you are supposed to draw the behavior.

Reset:

Resets the filter.

Invert:

Inverts the current filter.

Load / Save:

Press here if you want to save you filter or load an old one.

Frequency guide:

The 'Frequency guide' asserts a tone representing the current position of the mouse in the 'point-and-draw' diagram.

High quality / Quick:

In most cases it may be hard to tell these settings apart, but when boosting the very low bass, it may be appropriate to use the 'High quality' check box which will give the bass a more natural dynamics ("kick or punch").

12/18/24 dB check boxes:

Check either of this check boxes to determine the strength of the curve. ("+-12 dB" is recommended, and means that the filter can amplify/mute up to 12 dB at the given frequency)

<u>What could go wrong:</u>

You have, of course, the option to make very weird filters, and perhaps that is what you want? An FFT filter can produce extremely high Q-values ("sharpness in frequency") which may not represent a real-world behavior.

Beware of boosting the bass to much if the waveform is already near its maximum. It will inevitable be distorted when clipped.

12 band equalizer

Description of effect:

This is a 12 band graphical equalizer. It can be used for altering the frequency characteristics.

Main benefit:

The filter is highly configurable and inuitive.

Description of parameters:

Gain:

The gain slider will define the overall gain for each frequency bands. All bands will be affected by this parameter. The default value is +12 dB, meaning that each band can amplify / mute a maximum of 12 dB. Please note that the gain can also be negative, meaning that the behavior of the equalization is inverted.

Bands:

For each frequency band, there is an associated slider. The band is emphasized at the frequency’s identifier underneath the slider. The band is restricted to function only between the 2 beside it. It will not affect frequencies outside its own band. The end-bands (the first and the last) are a not “band” filters but will expand to the end points of the spectrum.

The frequency identifiers:

The frequency identifiers can be found under the sliders and define the frequency of the band. Double click on the identifiers and enter a new frequency of your choice. (Because the band cannot be expanded outside the neighboring bands, you must enter a frequency in between.)

Load / Save:

This enables you to load or save a setting.

Image noise suppression

Description of effect:

Reduces, or even removes, noise from a waveform.

Main benefit:

This effect uses a new approach to FFT noise suppression. Usually, this kind of effect is very sensitive to changes of the noise in a sample, but in Pristine Sounds, there is a built-in adaptive mechanism to work against that. This effect is also useful to remove hissing-noises. This approach has a very high time- and frequency resolution.

Description of parameters:

Note! The noise reduction effect is a two-step-operation! To begin removing noise, Pristine Sounds first have to know what it is considered noise and what is not!

Select a part of the wave that contains only what you want to remove. This will be considered an unwanted component. Press the right mouse button and select "Scan noise image" from the menu. The noise will be analyzed.

In the Noise suppression dialog box, proceed with the following parameters:

Image of noise spectra:

An image describing how the noise look like. This is the spectra you previously scanned in with the right mouse button.

Strength sliders representing 4 different frequency bands:

To preserve the properties of the original sound, several completely separate layers of noise removal are associated with each slider. Please experiment with the strength of the sliders to find the optimum result.

Release time slider:

The default setting is “90”, meaning that the release time should react on transients of normal sounds. If you are removing noise from drums or hi-hats should have a small release time. If you are removing noise from voices you can have a long release time.

Load / Merge / Save:

Enables you to save & load the current spectra to and from the database. When you merge a spectra, you use both the current and the new spectra. For example: You could merge a spectra of low frequency rumbling and white noise.

Carving, option and slider:

Some times, one can hear weak remains of noise being left after reduction. That noise was stronger than what the noise image predicted. By using the carving option you can hunt down those remaining pieces. This function tries to isolate noise from music, and making the noise a "little more tender" and the music "more rigid" and therefor affecting the noise more.

Overwork protection:

This parameter has 2 uses. Primarily, it used for helping the user not to remove too much audio content during heavy restoration. Also, as more noise is removed from a waveform, the greater the inaccuracy of the operation becomes. The protection algorithm is such that if "too much" noise is removed, the actual values of the strength sliders are lowered dynamically. During "real life" it is more likely that a frequency component is muted somewhat than completely removed. The definition of "too much noise removal" is that a number of the frequency components actually have been completely removed.

When to use this parameter? The usage of the overwork protection may trick the user to use more noise reduction. But the primary use is to bring back removed audio content and keeping the strength as is. An important suggestion to the user is to be careful not to remove too much audio content in the process. One can always live with a little noise left. It is harder to be without all the beautiful music!

Pencil editor

Description of effect:

With this function, you are able to draw directly in the waveform with a pencil tool.

Description of parameters:

Main graphical window:

This is a "paintable" area that represent a segment of the current waveform. If you have made a selection, this segment will be placed in the beginning of the selection. It is not possible to zoom in or out this picture, and the segment length is fixed (and happens to be 512 samples). Every pixel on the screen matches one sample in the waveform.

Use the mouse to draw with the currently selected function.

Draw-pencil / Smear tool / Undo brush:

If you choose the Pencil-tool you may draw directly on the curve. If you chose the smear tool then a low pass filter is applied. The Undo brush will gradually give you back the real waveform.

View channel Left / Right:

You can only look at one channel at a time if you are editing a stereo waveform. Viewing a channel has nothing to do with where you paint.

Draw on channel Left / Right:

If you are editing a stereo waveform, you may not want to draw on both channels with identical strokes. Select which channel to draw on. (If you are not viewing that channel, -you wont see anything.)

