

Sonnox Fraunhofer Pro-Codec

Operation Manual

Version 1.1 5th September 2012

1. Introduction

The Sonnox Fraunhofer Pro-Codec Plug-In is designed for the real-time auditioning, encoding and decoding of audio signals by the Fraunhofer codecs.

In the past it has not been possible to critically audition and then encode multiple formats in real time within a host DAW environment. The revolutionary Sonnox Fraunhofer Pro-Codec now makes this possible, and includes extensive monitoring tools and built-in encoding/decoding features.

With this 64-bit compatible plug-in, mixing engineers can produce compensated mixes optimised towards specific target codecs, thereby ensuring maximum fidelity. Similarly, mastering engineers may audition material in the final format, and produce compensated, optimised masters for final encoding and distribution.

2. Supported Codecs and Formats

Major codecs supported:

mp3, mp3 Surround, AAC-LC, HE-AAC, HE-AACv2, MPEG Surround,
Apple AAC (iTunes Plus; Mastered for iTunes) (Mac only)

Lossless codecs supported:

mp3 HD (up to 16bits/48kHz) and HD-AAC (16/24bits/192kHz)

Formats supported:

Mono: mp3, mp3HD, AAC-LC, HE-AAC & HD-AAC

Stereo: mp3, mp3HD, AAC-LC, HE-AAC, HE-AACv2 & HD-AAC,
Apple AAC (iTunes Plus; Mastered for iTunes) (Mac only)

5.1: mp3 Surround, AAC-LC, HE-AAC & HD-AAC, MPEG Surround

2.1 Additional Codec Details

Please see section 10 for Codec specifications.

2.1.1 Apple AAC iTunes Plus (Mastered for iTunes) (Mac only)

The addition of this Apple AAC codec allows real-time audition of files produced under the Mastered for iTunes initiative that are destined for the iTunes store.

This codec is available for Mac computers only. It uses the same Apple codecs and re-samplers that are used for the current iTunes Catalogue. It very specifically re-samples to 44.1 kHz if necessary, and encodes as AAC-LC, 256kbps, variable bit rate and maximum quality, which are the settings used for the iTunes catalogue. For online audition, the signal is then resampled back to the host DAW sample rate if necessary.

This option is available only for stereo configuration. There are no options to change the iTunes codec settings because this would not be representative of the iTunes Plus standard. The Apple codecs and re-samplers are components of the Mac OSX operating system, so can be upgraded during an OSX update.

On Windows computers, the closest approximation to the iTunes Plus codec is to select Fraunhofer AAC-LC codec set to highest quality VBR at 256kbps.

2.1.2 MPEG Surround (Fraunhofer)

MPEG Surround is a feature-rich open standard compression technique for multi-channel audio signals. Operating on top of any core audio codec — including AAC, HE-AAC and MPEG Layer 2 — the system provides a comprehensive feature set, including highest surround and stereo audio quality, multi-channel support at stereo bit-rates and virtualized binaural multi-channel sound reproduction capability for stereo headphones.

2.1.3 HD-AAC (Fraunhofer)

The HD-AAC codec now supports 16-bit files in addition to 24-bit data.

Please see section 5.1.1 for more details.

2.1.4 AAC-LC (Fraunhofer)

The AAC-LC codec now supports a 96 kHz session sample rate and a VBR mode at 256kbps.

3. The Sonnox Fraunhofer Pro-Codec

3.1 The User Interface and Workflow Overview

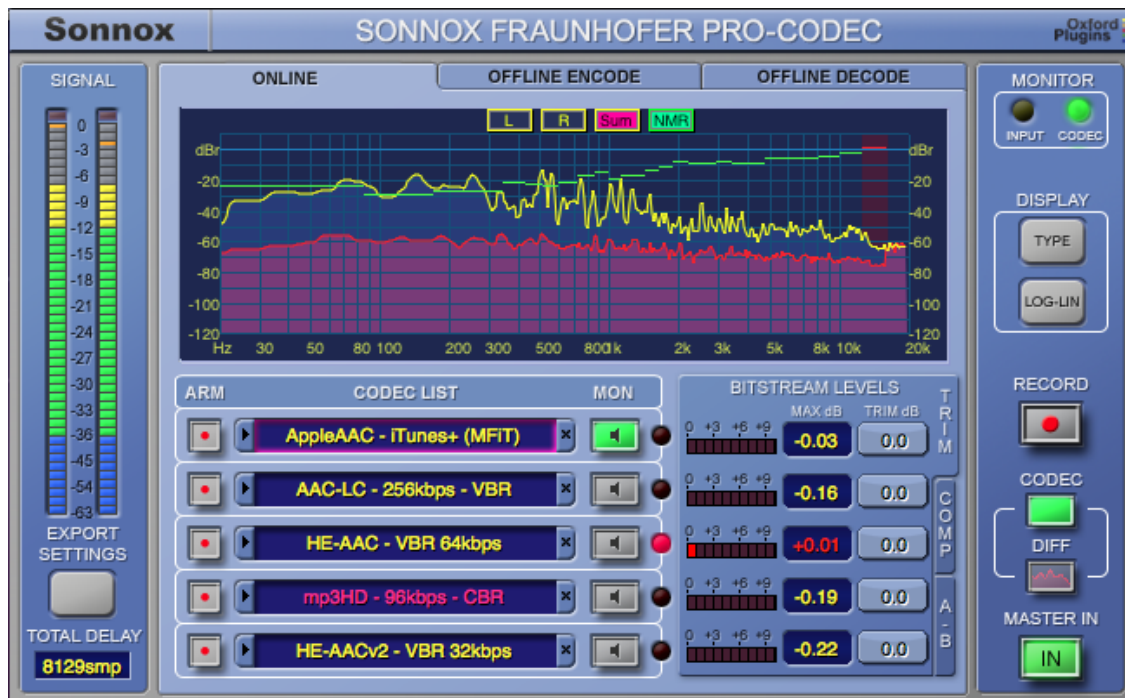
The primary purpose of the Sonnox Fraunhofer Pro-Codec is to provide a way of monitoring your host session material using the highly regarded Fraunhofer encoders, without the need for offline processing. By auditioning and comparing codecs online, it allows you to make executive decisions on the best codecs to use and, when all is complete, write the encoded audio to disk.

The latest version of this plug-in also includes access to the Apple AAC codec (Mac only), which allows users to directly audition audio that has been produced according to the ‘Mastered for iTunes’ format (also known as ‘iTunes Plus’).

Also contained within the plug-in is a fully featured offline encode window. You can compress multiple WAV or AIFF files for offline batch processing with the codecs and settings of your choice.

When you need to import an encoded file into your project, but the host does not recognise the file format, the plug-in also has an offline decode section. This eliminates the need for any proprietary applications to be used when importing encoded material into your session.

There are three main workflow modes: **ONLINE**, **OFFLINE ENCODE** and **OFFLINE DECODE**. The modes are accessible via tabs at the top of the main graphical user interface (GUI):



By default, the **ONLINE** mode is active, and the centre panel of the GUI provides a large FFT (Fast Fourier Transfer) display, showing the spectral content of the input signal in yellow, the **DIFF** signal in red (the difference between the input and output signals), and what we refer to as an **NMR** (Noise-to-Mask Ratio) display in green.

This display can be scrolled and zoomed in both frequency and dB.

For a stereo signal, you can select **LEFT**, **RIGHT** or the **SUM** of both left and right. The **NMR** signal can be turned off, if required.

Beneath the FFT window, the **CODEC LIST** enables the simultaneous selection of up to five codecs. The associated **MON** buttons allow monitoring of a selected codec. The **ARM** buttons are used when recording online to an encoded file.

To the right of the codec list are three tabbed windows; shown above is the **BITSTREAM LEVEL / TRIM** window. The **COMPRESSION** tab provides an indication of the data compression for the selected codecs, and the **A-B** tab allows A-B and A-B-X auditioning.

3.2 Input and Output Panels

The left hand panel of the plug-in features a large input signal meter. This provides 1dB per LED down to -36dB, and thereafter 3dB per LED. There is a peak hold display by default.

The **EXPORT SETTINGS** button allows access to a panel for configuration of the encoded file data (see below for more details).

A window at the base of this panel displays the **TOTAL DELAY** through the plug-in, and will indicate the maximum of each individual codec delay. Click in the window to toggle between milliseconds and samples. This field is immediately updated in the GUI when a new codec or other setting is selected. Depending on the host behaviour, a new value may not be available for delay compensation until the plug-in is removed and re-instantiated.

The output panel on the right of the GUI includes a **MASTER IN** button, used to switch glitch-free between the monitor output of the plug-in (usually a codec output) and the input signal. The signal you are **MONITORing** is indicated by two LEDs at the top of the panel, designated **INPUT** (yellow) and **CODEC** (green).

The listen section of the output panel defaults to the **CODEC** monitor output, but also contains an option to listen to the difference between the input and the processed output by selecting the **DIFF** button. This is particularly useful for monitoring the audio that is removed by the encoding process.

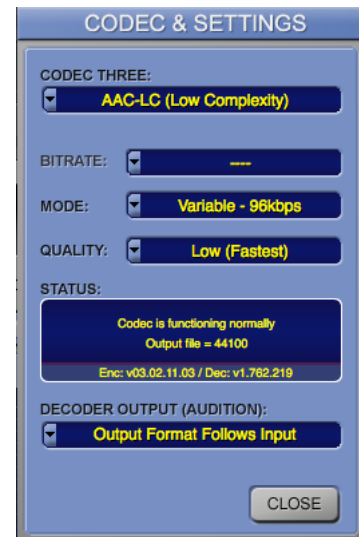
The output panel also contains buttons to control the **DISPLAY. TYPE** switches between the main FFT display and a Lissajous phase display, and **LOG-LIN** button switches between logarithmic and linear frequency displays.

A master **RECORD** enable will start recording to file for any codec that is **ARMed**.