Play Preview / Original, Stop playing:

You can preview the segment before you apply the changes.

Apply:

This will apply the current changes.

Apply and scroll buttons:

Press the < button to apply, and scroll a little bit to the left. Press << to apply and scroll much to the left. Press > to apply and scroll little to the right. Press >> to apply and scroll much to the right.

Undo:

This will undo the current state.

Scroll buttons:

Press the < << > >> buttons to scroll and ignore the current changes.

Presence

Description of effect:

Tries to mimic the behavior of the "presence" function in different kinds of amplifiers. The effect can be applied to all kinds of sounds though the result may vary.

Main benefit:

This will add a sense of "overdriven" amplification to the sound. Guitars and bass guitars will like this in particular as it in many cases will give a certain "power" or "distorted" look to the result. Note that this effect is not a true filter, even though the bass and treble will be greatly affected.

Description of parameters:

Valve vs. Sold state:

This will toggle the behavior of the effect. The general quality of a "valve" amplifier and a "solid state" amplifier has little to do with the relative quality of this effect. The two different settings can be very useful in different situations. You will just have to try it out...

Strength Slider:

This will adjust the actual quantity of the effect. Note that some sounds will be affected relatively little by this effect, and that is why the slider is very powerful, not to be set too high in most cases.

What could go wrong:

Actually, it very easy to over-do the effect. Some sounds are affected very little, and others extremely much. That depends on the actual frequencies in the sound. Most probably, it is very hard to exactly know what the result will be.

For you that want to know more...:

By saying that this effect "mimics an overworked amplifier" means that the signal in the waveform is restricted "not to move too quick". This is called slew rate, and in this effect it is implemented as restricting the "derivative" of the signal. First the derivative is taken and thresholded, and then integrated back.

Split & save in Audio CD format

Description of effect:

With this tool, you can split and save your waveform into a Audio-CD compliant format. The Audio-CD format demands that files are dividable by 588 samples and are longer than 4 seconds. The sample frequency must be 44100Hz, and the waveform in stereo. By placing markers into the waveform you can define appropriate points. Those points will be placed on CD-sector boundary (588 samples). You will be warned when you have defined segments that are shorter than 4 seconds.

All programs that burn to CD-R accept wave files of any length and characteristics, but there might be gaps between songs even though you burn DAO (Disk At Once).

Also, all programs that burn to CD-R will accept wave files that are saved with the ‘Save as...’ or ‘Export waveform as wave file’ function.

This function is not available under the Light license.

Description of parameters:

Start points:

A list of split points.

Delete:

Delete the selected split point(s).

Delete all:

Will clear all split points.

Add markers:

If you have markers in your waveform, they will be converted into split points. The markers will be quantized to 588 samples boundary, as if they where inserted within this Split & Save tool.

Play / Stop:

This will play the current window so you can see and hear where the split points are.

Upper graphical window:

A picture of the whole waveform.

Lower graphical window:

A picture of a small section of the waveform. This section is about 6 seconds long. Press the mouse in this window to create a split point.

Scroll bar:

By dragging the scroll bar you can set the lower graphical window to zoom in on any part of the waveform.

Use prefix on files:

Defines the name for the resulting files. For example, if you use the name "Album name" the files will be called "Album name-01.wav" "Album name-02.wav" etc.

Save files in directory:

Define which directory to save the resulting files. You can also “Browse...” to this directory.

Write to disk!

The resulting *.wav-files will be written to disk, only if you press this button.

Sound system profiles

Description of function:

With Sound System Profiles, you can calibrate your audio peripherals. All playback in Pristine Sounds will use this compensation so you will enjoy a perfect mixing situation at all times, even using very cheap equipment. Compensation implies frequency response, not noise/distortion reduction.

Enable Sound System Profiles:

Check here to enable a profile that you have already created. Use the “Browse...” button to find it.

Create / Edit profile:

The first thing you must do is to create a new profile for your sound system. A sound system includes: Soundcard, amplifier, speakers / headphones and your ears. The calibration is made on a subjective level using your ears as a final judge of calibration.

<u>How to create a Sound System Profile.</u>

In Pristine Sounds, a Sound System Profile, consist of 32 calibration points along the frequency axis, ranging from 40Hz to 10kHz. The calibration points have been strategically placed.

The calibration is subjective; it depends on who makes the profile. The user is supposed to decide whether or not the 32 test tones are louder or weaker than the reference tone. The reference tone is 1kHz and is played along with the test tones.

Initially, the test tones are set to be very low (24 decibels weaker than nominal). Going through the 32 tones, you are supposed to set the tones to be as loud as the reference tone (1 kHz). It is the "*as loud as*" definition that is subjective. It is very hard to compare the volume of a very low frequency tone, and high frequency tone. Especially the low frequency <200 are hard to estimate. What you will have to think about is: *Listening at the two pulsating tones - which one is the loudest - which one is the weakest?* It is normal not to be able to hear the difference of a decibel or two.

Press “Begin”. The first test tone will pulsate together with the reference tone. Use the slider to make the test tone "as loud" as the ref. tone. When you are content, press "Next tone" to calibrate the other test tones. When you are satisfied with all 32 points, press “Continue...” to continue with the wizard. Enter a description for your profile and save it. When you have saved it, you will have to enable it in Pristine Sounds by loading it using Enable Sound System Profiles and “Browse”.

What can you expect when calibrating your equipment?
If your sound system where perfect (as well as your subjective decisions...) you would get 32 tones calibrated to 0 decibel adjustments, all lying on the 0 dB line. But this never happens. Usually, a pair of headphones may have an extra strong 2-5kHz response so that you will lower the calibration in that band.

If you are unsure and cannot decide if the test tone is louder or weaker than the ref. tone, please set the calibration closer to the 0dB line.

Spectrum normalization (filter designer)

Description of effect:

This function creates filters. The automatic creation process is such that the current wave form will appear to have the same frequency density as a previously scanned wave form.

You should use this effect when you want one sound "have as much treble / bass etc., as an other sound"

Main benefit:

A filter can easily be constructed for a song or a sample that has been recorder improperly. It is not easy to know whether the recording has too much or too little bass / treble. These parameters can be adjusted to be “like another song” of your liking.

Why not make your song more Motown-ish?

Lost treble can now be found again.

Description of parameters:

Note! The spectrum normalization effect is a two-part-operation! To begin, Pristine Sounds first has to know just how the song “sounds like”, before a suitable filter can be constructed.

Select a part of the waveform that contains what you believe is characteristic for the song. Preferable, you want to select as much as possible, but at least 10 seconds. Press the right mouse button and select "Scan spectral image" from the menu. The song will be analyzed.

When done, you will be asked if you wish to add the analyze to the database, or go directly to the filter creation process.

In the Spectral normalization menu, proceed with the following parameters:

Load:

Load a normalization *target*. The current song will assume the same frequency characteristics as this target. It is vital that a target is loaded before you continue. Many targets have been included for you until you have some of your own in your database.

Export:

If you wish to tweak the created filter, you can export it to the ‘point and click graphical equalizer’. To be able to export, you must first have loaded a target. The exported filter can be loaded in the graphical equalizer to be edited just as any other filter. Beware that the “normalization process” is only related to the waveform, target waveform and this specific filter. (The created filter can of course be used in any way.)

Tremolo

Description of effect:

A tremolo will make the volume oscillate. It can be useful on guitars, synthesizer and sometimes on vocals.

Description of parameters:

Tremolo speed:

This is the oscillation speed. The range is 0-10Hz.

Amplitude:

The volume will be varying this much. If set to 100 the volume will swing from 100% to 0%, if set to 50 the volume will swing from 100 to 50%.

Phase difference:

If you are using a stereo waveform, the oscillation phase of the two channels will differ this much. If set to 0 degrees the tremolo will be identical in both channels. If set to 180 degrees, the channels are inversely mirrored; low volume on one channel => high on the other and vice versa. If set to 180 degrees, the overall impression will be that the waveform panorama is swinging from side to side.

Load / Save buttons:

To save and load settings and for usage in batch processing.

What could go wrong:

Too high speed combined with a too high amplitude will be very hard to listen to. Normal settings are speeds of 2-5 Hz combined with an amplitude of 20-30%.

Volume

Description of effect:

Used for altering the volume of the sound.

Main benefit:

To avoid distortion, you can find the maximum amount of amplification possible.

Description of parameters:

Join channels:

This option will lock the two sliders to the same volume.

Faders:

Lets you set the per-channel associated volume change. A range of +-48 dB would give you a wide range of choices. (A 16 bit sound has a dynamic range of 96 dB.)

Find maximum:

This will scan the current selection for the maximum amplification possible.

What could go wrong:

With a 48dB range it is quite easy to make mistakes. If you intend to "silence" a part, the best way is to use the "Force silence" function. (It is quicker). Also, for a 16 bit sound, 48 dB is NOT enough to create silence, as 48 dB only covers the first 8 bits. For complete silence, use the function twice, setting the faders to -48 dB. (That is mainly why a dedicated function is quicker.) It is also easy to be fooled to believe the "turning the faders to max" won't ruin the sound as much as it does. Never do that, if "Find maximum" hasn't suggested so.

Note: If you really want to amplify more than you are advised to by the 'Find maximum' property, you should use the "Super gain" plugin which lets you crank up the volume without clipping!

Normalize against RMS

Description of effect:

Used for normalizing the volume of a waveform. The RMS of the waveform will be calculated by averaging the volume of the two channels, if in stereo.

Main benefit:

This function does not have to use any input parameters. It will normalize the *average* volume of the waveform and will not be fooled be occasional peaks.

Description of parameters:

New RMS:

The default is -15 dB. This value is good for most types of material. Such a waveform would utilize most of the dynamic span without clipping too much.

What can go wrong:

The amount of gain applied to the waveform is directly proportional to its *average* volume. The mean that if there are a few loud transients in a soft environment, those peaks will be amplified in the respect of the average volume, not to the highest peaks. In such a situation, the peaks will be heavily clipped. Try to have a consistent volume throughout the song, for example by using compressors.

Auto trim silence

Description of effect:

This function can be used for deleting or silencing parts of the waveform that are below a predefined volume.

The auto trim function can only be used on the complete waveform; not on the current selection.

Main benefit:

This function is better and more versatile than the average “noise gate” as you can choose to either delete, denoise or silence the quite parts.

Description of parameters:

Volume threshold for “on”:

This will define the volume at which the function will “turn on” the audio from being turned off. This volume must be equal to or higher than the “turn off” threshold.

Volume threshold for “off”:

This will define the volume at which the function will “turn off” the audio from being turned on. This volume must be equal to or lower than the “turn on” threshold.

Sliding window length:

The volume is extracted from the audio by taking the mean of a small segment. This segment (or “window”) can be between 10-2000 ms.