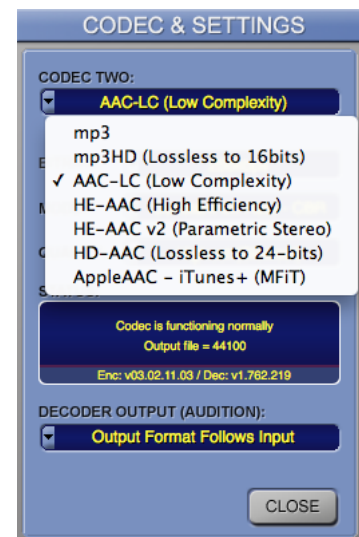
3.3 Online Auditioning and Configuring Codecs

The **ONLINE** encode tab enables Fraunhofer codecs to be auditioned, compared and the encoded audio written to disk, all while the host workstation is operating. It is the default section when the plug-in is instantiated.

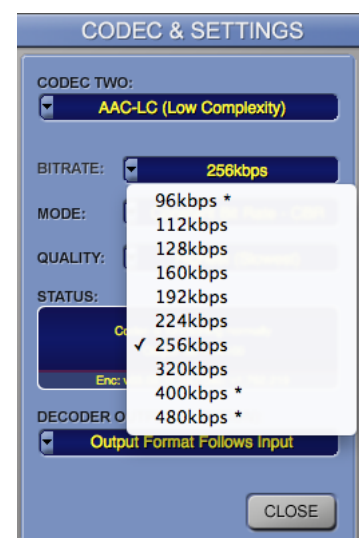
The default screen will open, displaying the online codec audition screen. To add a custom codec, click on any available slot in the **CODEC LIST**. This will bring up the **CODEC & SETTINGS** edit window (see right), enabling the selection of parameters specific to the selected codec.



The screenshot on the right shows the selection of available codecs.

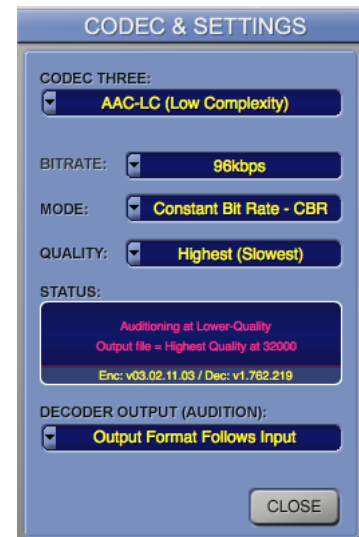


This screenshot shows the bit-rate selector. Some of the table entries are accompanied by an asterisk, indicating that the codec will be resampling the audio. This is conditional on the codec type, the session sample rate and the quality setting.



The **STATUS** window (see right) displays information about the codec, including its version number. In some cases (similar to the bit-rate table entries in the drop-down menu above), the text may be in red, indicating a special condition.

Some combinations of bit-rate, quality setting and sample-rate might cause an encoder to down-sample internally (especially at the highest quality settings). In this case, if the stream is decoded for auditioning, the number of samples per second produced will not match the session sample rate and there will be a pitch shift.



To work around this problem, when one of these special combinations is selected, the plug-in creates two versions of the codec; one operating at low quality for auditioning (because it does not internally down-sample at low quality), and one operating at high quality, which will be used only for writing to file, and only if the user drops into online record mode.

Red text in the **STATUS** window of the **CODEC & SETTINGS** pop-up indicates this condition.

At higher session sample rates, the **STATUS** window also informs the user if downsampling is occurring. In this case, the input audio is down-sampled prior to encoding, to bring the sample rate into the range supported by the codec. The decoded audition signal is correspondingly up-sampled.

NOTE: The encoded file's sample rate is derived from the session sample rate, the codec supported sample rate, the selected codec bit-rate, the selected codec quality setting and the channel configuration. Hence, it is not a simple process to choose a desired encoded file sample rate. For more details, see section **4.7 Codec and Encoded File Sample Rate**.

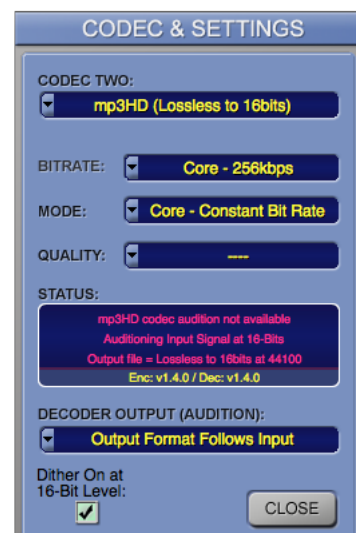
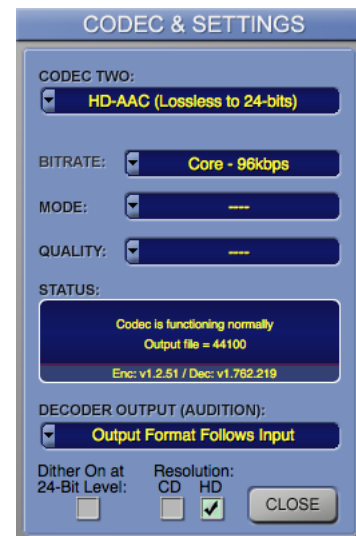
The **DECODER OUTPUT (AUDITION)** window determines the decoder format when playing back multichannel encoded signals.

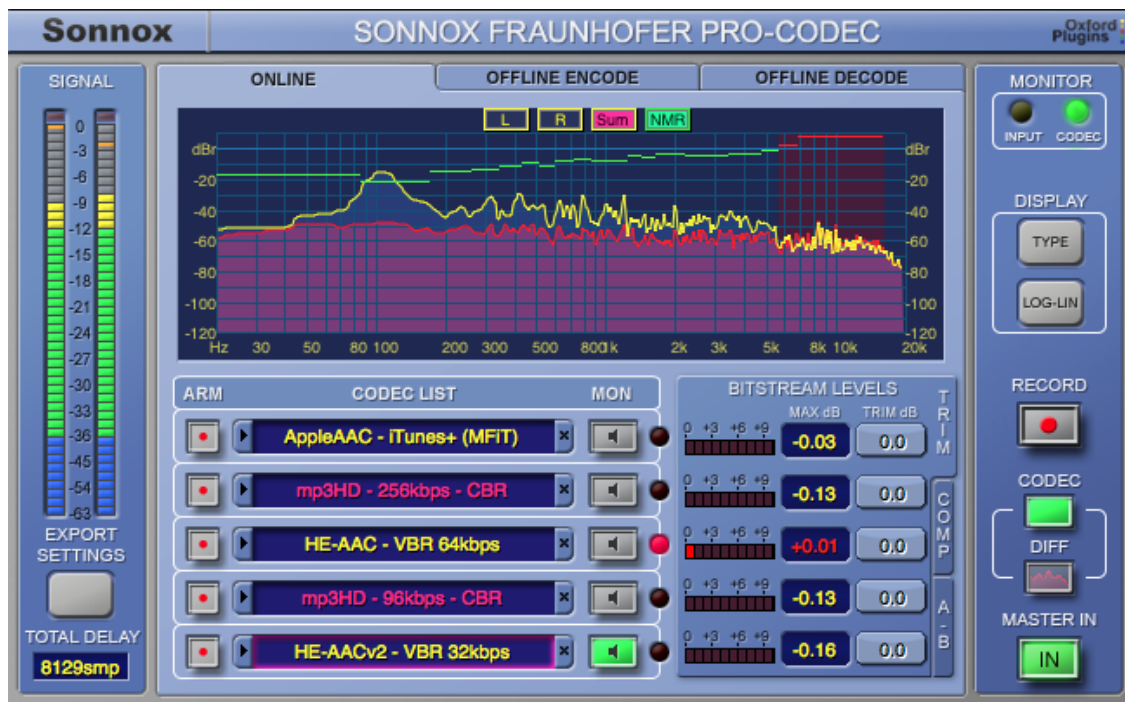
The two screenshots on the right show the codec configuration options for the two Lossless codecs. Note there are options available for the lossy cores.

Also notice that the **STATUS** of the mp3-HD codec has red text (see lower screenshot). This is another example of a special condition.

In this case, the user has been alerted to the fact that the audition signal has been forced to 16 bits (see section 5 on **Lossless Codecs** for more details).

Note the dither check boxes, only present for Lossless codecs.





In the example shown above, two of the **CODEC LIST** windows show red text, indicating that a special condition is present. Clicking in the relevant windows will produce the **CODEC & SETTINGS** window for more information, if required.

If you want to add a new codec that is similar to one that is already in the list, you can duplicate it by clicking an existing list element and dragging it to an empty row (Windows systems only).

To remove any of the codecs in the list, click the blue 'x' tab on the right hand side of the list element.

When you start playback through the host, by default you will hear the encoded audio (providing the **MASTER IN** button is selected).

To audition one of the other selected codecs, click on the **MON** button for any active codec in the list.

For glitchless bypass switching, you can toggle the **MASTER IN** button; note that this button is automatable. The monitor LEDs at the top right of the plug-in indicate **CODEC** or **INPUT** as the monitored signal. Clicking on either the input or codec monitor LEDs also provides glitchless switching between the two.

The FFT display shows the signal frequency content. The input signal is displayed in yellow, whereas the red line indicates the spectral content of the difference signal of the selected codec (ie. the difference between the input signal and the encoded/decoded output signal). This is referred to as the **DIFF** signal.

Alternating between the **CODEC** and **DIFF** buttons on the output section allows monitoring of the codec output or the audio that is removed by the encoding process.

The **LOG-LIN** button toggles the frequency scaling of the displayed graph between logarithmic and linear.

It is possible to both scroll and zoom parameters in the display area:

- For frequency, use the left mouse button to scroll left and right, and zoom up and down.
- For level (dBs), use the right mouse button to scroll up and down, and zoom left and right.

For a stereo signal the **LEFT**, **RIGHT** or **SUM** can be selected for display.

3.3.1 The NMR indicator

Also featured in the FFT display is the NMR indicator line, predominantly green in the screenshot on the previous page, which provides an indication of the frequency areas where the difference between the codec output and the original input might be audible. The NMR indicator can be disabled by clicking on the **NMR** button (above the display).

All lossy codecs will produce a very slightly different output from their input. Sometimes, if you listen to the difference signal it can sound significant; however, the very nature of a perceptual coder is that this difference signal should be inaudible (ie. masked by the output signal). You can choose to trade off more data compression (and smaller files) against increased audibility of artefacts and codec-induced noise. Theory states that this codec-induced noise should be inaudible when the NMR indicator is green. Under some circumstances (codec, frequency and input signal dependent) this **NMR** line turns red (see right hand side of FFT display on previous page), indicating that the induced noise is possibly audible. The training and sensitivity of the listener's ears is also an obvious and variable factor to take into account.

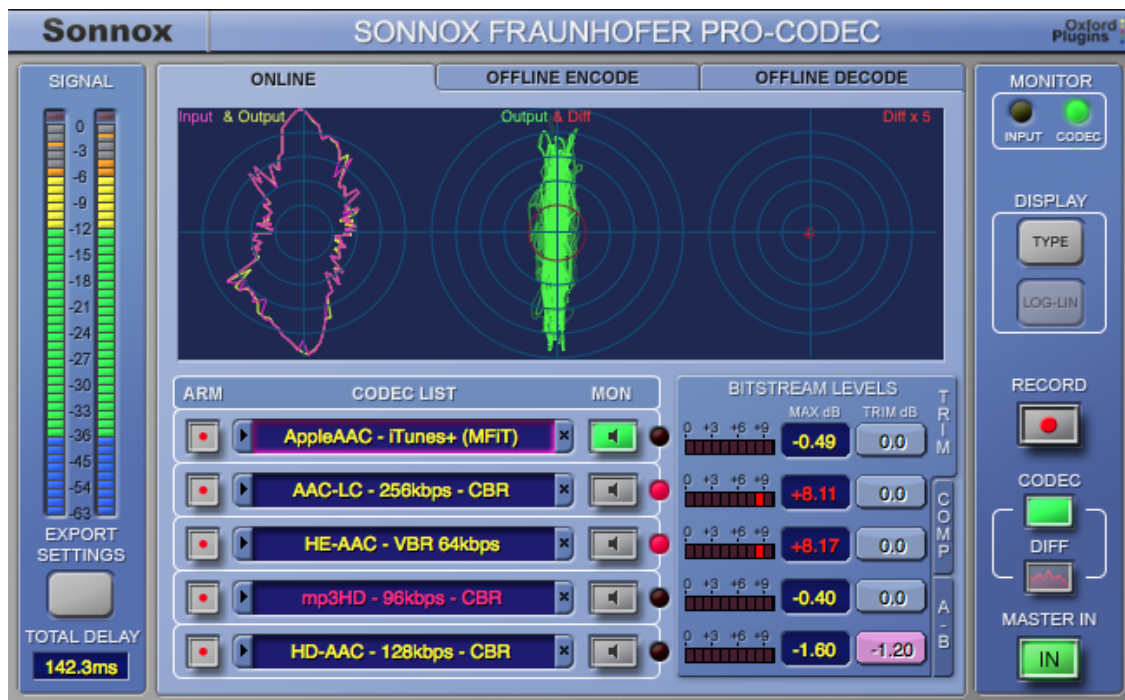
The NMR display is a useful indication when auditioning the mp3 or AAC codec families. If there are areas where the NMR line is above 0dB, then those are the frequency areas of interest. The display will highlight these frequency areas in red. It is possible, for example, to make very quick comparisons of a selection of different bitrates while auditioning music.

The NMR calculation is less accurate for parametric codecs (those that use enhancements such as Parametric Stereo or Spectral Band Replication). HE-AAC and HE-AACv2 use parametric enhancements to achieve very high compression ratios. The NMR indicator is enabled for these codecs by default, because it can still give an indication to the engineer of frequency areas that might require examination.

Note that the NMR indicator is not present if the selected codec is lossless (ie. mp3-HD or HD-AAC). Also, the NMR display is less accurate when **SUM** is selected; it is more accurate with single audio streams.

The difference signal used in the NMR calculation is absolute. For some low level signals the NMR might be relatively high, but any difference will be inaudible due to the low signal level.

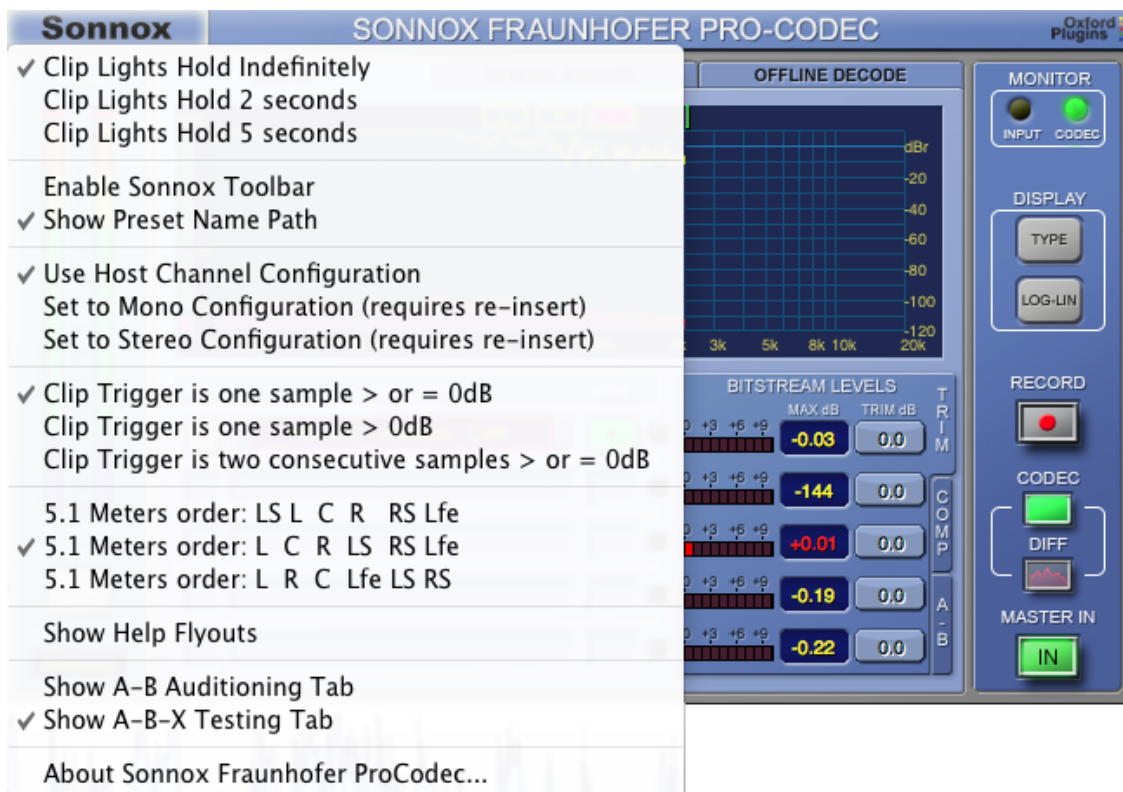
The NMR indicator is only valid for session sample rates 44.1, 48, 88.2, 96, 176.4 and 192kHz. In particular, it is not supported at 32kHz. The NMR display is not valid and is not displayed above 16kHz.



Clicking on the **DISPLAY TYPE** button on the right hand side of the plug-in will alternate between the FFT display and a phase display (see above). The centre of this display shows the left-right phase behaviour of the output of the codec in green and the phase behaviour of the difference signal in red. This difference phase is also displayed on the right, amplified by a factor of five. The display on the left shows the input and output envelopes.




If the feature is enabled, a help flyout panel can be displayed by hovering over a control or label (see above). Help Flyouts are enabled or disabled via the **Sonnox** button options menu:



This sub-menu also contains several other configuration controls for the plug-in, for example switching between the **A-B Audition** and **A-B-X Test** modes.

3.4 Bitstream Level (TRIM tab)

Any form of filtering can, under certain circumstances, *increase* the level of the signal. If your input is hot, this can potentially produce overloads when the signal is decoded. This effect can be very easily demonstrated with this plug-in.



The **BITSTREAM LEVELS TRIM** tab (see above) can be used to monitor this potential for overload. The post-decoder level is continually monitored, and a small meter indicates levels over 0dB. There is a peak-hold text display (**MAX dB**) of the maximum post-decoder level achieved, which is yellow if below 0dB and turns red on overload. Clicking on this display resets the overload condition.

The **TRIM dB** control allows you to trim the input level to the codec to accommodate these overloads. The gain reduction (headroom) is applied pre-encoding, and if non-zero the trim value is displayed on a pink background. In order to allow like-for-like signal level comparison and difference signal generation, the headroom is applied before encoding and removed after the decoder output. This way the encoded bitstream is reduced in level, but the auditioned signal remains at the input level.

In the screenshot above, the HE-AAC codec has generated an overload condition, and has had its input trimmed down by 0.90dB to accommodate this overload. The **TRIM dB** controls are automatable.

For convenient trimming, the value in the **MAX dB** window can be selected, dragged and dropped onto the **TRIM dB** cell. The **MAX dB** value can be repeatedly dropped onto the **TRIM** cell; each time the value is subtracted from the current **TRIM** value. In the screenshot above, there is a 0.30dB overload from a signal already reduced by 0.90dB. Selecting and dragging the value from **MAX dB** to **TRIM dB** will set the latter to -1.20dB.

This post-decode level overload monitoring is present even if the **TRIM** tab is deselected. There are permanently visible overload indicators to the right of each of the five codec selection windows. Clicking on an overload LED resets the overload condition for that codec.