Trigger offset in window:

Because of the nature of a sliding window, the turn on/off moment will be somewhere in the middle of the window. This offset can be changed for the moment of turning “on”.It can be useful if the beginning of a loud transient is lost (however quite improbable).

Use 100Hz high pass filter:

To make the auto detection process less susceptible to low frequency noises and hum, a filter of 100Hz can be applied to the detection process. (This will not affect the sound of the audio, but only the triggering of the detection process.)

Range:

The trimming of the waveform can be restricted to the following:

* Whole waveform (no restrictions)
* Beginning and end
* Only beginning
* Only End

The “beginning and end” are those moments that are directly attached to the first and the last sample.

Initial state:

The volume of the very first moment of the waveform is hard to determine because the detection process has no knowledge of what could be “virtually” found before the start of the waveform. This will help the detection process to determine this.

Action:

The following action can be defined for a “silent” part:

* Removal of part (a new undo layer for each part)
* Silencing of part (insertion of zeroes)
* DeNoising of part (3 band generic dehissing)

DeHiss

Description of effect:

This function will suppress low level noise. It is immune to hissing noises that has a varying volume.

Main benefit:

No “image” is needed for dehissing audio. This function does not use “FFT’s” but instead common “IIR” filters meaning that the transients are perfectly preserved.

Description of parameters:

Band sensitivity:

A unique dehissing is performed in 6 different frequency bands; 1 kHz, 2 kHz, 3 kHz, 5 kHz, 8 kHz and 9 kHz. The threshold for the definition of “noise” is determined with these sliders. The unit for the sensitivity is similar to dB, but is slightly different. The higher the sliders, the more sensitive the dehisser becomes,

Suppression:

This slider will define how much noise is removed in each frequency band. The range is 0-20dB. This will affect all bands.

Use reverb to lighten up:

A discrete reverb can be applied to make the result somewhat brighter. It has no user adjustable parameters.

Echo tool

Description of effect:

You can select a segment to use as a source for echoing. (Select a short beat or word, and echo it through the whole song without using any other part of the waveform.)

You can also chose to echo only within the current selection.

This echo can convert between beats per minute (BPM) and samples per second.

Description of parameters:

Use selection as source / echo inside:

You can either implement the echo inside the current selection, or you can echo a selection throughout the rest of the waveform.

Echo for at least...:

The minimum number of echoes when the selection is used as a source.

Gain:

The amount in percent of the signal used for echoing.

Feedback:

The amount in percent of feedback from the echo.

Delay:

You can state a value either in samples or in BPM’s. When playing the waveform, you can tap the ’B’ key for an estimate of the BPM in the status bar.

Electric guitar reverb

Description of effect:

A quick and handy reverb, specially designed for the guitarist in mind.

Description of parameters:

Amount of effect:

This will define the “wet” mix.

Width of room:

The reverb will mimic a room, and this will be the approximate size of that room.

Wall hardness:

The room can be set to absorb or reflect different amount of sound. A "harder" room has a brighter impression. This is basically a low pass filter where a "hard wall" reflects more treble and a "soft wall" reflects less treble.

Presence:

Different sounds react differently to this property, so you must try to find a good value yourself. This setting can in some way mimic the properties of a hard working amplifier. Please have a look at the “Presence” effect for more details.

Load / Save:

You can load or save the reverb to disk.

EZ compression

Description of effect:

A fast and easy-to-use compression tool, with only a few parameters.

Description of parameters:

Compression:

This will define the amount of compression. Vocals and slow transient sounds may preferably be compression more than drums and sounds with fast attacks.

Release time:

The release time defines the reaction time of the compression. Put simply: If the volume that you wish to affect is varying slowly, then set this parameter to high. This would probably be a mix of instruments. Single instruments would probably need a faster reaction time.

Stereo / Dual mono:

If the waveform is in stereo, the two channels may be compressed individually or “as one”. Preferably, compress them as one (Stereo mode) as this will keep the sound from flutter from side to side.

Hi-fi Chorus

Description of effect:

This is a mono and stereo chorus. A chorus varies the playback rate of the waveform. A mix between the original and a slowly varying copy may produce great stereo effects. It may be suitable for backing vocals, for electric guitars, and on full mixes. The overall impression of a chorus is a "fatter" sound.

Description of parameters:

If the waveform is in stereo, these parameters go for both channels.

Chorus speed:

The speed of the oscillation. You may not want to use a speed that is too great. A high speed oscillation will make the sound "vibrate" or "flutter".

Chorus depth:

The “width” of the oscillation. This parameter should be inversely proportional to the speed: *high speed* -> *small depth, low speed* -> *greater depth*. The depth of the chorus determines the maximum playback variation of the waveform.

Chorus delay:

This will apply a delay between created chorus sound and the original sound. When creating stereo effects, try to use different delays for each channel.

Load / Save:

This will load and save the setting to disk.

Room simulator

Description of effect:

Create a "virtual room" for your recordings. You can easily design your own room. Define its walls and their absorption. Use the two channels of the waveform, symbolized as two speakers, and resample the result, symbolized as two ears.

The difference between a room simulator and a reverb is that a reverb is basically designed to create the "audibly nicest" effect whereas the room simulator is designed to create the most "mathematically correct" model of a room.

Description of parameters:

Length / Width of room:

You can set the size of the room by dragging the sliders. The range can be from 1 meter to 100.

Wall parameters:

Check box Enabling or disabling the wall.