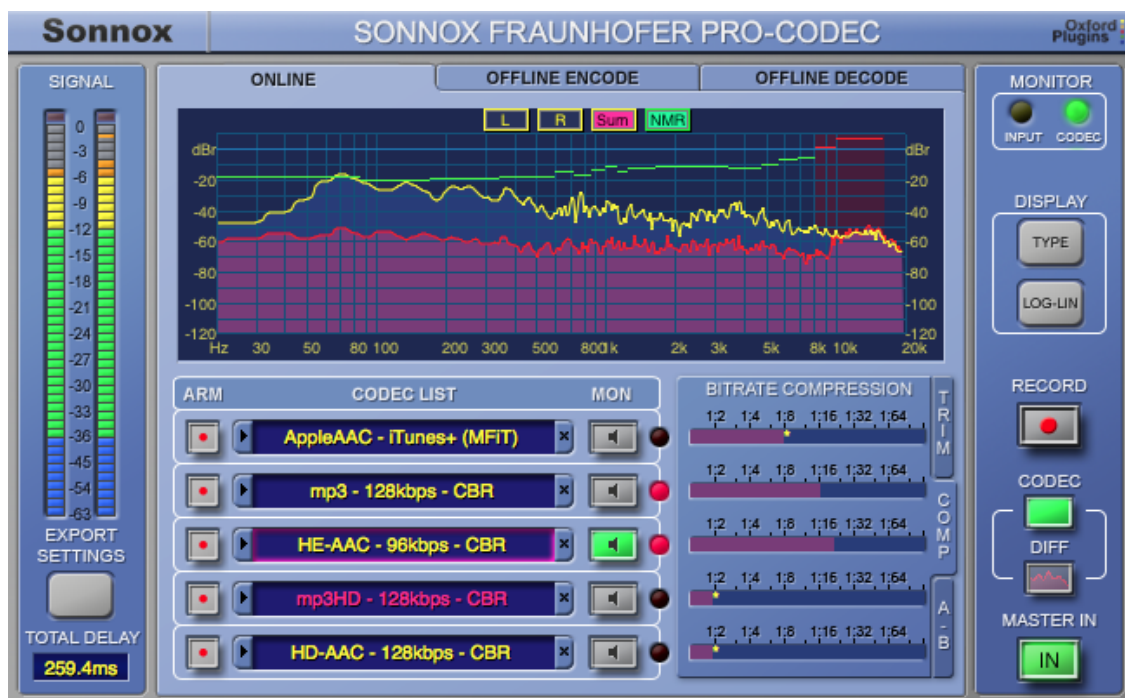
NOTE: See section **4.4 Bitstream Level Monitoring** for more details.

3.5 Data Compression Factor (COMP tab)

All of the available codecs compress an audio data bitstream. Different codecs compress using different algorithms, and will compress to very different degrees.

HE-AAC is optimized for the best quality at very low bit-rates, and will produce very small files (potentially around one hundredth of the input file size). Lossless encoders, in contrast, are not so effective at compressing data, and will produce files that are around half the size of the input. The familiar mp3 codecs produce compression factors of typically around 12.

Selecting the **COMP** tab (see below) gives a visual indication of the compression ratios pertinent to the selected codecs:



A yellow asterisk (if present) indicates that the figure is approximate. This will be the case for lossless encoders and variable bitrate (VBR) encoders.

Remember that using perceptual coding is always a trade-off between more accurate reproduction of the original on the one hand, against a higher compression ratio and smaller files on the other hand.

3.6 AB Auditioning and ABX Testing (A-B tab)

NOTE: The selection of AB Auditioning or ABX Testing is made via the **Sonnox** menu button.

3.6.1 AB Auditioning

For simple and quick auditioning use AB Mode. Control is via the **A-B** tab (below the **TRIM** and **COMP** tabs):



Clicking in the **A** or **B** columns will assign a codec to the **A** or **B** audition channels. Clicking in the **B** column adjacent to a selected **A** codec will automatically set the **B** column to the input signal (**I/P**). To revert to a codec for the **B** column, click on a cell non-adjacent to the **A** codec.

The **A-B** button will toggle between the codec selected in the **A** column and the codec (or input) selected in the **B** column. This is ideal for critical listening tests.

The **HIDE** button makes the current selection anonymous. The **A-B** button still toggles between the two column selections.

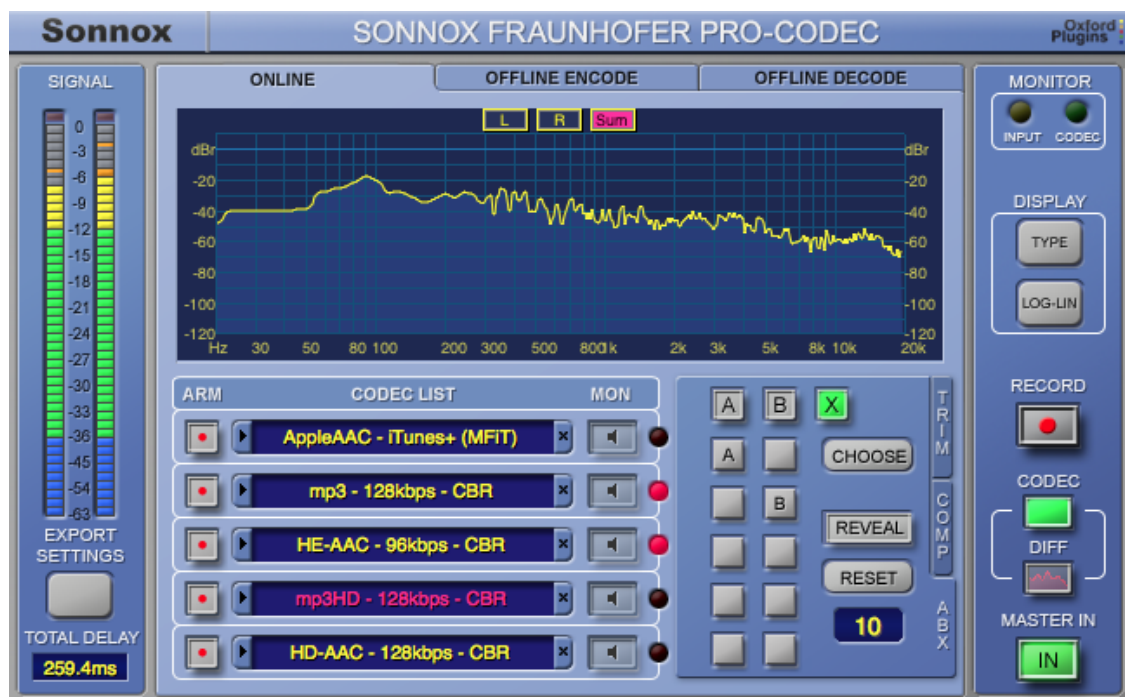
The **BOTH** button displays both of the two difference signals (the currently auditioning signal is coloured green).

3.6.2 ABX Testing

For more detailed and statistical listening tests, use ABX mode (for the principles behind ABX testing, there is a good Wikipedia entry). Select the ABX tab via the **Sonnox** menu button.

Unlike hardware ABX systems, there is no switch delay or switching glitches in this implementation.

The two **AB** assignment columns are used in a similar fashion to **AB** auditioning, but the column headers are now buttons. There is also an additional button labelled '**X**':



Although the three **A B X** buttons are like mechanical 'radio buttons', repeatedly pressing the same button will alternate between that button and the previously selected button. This means you can repeatedly alternate between **A & X** with your eyes closed (or **B & X**, or **A & B**) because you don't need to move the mouse pointer. Eyes-closed auditioning might improve your listening focus!

As with **AB** mode, there is access to the input signal (**I/P**) by clicking in the **B** column adjacent to a selected **A** codec. **RESET** will initialise the internal counters. The display window beneath the **RESET** button displays the number of completed test cycles.

X will have been set to a random selection of either **A** or **B**. Listen critically, switching between **A**, **B** and **X**; when you think you have identified **X** as either **A** or **B**, leave the selection on the identified **A** or **B**, and press **CHOOSE**. This will increment the number of test cycles but not display whether your choice was correct. The act of pressing **CHOOSE** also randomly re-allocates either **A** or **B** to **X**. For a thorough and statistically significant test, repeat this audition and **CHOOSE** cycle up to 20 or 25 times.

When you have completed a satisfactory number of cycles, press **REVEAL**, which will show the percentage number of correct guesses:




A 95% confidence level is regarded as statistically significant (so anything less than that means you have **not** reliably been able to distinguish between A and B — see above!). Note that a 95% confidence level is **NOT** the same as a 95% correct score! (Think of a single cycle; you might by chance guess correctly, and obtain a score of 100%, but with a confidence level of only 50%.)

As a quick guide, for a 95% confidence level you need a minimum of 8 correct guesses from 10 cycles (80%), 11 from 15 (73%), 14 from 20 (70%), or 17 from 25 (68%).

3.7 Online Encoding

To encode a file in real time, first you have to record enable the codec. This is done by selecting the **ARM** buttons to the left of the codec list. Multiple files can be selected for simultaneous recording.

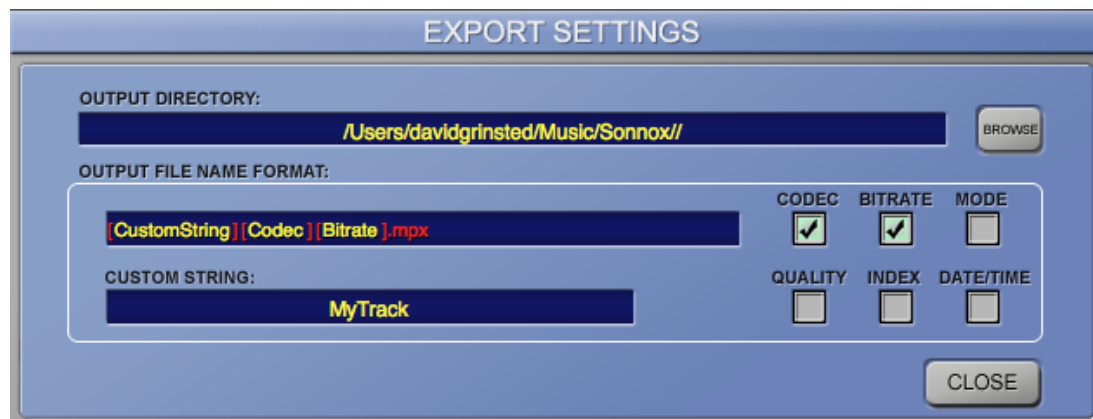


To record straight away, press the **RECORD** button when the host is playing back audio; this will write the output of your selected codecs to disk. You cannot change the **RECORD** **ARM** settings while dropped into **RECORD**:



You can also **ARM** the codec when playback is stopped, so that codecs will start writing the instant you start host playback. Likewise, stopping playback will terminate the writing procedure. Record duration and exact data compression ratios are displayed (see screenshot above).

The **EXPORT SETTINGS** window (select the button in the lower left of the plug-in) allows the user to specify the destination folder for the recorded files and the output filename:



The filename can be composed of elements including codec type, bit-rate and date of recording. It is possible to re-order these elements to form a flexible and informative filename. This can be done by clicking on any element within the complete string and dragging it either to the left or the right to a suitable position.

To create a custom string within the list, edit the **CUSTOM STRING** text box underneath the complete string window. Should you want to add or remove any codec data information from the string, there are check boxes for the **CODEC**, **BITRATE**, **QUALITY**, **DATE/TIME**, **INDEX** and **MODE**.

NOTE: Care should be taken to make filenames distinct, to avoid confusion for files written concurrently. For example, if five mp3 codecs were selected with different bitrates, it would be wise to enable the **BITRATE** option for the resultant filenames. The **INDEX** option is useful to make filenames unique (**INDEX** indicates the codec window number). Furthermore, if repeated renders are likely, the **DATE/TIME** can be used to distinguish subsequent actions (**DATE/TIME** includes a field for seconds).

The **CUSTOM STRING** is restored on a project/session load. When loading Presets, and if you delete the plug-in and re-insert it, the **CUSTOM STRING** will default to ‘My Track’.

The **OUTPUT DIRECTORY** is restored on a project/session load. Otherwise, it remains unaltered, even when loading Presets, and even if you delete the plug-in and re-insert it.

See section 7 of this manual for more details regarding Presets and session loading.

There is limited provision for ID3 metadata. The plug-in inserts the following metadata into an mp3 stream:

- Date and time

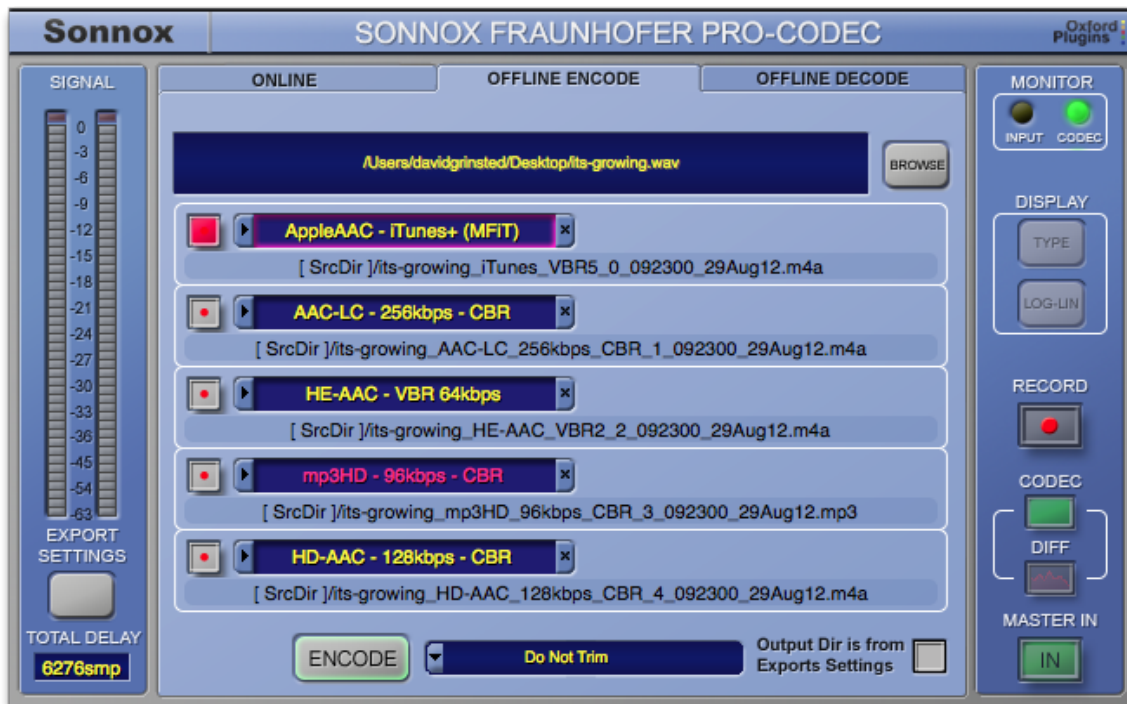
- Codec type, bitrate and other codec settings

- Encoded by manufacturer message (‘Sonnox Fraunhofer Pro-Codec Plug-in’)

Currently, AAC metadata (eg. for iTunes) is not supported.

3.8 Offline Encoding

It is possible to encode files offline. To do this, select the **OFFLINE ENCODE** tab:



Select **BROWSE** and navigate to a desired file. You can also drag and drop chosen files onto the window to the left of the **BROWSE** button.

By default, the output file(s) will be written to the same directory as the source file. Selecting the **Output Dir is from Export Settings** check box will force the output file(s) to be written to the directory specified in the **EXPORT SETTINGS** window (see previous section).

You can select Trim behaviour via the drop-down selector next to the **ENCODE** button. The choices are:

- Do not trim.

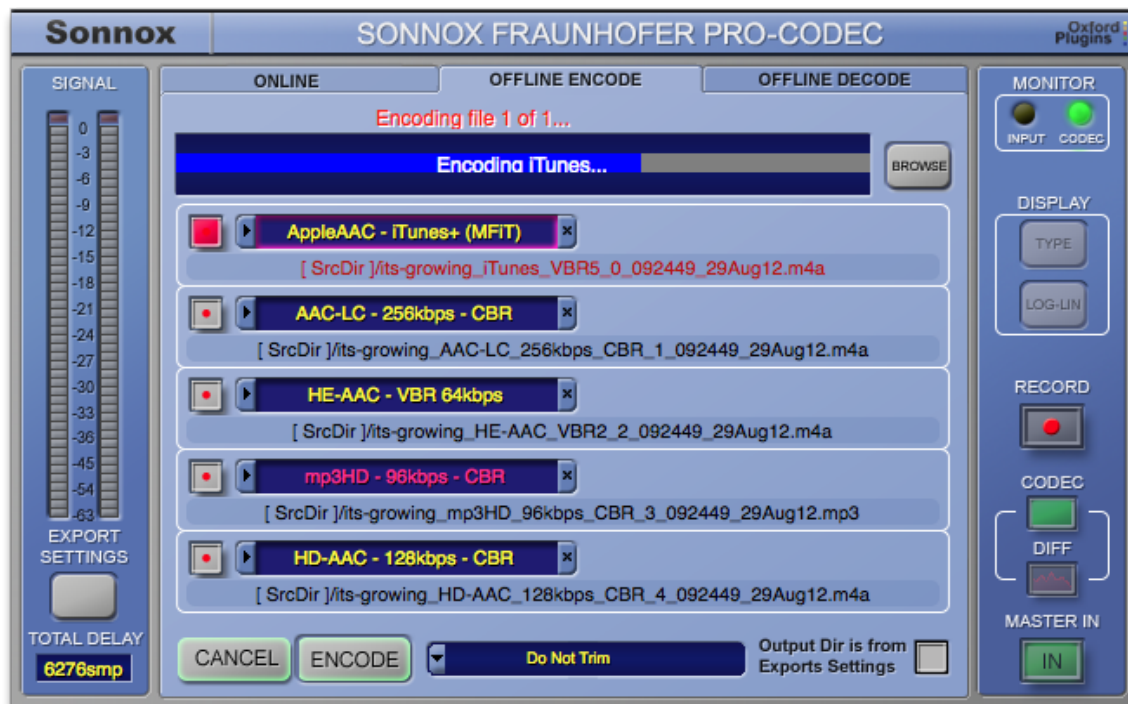
- Use online trim values.

- For each codec, prevent decoded levels from clipping by pre-processing the track.

- For each codec, normalise the level by pre-processing the track.

Write enable each codec in the list by selecting its 'arm' button on the left: when armed, the buttons will start to flash.

Each list element shows its respective output name and file location. Click **EXPORT SETTINGS** to edit it.



Once ready, clicking **ENCODE** will batch process all write enabled codecs (see above).

Only one source file can be processed at a time, but it can be written to multiple output files as a batch job with different codec settings.

NOTE: Offline encoding is only possible for files with a sample rate the same as the current session.

NOTE: Offline encoding only supports interleaved stereo files (not multi-mono _L and _R pairs).

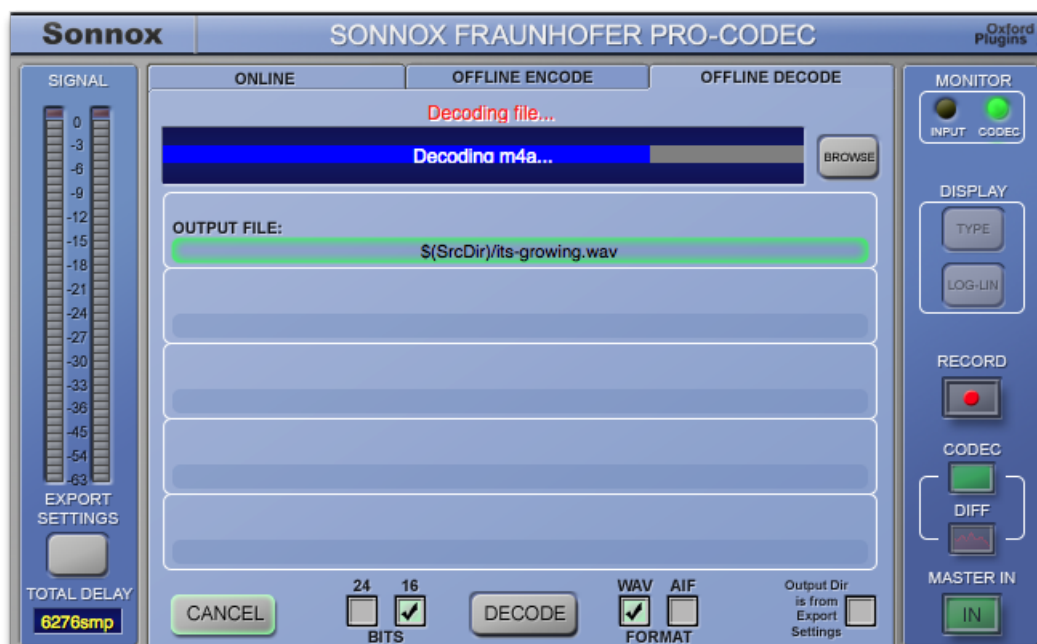
3.9 Offline Decoding

The Sonnox Fraunhofer Pro-Codec can also decode compressed files for import into the host sequencer (or for general purpose decoding). To do this, select the **OFFLINE DECODE** tab:



Click the **BROWSE** button and navigate to a compatible encoded file. You can also drag and drop a file into the window to the left of the **BROWSE** button. Output files can be either wav or AIFF, and have a bit depth of either 16 or 24 bits. Stereo files are interleaved only.