When a wall is disabled it is "physically" removed and the sound “passes out” of the room.

Reflection slider:

The amount of reflection of the wall. 100% means that all sound is reflected back.

Material slider:

There are 10 different materials to every wall. The materials represent different kinds of absorption.

Load / Save buttons:

You can add and get your favorite room from disk.

Super gain

Description of effect:

Raise the volume “above” the 0 dB level without clipping the waveform. The 0dB level is the maximum theoretical volume a digital signal can have. This function can apply an algorithm on parts of the waveform that are amplified above this level so that they are quickly compressed instead of clipped.

Description of parameters:

Find maximum:

This will search through the waveform, for the maximum possible amplification *without* altering any samples. If the slider is raised above this level, an anti-clipping algorithm is applied.

Slider:

Sets the volume of the resulting waveform.

Vinyl restoration tool (VRT)

Description of effect:

This function can remove impulse noises (pops & clicks) in vinyl recordings. It also has settings for simple noise reduction.

The philosophy behind this tool is to favor *intact music* instead of too much restoration.

The effect is divided into 4 modules; groove & rumble reduction, pop & click reduction, white noise reduction, room ambience. Each of these modules use up a specific amount of processing time, and when they are all enabled, it may take up to 3 times the playing time to restore a waveform on a 200MHz computer. Each module takes a fixed amount of processing time.

Description of parameters:

Attenuate grooves & rumble module

Strength slider:

This will set the amount of low frequency noise suppression.

Pops & click removal module

Primary algorithm:

Sensitivity slider:

This will set a threshold for *how to* detect impulses (=pops & clicks). If this slider is set too high some musical content will be "detected" as being impulse-noise as well. If set too low, some obvious pops may go undetected. Default setting is 50%.

It may be possible that some clicks can never be removed, whatever sensitivity is used. The reason for this is that if 2 or more clicks are too close to each other only 1 will be removed. This is a safety precaution.

Strength slider:

This setting should only be used, if some of the detected clicks / pops have been detected, but not entirely removed. Default setting is 50%.

The algorithm by itself is clever enough to adjust to pops of different strengths. It is advised to manually remove clicks with the “Frequency space editor” instead of rising the strength slider too much.

Double pass:

The algorithm that is searching for impulse noises may sometimes be fooled when several impulses being too close together, the result being that only the largest pop is removed. If that is the case, then use the double pass option to re-scan around that area. This will effectively run the function twice.

Use internal S-Comp object:

The S-Comp feature will reduce artifacts when removing too much pops & clicks. The most common artifact is low frequency distortion, and distortion in general. It is advised to always have this feature enabled.

Secondary algorithm:

Use secondary algorithm:

The pop & click removal tool consists of two independent algorithms. Try this approach if you can hear clicks being left in the result after the main algorithm. This option enables a method that is less selective than the main approach, therefor leaving less clicks behind also implying that the result becomes less pristine.

The secondary algorithm is mainly intended for use with very old material.

White noise reduction module

Strength slider:

This function is very much alike the “generic noise removal” function in the 'Special effects | Generic noise removal' menu.

Room ambience Module

This module will place the restored material in a simulated room. You may use the rooms from the room simulator plugin in the 'Add on | Room simulator' menu as the file format is identical and that means that you also can design your own rooms. A recommendation is to only use this module after really heavy restoration. It dampens the noise and clicks as well as creating a nice ambience, but you will find that the natural ambience of a recording may be preferable. Use this function only when the restoration has damaged the natural ambience.

Use global S-Comp object:

The S-Comp feature will reduce artifacts when removing too much pops & clicks. The most common artifact is low frequency distortion, and distortion in general. It is advised to always have this feature enabled.

An S-Comp object offers a functionality that enables any effect to do a "before & after" comparison. This allows the effect to actively "know" what turned out good and what turned out bad after processing. Using an S-Comp object may take up to 10% more processing time but enables the effect to "correct itself" for artifacts and mistakes. The default is enabled.

Material was originally in mono:

This could be a useful option for really old records! It is a feature of the global S-Comp object which therefor must be checked. The S-Comp object uses the fact that music that was recorded in mono but sampled with a stereo pick-up must have two stereo channels that are very much alike (perfect match in theory) while the vinyl noise and impulse noises are effectively random. If you listen to a mono record that was sampled by a stereo pickup you can hear the pops & clicks all around in the panorama, while the music lies still in the middle. The S-Comp object tries to combine the common factor of the two channels. The result will be a stereo waveform where the two channels are identical.

As a reminder; do not use this option on a waveform with stereo properties.

Voice multiplier

Description of effect:

This effect will create multiple copies of the waveform, tuned to different notes. It can be useful for creating “choirs”, or simple backing vocals.

Description of parameters:

Voice check box:

Enable and disable the current voice. You can use up to 12 at a time.

Load / Save:

You can load and save the current choir from / to disk.

Set button:

Press this button to bring up the voice properties dialog box. These are the settings you can control:

Note:

You can transpose the voice up and down 12 halftones compared to the original.

Volume:

The volume of the new voice. It can be 0-100% compared to the original voice.

Off-pitch:

Sometimes it could be useful to give a voice an "inaccurate" or a "random" impression, especially if you have several voices singing at the same transposed note. *But*, the more voices you have, the less off-pitch you should use, because it will not sound great if the whole background choir where intensely inaccurate... Use with caution.

Create voice from:

When using a stereo waveform you need to specify from where in the panorama to create the voice.