To decode the selected file, select **DECODE**. A progress bar will appear to indicate the status of the decode process:



Once decoding has finished, the filename text reverts from red to black, and the output filename box shows a green highlight. This file can then be dragged directly to your host sequencer (Windows only).

4. Additional Information on Workflows

4.1 Processing Sequence

The concept underlying the development of this plug-in is was that it should be the very last item in the mastering chain. A 16 or 24-bit master has been finalised and possibly rendered to a wav file. What is now required is to take that finalised master and produce an encoded, data-compressed version (or versions) for end-user playback.

In other words, the mix would normally be fed via a limiter to control peaks, and then dither would be applied before the data is truncated and written to the 16 or 24-bit wav master.

Again, under normal circumstances, this wav file would be passed via an encoder to derive a final mp3 or AAC file for end-user distribution.

However, please note that this final step is not necessarily carried out currently by the mastering engineer. Frequently, only the wav file (or losslessly encoded audio) is sent to the distributor or aggregator. This means that the mastering engineer has no control over which encoders are used, which encoder parameters are used, or whether the final encoded file will clip and/or distort when played back on a consumer device. Also, there is little point in producing a mix optimised for a particular codec or bitrate.

The mastering engineer can now achieve this control by using the Pro-Codec plug-in. Each mix can be auditioned exactly as the end-user will hear it; and consequently the mix can be tuned appropriately by the mastering engineer for a selected codec.

It is important to note that the Pro-Codec plug-in should be used after the limiter-dither-truncate process. If a dithered master wav file is required, it should be rendered with the plug-in monitor set to input or with the plug-in bypassed.

4.2 Overloading Post Codec (Audition Level Matching)

As discussed above, any form of filtering (including encoding and decoding) can increase the peak level of a signal. Some codecs such as the HE-AAC can add several dBs, hence the provision of Bitstream Level Trim to avoid post-decoder clipping in the audition chain. Up to five codecs can be selected for audition, and each could require a different trim setting. To keep the audition process palatable, the trim is reversed on each channel before the audition point, so that switching between codecs does not introduce audition level changes.

However, this means that you are not auditioning the true level at the input of the encoder. If there are significant trim values involved, it is possible that the encoded files will demonstrate differing average levels when played back, and this may be undesirable.

The Pro-Codec plug-in provides the ability to observe and monitor bitstream level behaviour, and provides the tools for the user to mitigate this behaviour. It is convenient to monitor many codecs simultaneously for comparison, but it might be best practice to treat each codec individually when producing a final master. In this way any level change can be kept in the master mix, rather than in the plug-in.

4.3 Overloading Pre Codec (Codec Headroom)

It is not desirable, but it is possible, that the plug-in will be fed by an overloading input (this is separate to the possibility that a lossy codec could overload from a non-overloading input). The input meters will show clipping.

If an overloading input is presented to the lossy codecs such as AAC-LC, the codec output will endeavour to follow the input as accurately as possible, using internal headroom in both the encoder and the decoder. The Bitstream Level max dB will accurately indicate the maximum level of the signal, and the monitor output will not clip in the digital domain.

This is not true with the lossless codecs, which do not have internal headroom. This applies to both the mp3-HD and the HD-AAC codecs.

In this case, monitoring the lossless codec can indicate the presence of a significant difference signal under some circumstances, but the Bitstream Level metering will read zero. In this instance, the ‘lossless’ output will be clipping. This situation can be mitigated using the Bitstream Level trim; of course it is always better practice to avoid an overloading input.

In contrast, a clipped wav file (recorded with overloads) will be replicated losslessly; and if the lossy core of a lossless codec overloads from a non-overloading input, the resulting file will also be lossless.

4.4 Bitstream Level Monitoring

Section 3.4 of this manual describes how any form of filtering can, under certain circumstances, *increase* the level of the signal. If your input is hot, this can potentially produce overloads when the signal is decoded. The section goes on to describe how the **TRIM** tab can be used in real time to monitor these potential overloads, and discusses the tools provided to alleviate the problem.

It is important to note that the indicated Bitstream metering and overload LEDs are provided *as a guide only*. If a segment of audio is being cycled, these maximum values are not exactly repeatable from one cycle to the next. Furthermore, in some host applications, repeatedly rewinding to the beginning of the track will also not result in repeatable maximum values. The discrepancies in both cases should be minor, and not all audio and not all codecs will suffer this. However, it is not always possible to trim exactly while auditioning, because the subsequent pass to online render might have a slightly different maximum value.

There is a highly technical reason for this, stemming from the difficulty of synchronizing the start point of the buffer that is used to input data to the codecs.

In some situations this will be important and, for those occasions, we recommend using offline encoding. Offline encoding is exactly repeatable however many times the file is encoded (because it is a file-based process, not a stream-based process).

Offline encoding has four headroom recording options:

- Do not trim.
- Use online trim values.
- For each codec, prevent decoded levels from clipping by pre-processing the track.
- For each codec, normalise the level by pre-processing the track.

If a track/encoder combination causes non-repeatable overload problems, and yet the desire is to encode as hot as possible, the third option can be used. This will pre-pass the file, calculate the maximum overload, and trim accordingly. As this is a file-based operation, this will give a repeatable result.

A fourth option is included, which will pre-pass then normalise the track. This can be used when a track is at a relatively low level, to boost the level just short of overloading.

4.5 Down-sampling Issues when Using Higher Sample Rates

There are several circumstances when the plug-in internally down-samples. One example is the auditioning of certain codecs, possibly because the session sample rate is higher than the highest sample rate that the codec supports. In this case, the audio is down-sampled for the codec, and then up-sampled for audition. In this case, a down-sampled signal is available in the codec for writing to a file in online encode mode.

In offline mode this is not the case. The user cannot offline encode a wav whose channel configuration and sample-rate is not a match for the codec settings in the plug-in. This is because of possible indeterminate behaviour if the configurations are not compatible. This is particularly important if you are working with a higher sample rate session and choose to offline encode.

If you are working at higher sample rates, the options are exemplified by:

1. Master at 88.2kHz and produce a gold master, then audition codecs at 88.2kHz, and online encode at 88.2kHz.
2. Master at 88.2kHz and produce a gold master. Downsample gold master to CD mix at 44.1kHz. Load CD mix session, audition codecs, create wav file, and offline encode.

4.6 Online and Offline Encoding

Following the above discussion on bitstream overload repeatability and downsampling, we can now discuss the advantages and disadvantages of online and offline encoding.

4.6.1 Summary of the Differences Between Online and Offline Encoding

Offline Encoding is intended to be used primarily as an alternative to Online Encoding.

The advantages of Offline Encoding are that:

- a) The bitstream overloads are consistent,
- b) The original filelength features of the files can be employed (meaning consistent start and end points).

In other words, the offline encode mode is intended to be used in the same session as when any codec-specific mix adjustments are made. It is not currently intended for large scale batch processing.

The disadvantages of offline encoding are that:

- a) An extra step in the workflow is required (compared to online encode, where the coded files are created during the bounce),
- b) There is no automatic downsampling for offline encode, so the higher sample rate master has to be bounced and downsampled first before working on codec-specific enhancements. With online encoding, the mix is automatically downsampled, and a codec file can be generated directly from this downsampled signal.

One disadvantage of online encoding is increased CPU loading. Many of these codecs are processor-intensive and a powerful computer is required for best performance of this plug-in. On-line encoding does not necessarily add DSP processing as the encoders are already in action; however, additional data buffering and disc writing will add significant overhead, and might cause CPU overloads.

4.6.2 When to use Offline Encode rather than Online Encode?

In general, offline mode is designed to be used when a pre-rendered wav file exists. If there is no wav file (perhaps instead there is a work-in-progress mix) online mode is appropriate. This would be the usual case when auditioning codecs.

It is possible with some hosts to render the mix to wav at the same time as online encoding; the codec might be inserted into an ‘aux’ channel for this. If the plug-in is in the master output channel, you will almost certainly not want to render the plug-in monitor output. Put the plug-in into bypass to render to wav, then either offline encode or use a second pass to online encode.

If the targets are both a compensated wav and a bitstream for a single codec, then it’s best to trim the input level (not the bitstream trim level) and use online encode while bouncing down to wav. An alternative here would be to additionally encode to a lossless codec rather than wav.

If many different codec files are required with no optimisation for any given codec/bitrate/quality combination, then offline encode is appropriate from a rendered master wav. To ensure the bitstreams are perfectly encoded to avoid decoded bitstream overloads, use the Offline Encode tab with either the ‘Normalise Decoded Level’ or ‘Prevent Decoded Levels > 0dB’ options (these options are not available in online mode).

Finally, on-line encoding should be avoided if many processor-intensive codecs are selected.

4.7 Codec and Encoded File Sample Rates

The sample rate of an encoded file is derived from the session sample rate, the codec supported sample rates, the selected codec bitrate, the selected codec quality setting and the channel configuration. Hence, it is not a simple process to choose a desired encoded file sample rate. The codec sample rate is shown in the Offline Encode **CODEC & SETTINGS** status window.

The following tables give examples of the encoded file sample rate that result from various selection conditions (full tables can be found in Section 10 of this manual — **Specifications**). The codecs have an internal ‘preferred sample rate’, which will give the highest quality reproduction (but, quite frequently, higher CPU usage).

In general, if the highest quality setting is selected for a codec, then the encoded file will have the sample rate as listed in the ‘Preferred Sample Rate’ column (matching the session sample rate, if this option is available).

In all other cases of quality setting, a codec uses the same sample rate as the session, except at higher sample rate sessions when the codec will use the highest simple division of the session sample rate that the codec supports.

4.7.1 Mp3 stereo, session 44.1kHz or 88.2kHz

Bitrate (kbps)	Supported Sample Rates (kHz)	Preferred Sample Rate (kHz)	Quality Setting	Actual Sample Rate Used (kHz)
96	32 / 44.1	32	Lower	44.1
96	32 / 44.1	32	Highest	32
128	32 / 44.1 / 48	44.1	Lower	44.1
128	32 / 44.1 / 48	44.1	Highest	44.1
320	32 / 44.1 / 48	44.1 / 48	Lower	44.1
320	32 / 44.1 / 48	44.1 / 48	Highest	44.1

4.7.2 Mp3 stereo, session 48kHz or 96kHz (96kbps not supported at this sample rate)

Bitrate (kbps)	Supported Sample Rates (kHz)	Preferred Sample Rate (kHz)	Quality Setting	Actual Sample Rate Used (kHz)
96	32 / 44.1	32	Lower	n/a
96	32 / 44.1	32	Highest	n/a
128	32 / 44.1 / 48	44.1	Lower	48
128	32 / 44.1 / 48	44.1	Highest	44.1
320	32 / 44.1 / 48	44.1 / 48	Lower	48
320	32 / 44.1 / 48	44.1 / 48	Highest	48

4.8 Codec and Session Sample Rate

The plug-in and many of the codecs are sensitive to the session sample rate and channel configuration. The plug-in supports a wide variety of host DAW applications, and understandably many of these applications provide differing amounts of session configuration settings to a plug-in. Some applications only present session configuration information to a plug-in when there is an active audio data stream (ie. the transport is set to play).

As a result of this, care must be taken when:

- selecting codec settings after switching a session sample rate or channel configuration
- or*
- when loading a Preset that was saved from a session that used a different configuration.

Session loads/restores in Pro Tools, Logic, Cubase and Nuendo are immune from this problem.

It is possible in some applications that a newly instantiated Pro-Codec has not been provided with current session data. If this is the case, when a codec configuration is selected (manually or via a Preset load) that is incompatible with the real session configuration, then the plug-in could receive incompatible data when PLAY is pressed. In all cases, the plug-in will endeavour to match the requested codec settings, and where possible will select the closest compatible settings. By preference, the substitute codec will be of the same type, but using the closest compatible bitrate. Where the codec doesn't support the sample rate at any bitrate, the codec will default to AAC-LC at the closest bitrate.

To illustrate this behaviour, consider a stereo session at 48kHz with three codecs selected. If the session configuration is changed to mono at 32kHz, the plug-in will make the following codec selections:

	Selected codecs Session Stereo 48kHz	Substituted codecs Session Mono 32kHz	Reason
Codec 1	HE-AAC 128 kbps	HE-AAC 56 kbps	HE-AAC 128 kbps not supported at 32kHz
Codec 2	HE-AACv2 56 kbps	AAC-LC 56 kbps	HE-AACv2 is stereo only
Codec 3	HD-AAC 160 kbps	AAC-LC 96 kbps	HD-AAC not supported at 32kHz

To avoid these problems after changing the session configuration, it is advisable when using affected host applications to press PLAY, then STOP, before instantiating the plug-in or loading a plug-in Preset. This will ensure that session sample rate and channel configuration data are flushed through the plug-in correctly.

It is also advisable to save Presets for the plug-in at a variety of session sample rates and channel configurations.

To illustrate this further, consider the following examples. WaveLab7 is used in these examples because it has the powerful ability to easily select and play files with different sample rates. Several other applications exhibit the same behaviour.

4.8.1 Example 1. Different session sample rates and compatible codec bitrates

WaveLab 7 is an example of an application that only provides session configuration information to a plug-in when there is an active audio data stream (ie. the transport is set to play).

Using WaveLab 7, make three mono 24-bit wav files available via the tabbed windows in the audio file workspace. One is 32kHz, and is currently selected and playing. The unselected tabs have a 44.1kHz wav file and a 22.05kHz wav file available.

A Sonnox Fraunhofer Pro-Codec is instantiated and shows the default codec as AAC-LC 192 kbps. Note the usual default of 256 kbps is not valid for a mono wav file (see section 10.2 for the full table of codec supported sample rates).

If the 22.05kHz file is selected, and play stopped and re-started, the codec will automatically select 28 kbps (which is the closest to 192 kbps that the codec supports at 22.05kHz).

If the 44.1kHz file is now selected, and play stopped and re-started, the codec will automatically select 56 kbps (the closest to 28 kbps that the codec supports at 44.1kHz).

4.8.2 Example 2. Different session sample rates and incompatible codec bitrates

Using WaveLab 7, set up two wav files in the tabbed windows in the audio file workspace. One is stereo 24-bit 32kHz, and is currently selected and playing. The unselected tab has a stereo 24-bit 48kHz wav file associated with it.

A Sonnox Fraunhofer Pro-Codec is instantiated. As the HD-AAC codec does not support 32kHz, the HD-AAC option does not appear in the codec selection window.

The transport is stopped, and the 48kHz file is selected. The user now tries to select the HD-AAC codec, but it is still not available as an option. After the transport is set to play, the codec can be changed to HD-AAC. The file plays normally and the plug-in and codec perform normally.

The transport is stopped, and the 32kHz file is re-selected. Now, if the transport is set to play, the codec will receive invalid audio data, and the plug-in will substitute an AAC-LC codec.

Quit and re-start WaveLab 7. Select a 32kHz wav but do not press PLAY. Instantiate the Pro-Codec. At this moment the plug-in has not picked up the correct session data because WaveLab has not been set to play. The plug-in has been defaulted to 44.1kHz. Note it is possible now to select HD-AAC but, if this is done, the plug-in will substitute an AAC-LC when PLAY is pressed (as above).

Now try to load a Preset saved from a 32kHz session. A warning message indicates that there is an incompatible sample rate (the plug-in thinks it is 44.1kHz). To progress, press PLAY, then STOP, and then load the Preset.

4.8.3 Summary

These complications arise mostly when low sample rates (32kHz) are combined with low bitrate codecs; or when high sample rates are combined with high bitrates.

Generally Pro Tools, Logic, Cubase, and Nuendo are among the hosts that are immune to this behaviour, but Sequoia, Samplitude and WaveLab are examples of hosts that can exhibit these issues.

To avoid these problems after changing the session configuration, it is advisable when using affected host applications to press PLAY, then STOP, before instantiating the plug-in or loading a plug-in Preset. This will ensure that session sample rate and channel configuration data are flushed through the plug-in correctly.

It is also advisable to save Presets for the plug-in at a variety of session sample rates and channel configurations. It is advantageous to categorise the Presets by channel configuration and session sample rate.

5. Lossless Codecs (mp3-HD and HD-AAC)

A lossless codec is intended to produce an output data stream that is an exact duplicate of an input audio data stream. In the simple case of a 24-bit wav file as input, the output of the HD-AAC lossless codec will be a wav file with identical audio data. The behaviour of the mp3-HD lossless codec is the same when presented with a 16-bit wav file as input.

When these lossless codecs are used in a DAW, the situation is slightly more complicated because usually the input is not a simple 16 or 24-bit input file. Most DAWs use floating point processing, and so do most plug-ins. A floating point datum will have many more states than an audio integer as taken from a 24-bit wav file (this is another way of saying that the floating point format has higher resolution than the fixed-point wav format). Whenever processing is performed in a DAW or plug-in, the processing is performed at the higher resolution of the floating point system, and the resulting audio data cannot be represented perfectly on a 24-bit scale any more. The data stream will contain floating point values between the quantised values of the 24-bit scale, and these in-between values are called ‘inters’. This interactivity can be readily seen using the **DIFF** signal, apparent up to -144dB for HD-AAC and up to -96dB for mp3-HD. (One of the main reasons the **DIFF** signal is included is to aid visibility of this situation.)

The lossless codecs will not reproduce these inter values as they only use a 16 or 24-bit scale, and the correct terminology is ‘lossless to 24 bits’ for HD-AAC, or ‘lossless to 16 bits’ for the mp3-HD.

Therefore, if the audio channel includes processing, it is important to present correctly truncated data to the input of the lossless encoder. When an audio data stream is truncated, it must first be dithered at the correct level. More on this below.

5.1 Auditioning Lossless Codecs

The two lossless codecs from Fraunhofer have a very clever feature; the single compressed lossless file includes a lossy core channel. It therefore acts as a lossless archival format, a lossless distribution format for the masters, *and* a final playback format for both lossless and lossy decoders — and all of this in a single file.

In the case of HD-AAC playback, for example, if the decoder in your player has full HD-AAC capability, you will be able to listen to perfect replication of the original wav file. However, the same HD-AAC file will still play through a decoder that doesn't have HD-AAC capability, and instead you will be listening to the lossy embedded AAC channel.

5.1.1 HD-AAC (24-bit / 16-bit lossless)

It is possible to audition in real time the HD-AAC codec. If the HD-AAC codec is selected, and its **MON** button enabled, the audition point is post the HD-AAC encoder-decoder.

You will notice that the **DIFF** signal in the FFT display will drop below -144dB. (HD-AAC is lossless to 24 bits, so the **DIFF** signal has an amplitude smaller than 24 bits down. If the input signal to the codec is quantised to 24 bits, either by dither and truncation, or because it is unprocessed data from a 24-bit input file, the diff signal will drop below -180dB).

While encoding/auditioning an HD-AAC signal, attention should be paid to the lossy core. We recommend selecting the appropriate core codec in an adjacent codec window, which will allow correct auditioning of the lossy codec in the HD-AAC channel, and will also allow correct overload monitoring and mitigation.

The HD-AAC lossy core takes the parameters:

Bitrate	Selectable from 96kbps to 192kbps
Mode	CBR
Quality	Medium

However, remember that if the **BITSTREAM LEVELS TRIM** function is used to mitigate the core codec overload, and if that same trim value is applied to the HD-AAC encoder input, then *both* the lossless and the lossy data streams will be trimmed.

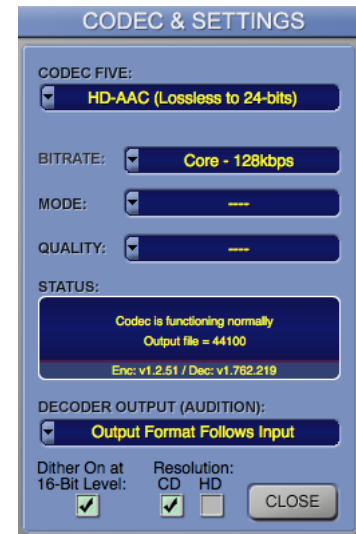
Also note that, under normal circumstances, if the lossless codec's Bitstream Level Trim is non-zero to avoid clipping in the lossy chain for an offline encode that uses the online trim values, then the resulting decoded output file will not be lossless because the level has changed.