Placement of voice:

When using a stereo waveform, you need to place the new voice somewhere in the panorama.

R2000 reverb

Description of effect:

This reverb has a lot of parameters. The included presets should be used for learning the different kinds of behaviors the reverb has.

The walls in the simulated room of this reverb can be set to mimic a variety of materials.

This reverb can for the moment *not* be previewed in real-time.

Description of parameters:

Input gain:

The amount of signal to feed the reverb. The input should be set between 100% and 200%.

Original / effect mix:

This parameter defines the dry / wet mix. Set this to 90% / 10% for normal use. This would give a signal of 90% original and 10% effect.

Reverb from selection:

If this option is checked, the current selection will be used as a source of reverberation. The actual “wet” signal will be applied to a little bit outside the selection.

Quality:

Set this to 10 for normal use. You will probably never have to go higher than this if you do not have a waveform with much treble or with a lot of quick transients. The quality parameter is set too low if you can hear a faint echo of "needles falling on the floor" after a quick transient. The setting determines the density of the echoes.

Size of room:

The size of the virtual room that the reverb tries to simulate. Range 5-100m.

Stereo spread:

If the waveform is in stereo, you can add a “broadening” effect. In this effect that is accomplished by simulating *asymmetries* in the room. You can set this to a 100% at most times.

Use walls with...:

Check this option to enable the "equalizer-looking" filter. This will enable you to set the properties of the walls in the room. When a slider is set "high" it will reflect the frequencies in question, and when the slider is set "low" the wall absorbs the frequencies in question. Most often, walls tend to absorb high frequencies.

Reset:

This will reset the wall absorption filter to the most basic setting.

Type:

Select a reverb type of your liking. You will just have to experiment to find your favorite.

Load / Save:

You can save or load favorite reverbs to and from disk. This reverb has been accompanied with a lot of examples, if you find that it takes too long time to experiment.

The Action Navigator

Description of function:

The action navigator is a graphical front end to the undo-system. With this tool you can navigate freely through different versions of your editing

Description of parameters:

The main graphical window:

The graphical window hosts the states of your editing. The states are displayed in the form of colored boxes. The current state is displayed in red. The rest of the states are displayed in blue.

The relationship between the states is symbolized with a line going from the parent state to the child state. On parent state can have any number of child states.

The sliders:

You can use the sliders to move the graphical display to show any existing state.

The ‘Scale’ slider:

To be able to display more (or less) states in the graphical window, you can zoom in and out in the hierarchy.

Double-clicking on any state will make this the currently active state.

By pressing with the right mouse button on a state a floating pop-up menu will be displayed:

Activate:

This is the same as double clicking on the state. It will be selected as the current state.

Destroy branch:

A “branch” consist of a state and all states descending from this. It is not possible to destroy a branch if the currently selected state is descending from the state in question.

A state has the following properties displayed:

The color red mean the state is currently active. The color blue means it is not active.

A number in parenthesis defines the “age” of the state. All new states get a new number counting from 1 and upwards.

A description of what lead to the new state; for example “loading of waveform”, “Normalization”, “Reduction of noise” etc.

Real-time parameters toolbox

The real-time parameters toolbox can be used to control & view parameters in real-time.

The soundcard can be controlled with these parameters:

Volume:

This will set the output volume of the soundcard.

Gain:

This is a *software* gain. You can amplify the output up to 20 times so you can hear small details that previously where hard to hear.

Please note that the output may be distorted at times when the signal will be clipped. Use software gain only during silent passages.

Seek speed:

To make navigation easier, you can seek through the playing at different speeds, A quarter, half speed, normal, double and quadruple speed.

Sound System Profile:

You can enable or disable your current Sound System Profile. (You cannot choose which profile to use. This must be done in the ‘File -> Sound System Profiles’ menu)

<u>The hard disk can be monitored with these parameters:</u>

You have two directories for storing temporary files. A tip is to put each directory on unique hard disks if you have a SCSI system, this will speed up performance. The two directories are monitored for storage space. You can easily see if you are starting to run out of space. Note that early Windows 95 (build 950a) may not properly report storage space properly.

<u>The waveform can be displayed with these parameters:</u>

Current Zoom:

This is the number of samples per pixel on your display. Use the controls to change zoom on the waveform. If you zoom out more than 512 samples per pixels, waveform-caching will be used. This mean that the drawing of the waveform will be *quicker*, but to the expense of *accuracy.*

Current magnification:

This is a magnification of the y-axis if you use the waveform view, or a magnification of the intensity in the frequency view.

Selection length:

This is the length of the current selection (in time).

Length:

This is the length of the waveform (in time).

Playing xx.xx:

This is the time of playing since start.

Real time frequency analyzer

The real time frequency analyzer can be used for analyzing the frequency content of a waveform or the inputs of your soundcard.

The analyzer accept 512 samples of data per analyzing moment. This mean that there is a limited resolution in the frequency range. Typically, a file with a sampling frequency of 44.1kHz will be displayed with a resolution of +-86 Hz. This resolution can not be changed.

Analyzers window:

This window will contain the frequencies of the content being analyzed. The axis spans the whole frequency range of the waveform /soundcard input, and 80 dB.

Scaling:

The scaling knob can be used for changing the range of the dB scale. It can be adjusted +-30 dB.

Left / Right:

For playback, recording and processing events, you can choose which channel to analyze. If the waveform is in mono, this is ignored. The channel is also ignored if you click on a certain channel with the mouse in the waveform window.

Mouse click, Playback, Recording, Processing:

The analyzer can be set to accept any of these inputs. Only one input will be active at once.

When enabling “Mouse click” you can click the mouse in the waveform window and the waveform will be analyzed that exact point in time.

When enabling “Playback” the analyzer will accept data that is heard through the speakers at the moment. This include regular playback, but also previewing and real-time previewing.

When enabling “Recording” the analyzer will accept data from the soundcard. What is seen in the analyzers window will be the data saved to disk.

When enabling “Processing” the analyzer will accept data that is currently being processed. This is turned off by default because it will slow down processing efficency.

Clicking the mouse on the analyzers window:

This will display the frequency and dB of the mouse click. If you release the mouse button, it will follow the largest peak and therefor displaying the amplitude of the largest peak in a small region.

Preferences

<u>Description of parameters:</u>

Fit wave to window when opening:

When your waveform is opened, it will be zoomed out to fit the screen.

Use wave caching:

To save time while displaying the waveform when zoomed out more than 512 samples per pixel, a faster, less accurate, display mode is used.

Position guide when playing:

When playing the waveform, it could be useful to see exactly where you are. A vertical line displays the current position.

Track window when playing:

If you use a position guide, you can also let the window track the current position on the screen when playing, so you never loose track of where you are.

Spectrum meters:

If the power meters are enabled, they will also show a spectrum distribution.

Visual update time (ms.):

"Visuals", are the power meters, the position guide and the window tracking. For absolute smoothness, use 20ms if you can, but some video cards / CPU’s handle at most 50ms (20 frames per sec.) which is the default.

Don’t show "real-time parameters" toolbox on startup:

The toolbox containing real-time parameters such as volume & gain, seek speed, zoom, Sound System Profiles.

Startup splash screen:

Save some time. The splash screen is the picture displayed when loading the program.

General exception handling:

Disabling exception handling is a feature for 3:rd part developers.

Release soundcard...:

The soundcard will be released (stopped and closed) when another application becomes activated.

Windowing functions:

When using the frequency view of a waveform, it could sometimes be useful to use different "windowing" functions. If you do not know what they are, please try the Hamming window.

Use playback devices...:

If you have more than one soundcard, you can choose which one to use when playing sounds.

Browse paths:

These paths represent the installation of Pristine Sounds. You may change any of these if you are running out of hard disk space. When editing large files, the undo information may take up a substantial part of the disk. Please feel free to change either the First undo path or the Second undo path to any available hard disk / partition. There is a slight increase in performance if you choose 2 different drives for the two directories.

Menus

The following menus can be found in Pristine Sounds:

The Toolbar

File

Edit

Special

Functions

Effects

Add-on effects

All DirectX plugins from 3rd party vendors will show up under this menu. Some of the native effects in Pristine Sounds are also listed here.

View

Windows

The tool bar:

Pristine Sounds has a toolbar under the regular menu so you can reach the most important functions more easily.

The Open and Save buttons:

The open and save buttons have the same functionality as in the menu.

The waveform playing navigation buttons:

These buttons let you control the playback & recording of waveforms.

Stop:

Pressing this button will stop all playback. If you currently are playing, the cursor will be placed from where it started before playback, but if you go from a paused state to a stopped state, the cursor will remain where it currently is.

Pause:

The pause button will temporarily stop the playback. The playback will continue from the same point if you again press the play button again. If you press the stop button in a paused state, the cursor will remain where it is.

Play all:

The play all button will play from the beginning of the waveform to the end.

Play selection:

The play selection button will play your current selection; or if there is no selection, from the cursor to the end.

Record button:

The record button will not record a new waveform directly. It will start the Record tool from which the recording can take place.

The cursor navigation buttons

The cursor navigation buttons will let you jump back and forth between markers.

Go to start:

The window is scrolled to the beginning of the waveform and the cursor placed at the first sample.

Go to previous marker:

If you have markers to the left of the cursor, the cursor will be placed there. If there are no markers, the cursor will be placed at the first sample.

Go to next marker:

If you have markers to the right of the cursor, the cursor will be placed there. If there are no markers, the window will be scrolled to the end, and the cursor placed at the beginning of the window.

Go to end:

The window is scrolled to the end of the waveform. The cursor will be placed at the beginning of the window.

The real-time parameters toolbox button

Pressing this button will enable or disable the real-time parameters toolbox. The toolbox will automatically be enabled on startup if not stated otherwise in the Preferences window.

File:

Open:

Opens an existing waveform in a new window. The internal format of every waveform is 16 bits.

Record new waveform:

The same as pressing the record button on the toolbar. A new waveform will be recorded and opened.

Close:

Closes the current waveform, asking you to save it first.

Close all:

Closes all waveforms, asking you to save them first.

Save:

Saves the waveform with the name by which it was opened.

Save as...:

Saves the waveform with a new name and/or file format.

Save in Audio-CD format...:

(Not available in the Light license) Opens up an editor that enables you to split you waveforms into pieces compliant to the Audio-CD format. Wave files will be created starting & ending at CD-sector boundary (588 samples) and being at least 4 seconds long.

Export selection as...:

Will save your current selection as a *.wav waveform.

Save all:

Saves all waveforms.

Waveform properties...:

Display information about the number of samples, channels, file sizes etc.

Batch processing - Automatic treatments...:

(Not available in the Light license) Use up to 128 waveforms and apply up to 8 effects. Ideal when you plan to process many waveforms in the same way. No user interaction will be needed when started.

Preferences:

Definitions of Pristine Sounds properties.

Sound System Profiler...:

Calibrate your audio peripherals and use a real-time compensation for all waveform playing in Pristine Sounds.

Exit:

This will close all waveforms and exit Pristine Sounds.

Edit:

Undo action:

Redo action:

Action Navigator:

Cut:

Copies the selected interval to Pristine Sounds clipboard and deletes that part from the waveform.

Copy:

Copies the selected interval to Pristine Sounds clipboard.

Paste:

Inserts the copied interval into the waveform. You will be asked if you wish to "paste mix" (overlaying the copy onto the waveform) or "paste replace" (replacing with the copy). Please note that you can use the ’Mix-property’ for mixing and fading.

Paste as new...:

A new window will be created containing the content of the clipboard.

Delete interval:

The selected interval is removed. The waveform becomes shorter.

Insert silence:

Adds empty space.

Deselect interval:

This will release any selection you might have.

Crop:

This will delete everything but the current selection.

Convert to Mono / Stereo:

If your waveform is in mono, it will be converted into stereo and vice versa.

Special:

Statistics:

The selection (or the whole waveform if there is no selection) will be scanned and analyzed. The statistics of the waveform will be displayed: DC-offset, max sample value and noise level.

Real-time parameters:

This will enable or disable the real-time parameters toolbox.

Create duplicate waveform...:

A new window will be opened, having the same content as the currently selected waveform.

Scan menu popup:

Scan noise for noise suppression:

Same as the right mouse button menu. The selected interval will be considered "noise" when entering the Image noise suppression function.

Scan spectrum for normalization:

Same as the right mouse button menu. The spectrum will be used when normalizing the waveform spectrum in the Normalize spectrum function.

Scan statistics:

Same as the right mouse button menu and the Statistics menu.

Markers properties...:

Enables you to rename and delete markers.

Go to marker popup menu...:

This will place the cursor at the selected marker position.

Create selection between markers...:

This will create a selection between two of your markers or the start & end points.

Drop marker at cursor...:

This will create a marker at the current cursor position.

View markers list:

This will open up a small toolbox where you can jump to and play from the currently available markers.

Viewpoint properties...:

Enables you to delete and rename viewpoints.

Create new viewpoint...:

This will create a new viewpoint with the current properties.

Viewpoints list:

This will open up a small toolbox where you can recall the currently available viewpoints.

Re-sample waveform:

This will re-sample the waveform to a new sample rate.

Use other playback rate:

This will change the sample rate of the waveform, *but without re-sampling.*

Functions:

Volume...:

Changes the current volume of the waveform.

Envelope...

This will allow you to “draw” the volume with a pen.

Normalize volume:

This function takes no parameters. The waveform will be analyzed, and the maximum amplification will be applied without the waveform being clipped.

Fade in / out:

Applies a sliding volume change to the interval currently selected.

Force silence:

Sets the selection to zero.

Invert phase:

Inverts the phase of the waveform. Every sample is turned 180 degrees.

Swap channels:

Left channel becomes the right, and vice versa.

Reverse:

The selection will be reversed in time.

DC-offset...:

This function takes no parameters. The DC-offset will be removed.

Frequency space edit...:

This enables editing of the waveform in the time & frequency space. You will gain control over all tones of the waveform at all points in time.

Pencil edit...:

Draw the samples with a pen.

Effects:

Image noise suppression:

Implements a FFT noise reduction. Before you enter this menu, please scan a noisy part of the waveform with the right-mouse button menu-choice ’Scan noise image’.

Generic noise suppression:

Generic FFT white noise suppression. No scanning needed.

Graphical equalizer:

This enables you to ’point-and-draw’ an equalizer curve.

3D graphical equalizer:

This is an equalizer/compressor hybrid.

Normalize spectrum (Filter designer):

First analyze the waveform with the right mouse button menu ’Scan spectrum’. Then enter this menu and load *another* spectra from the database so a filter can be constructed.

Graphical compression:

Applies a dynamic compression / expansion to the waveform.

Deep bass:

Add under-tones to the waveform.

Presence:

Adds presence to the waveform. Imitates an overworked amplifier.

Tremolo:

Applies a tremolo effect to the waveform.

View:

Fit wave in window:

Zooms out the waveform so it fits in the window.

Zoom in selection:

Zooms in the selection so that it fits in the window.

Zoom in / out little:

This will zoom in and out a small factor. The window will still be displaying the current position.

Viewing mode

This will displays the *waveform or the frequency components*.

Normal view: This is the default "waveform view".

Frequency view: This is a more powerful way of looking a the waveform. You can see the exact frequency components of the sound. At any time of the time-line (left / right on the screen) you can see the frequency components (up /down on the screen). If a stereo waveform is viewed, the window will be split into two.

The strength of each frequency component is symbolized with the strength of the color.

Submenus, Full frequency range, lower range, extremely low range:

You can "zoom in" on the lower frequencies by selecting the appropriate choice:

Full range: full range of frequencies.

Lower range: up to 1/4 of the full range.

Extremely low range: up to 1/16 of the full range.

Submenus, Stereo / differential view:

When viewing a stereo waveform in frequency view, the two channels can be view separately (stereo) or "overlaid" (differential).

-When displayed in stereo mode, the two channels are viewed separately.

-When displayed in differential mode, the left channel is displayed in a red color, and the right channel is displayed in a green color. This could be useful for understanding the differences of the two channels.