5.1.1.1 HD-AAC 16/24-bit modes

The HD-AAC codec has a 16-bit mode that is selectable in the **CODEC & SETTINGS** window by selecting CD resolution rather than 24-bit HD resolution

For online HD-AAC Operation there are three volatile buttons that appear at the bottom of the Codec Settings window for an HD-AAC codec (see right):

1. A button to enable dithering at whichever resolution is selected by the other buttons.
2. A **CD** button to select 16-bit resolution
3. An **HD** button to select 24-bit resolution.



Together, these buttons allow the selection of how much resolution to use for the HD-AAC encoding, and whether to dither before truncation to that resolution.

These settings are reflected in the bit meter operation when viewing HD-AAC codecs, and they are reflected by all HD-AAC codecs during online DSP. The settings are saved into presets and project settings.

During offline encoding, the resolution used for the writing follows the PCM width of the source file. You cannot reduce the resolution as part of the offline HD-AAC encoding process.


For offline decoding, the HD-AAC decoder detects the original source file PCM width that was used by the HD-AAC encoder. There is also a general-purpose 16/24-bit output file selector on the **OFFLINE DECODE** tab.

If the switch is set to 24 bits, the output file PCM width will follow the original file PCM width: 16 -> 16 or 24 -> 24. This is the recommended (and default) setting. If the switch is set to 16 bits, then the decoded output PCM width will always be 16 bits.

5.1.2 mp3-HD (16-bit lossless)

Currently it is *not* possible to audition in real time the mp3-HD codec. (The codec library itself includes writes-to-file that cannot be used in the DSP thread of a plug-in.) If the mp3-HD codec is selected, and its **MON** button enabled, the audition path is switched to a synthesised version of the codec output. This synthesised signal mimics the behaviour of the mp3-HD codec, in that it will replicate the input but ensure that the signal is only 16 bits deep, whatever the input word size. The audition path might include dither (see below).

The **MONITOR** led (see right) indicates **INPUT 16**, to indicate that the signal is a modified version of the input that is intended to synthesise the codec output.



You will notice that the **DIFF** signal in the FFT display will drop below -96dB. mp3-HD is lossless to 16 bits, so the **DIFF** signal has an amplitude smaller than 16 bits down. If the input signal to the codec is quantised to 16 bits, either by dither and truncation, or because it is unprocessed data from a 16-bit input file, the **DIFF** signal will drop below -180dB.

The mp3-HD codec is lossless to 16 bits only. With 16-bit sessions, the files will be truly lossless, and auditioning the input will accurately replicate the post-decoder audio.

This is not true for a 24-bit session. As the mp3-HD codec is only 16-bit capable, the 24-bit input can be dithered and truncated internally as part of the encoding process. Thus the mp3-HD encoded-decoded file will not be bit-identical with either a 24-bit rendered wav file, or an externally dithered and truncated 16-bit wav file. (Note an mp3-HD file can be decoded to either 16 or 24-bit audio files.)

While encoding/auditioning an mp3-HD signal, attention should be paid to the lossy core. We recommend selecting the appropriate core codec in an adjacent codec window, which will allow correct auditioning of the lossy codec in the mp3-HD channel, and will also allow correct overload monitoring and mitigation.

The mp3-HD lossy core takes the parameters:

Bitrate	Selectable from 96kbps to 320kbps
Mode	Selectable from CBR to VBR Very High
Quality	Low/Fast

However, remember that if the Bitstream Level Trim is used to mitigate the core codec overload, and if that same trim value is applied to the mp3-HD encoder input, then *both* the lossless and the lossy data streams will be trimmed.

Also note that, under normal circumstances, if the lossless codec Bitstream Level Trim is non-zero to avoid clipping in the lossy chain for an offline encode that uses the online trim values, then the resulting decoded output file will not be lossless because the level has changed.

5.2 Dither, Truncation and Lossless Codecs

When encoding online, or preparing a session for render through the DAW, it is essential to provide the input to the lossless codecs with a correctly dithered and truncated audio data stream. As stated above, the plug-in should be the last element of the processing chain, and a dither/truncation plug-in should immediately precede it, if the lossless codecs are to be used.

When a lossless codec is selected for audition, the FFT display area changes slightly to include an input bit-meter for the left-hand scale. This bit-meter is included specifically to help select the correct dither/truncation settings for the signal into the codec.

There is one segment per bit for bits 1 to 24 (see right). One large segment at the bottom indicates any activity in bits higher than 24.




Naturally, with a 16-bit input you would expect activity in bits 1-16, and with a 24-bit input you would expect activity in bits 1-24. As noted in the introduction to this section, if there is processing before the plug-in, you would expect ‘inters’ and activity on all bits, including bits higher than 24 (even if it’s a 16-bit input signal).

The following diagrams show the plug-in FFT displaying the **INPUT** and **DIFF** signals with varying degrees of dither and truncation.

First, in order to see the effect on the noise floor we input 1kHz and 2kHz tones at -60dB. As this is a low-level signal at -60dB, there is no activity on the top ten segments of the bit-meter.

Note the magnified scale down to -180dB. (Swipe left with the right-hand mouse button pressed.)

5.2.1 HD-AAC




HD-AAC 24-bit dither/truncation (correct dither/truncation)

The noise floor of the input signal is displayed as around -175dB (~30dB down from -144dB).

The diff signal is indicating -180dB (the 24-bit integer input is identical to the 24-bit integer post-codec, with no ‘inters’).

The bit meter shows no activity in bits greater than 24.




HD-AAC 20-bit dither/truncation

The noise floor of the input signal is displayed as around -150dB (~30dB down from -120dB).

The diff signal is indicating -180dB (the 24-bit input has the bottom four bits zero'ed).

The bit meter shows no activity in bits greater than 20.




HD-AAC 16-bit dither/truncation

The noise floor of the input signal is displayed as around -125dB (~30dB down from -96dB).

The diff signal is indicating -180dB (the 24-bit input has the bottom eight bits zero'ed).

The bit meter shows no activity in bits greater than 16.


5.2.2 mp3-HD



mp3-HD 16-bit dither/truncation (correct dither/truncation)

The noise floor of the input signal is displayed as around -125dB (~30dB down from -96dB).

The diff signal is indicating -180dB (the 16-bit integer input is identical to the 16 bit integer post-codec, with no 'inters'). The bit meter shows no activity in bits greater than 16.




mp3-HD 20-bit dither/truncation

The noise floor of the input signal is displayed as around -150dB (~30dB down from -120dB).

The diff signal is indicating -125dB with overtones at around -100dB (this is an error condition because the 20-bit input contains information that the 16-bit output of the codec cannot reproduce).

The bit meter shows incorrect dither activity in bits 17 to 20.



mp3-HD 24-bit dither/truncation

The noise floor of the input signal is displayed as around -175dB (~30dB down from -144dB).


The diff signal is indicating -135dB with overtones at around -100dB (this is an error condition because the 24-bit input contains information that the 16-bit output of the codec cannot reproduce).

The bit meter shows incorrect dither activity in bits 17 to 24.

5.3 Normal Signal Treatment with Lossless Codecs


Following on from above, we can now investigate the behaviour of the lossless codecs under normal signal conditions. A typical audio signal is shown with the display zoomed in. As the signal level is high, we would expect activity at the top of the bit-meter.

5.3.1 HD-AAC




HD-AAC input from 24-bit wav file

The diff signal and the bit-meter are indicating no presence of audio data input at bits greater than 24.



HD-AAC input from 24-bit wav file, with channel processing


The diff signal and the bit-meter indicate that processing in the channel has resulted in 'inters' up to a level of -144dB. The input to the codec is not correctly dithered and truncated.



HD-AAC input from 24-bit wav file, with channel processing, dithered and truncated to 24 bits

A dither/truncation plug-in has been added to the signal chain just before the Pro-Codec. The diff signal is indicating correctly dithered and truncated audio data input to 24 bits.


5.3.2 mp3-HD:



mp3-HD input from 16-bit wav file

The bit-meter indicates no activity higher than 16 bits.


The diff signal is indicating no presence of audio data input at bits greater than 16.



mp3-HD input from 24-bit wav file

The bit-meter indicates no activity higher than 24 bits. The red shading indicates that dither is required.


The diff signal is indicating the presence of audio data input at bits 17-24.



mp3-HD input with channel processing

The bit-meter indicates activity in all bits. The red shading indicates dither is required.


The diff signal is indicating the presence of audio data input at bits greater than 16.



mp3-HD input with channel processing, dithered and truncated to 16 bits

The bit-meter indicates no activity higher than 16 bits. This input is correctly dithered and truncated.

The diff signal is indicating correctly dithered and truncated audio data input to 16 bits.



mp3-HD input with channel processing, dithered and truncated to 24 bits

The bit-meter indicates activity higher than 16 bits. The red shading of bits 17-24 indicates incorrect dither.


The diff signal is indicating the presence of audio data input at bits 17-24.

5.4 Internal Dither and Truncation in Lossless Codecs

The previous examples include the insertion of a dither/truncation plug-in in the audio path to provide correct treatment of the signal prior to the codec.


For convenience, if a dither plug-in is not available, the **Dither at XX-bit Level** option can be ticked in the **CODEC & SETTINGS** pop-up for either of the lossless codecs. However, remember that if the plug-in is subsequently bypassed to render to wav, that resulting output will not be dithered and truncated correctly.

The following principles apply to both mp3-HD and HD-AAC; the diagrams refer to HD-AAC only.




HD-AAC non-truncated input

The diff signal and the bit-meter are indicating presence of audio data input at bits greater than 24. The bits greater than 24 are indicated in red to show an error condition.



HD-AAC input using the internal dither/truncation

The bit-meter indicates that there is no activity in bits greater than 24. The input to the codec is correctly dithered and truncated. The border for the bottom segment is green to indicate that the dither/truncation is provided internally.



HD-AAC input dithered and truncated twice


Here a correctly dithered/truncated signal enters the plug-in, but the internal dither is also applied. One of the operations is redundant, and this is indicated by a red border round the bottom segment.

Furthermore, there is a situation when it is desirable to use the internal dither/truncation, even if the signal presented to the plug-in input is correctly treated.

This case is when the lossy core has been observed to cause overloads on decode, and the remedy has been to use the internal trim control. Here, the trim is applied internally to the plug-in, and any external dithering is not helpful. The signal should be dithered using the internal dither option.



The GUI above shows a correctly dithered signal input to the plug-in, but the internal trim is non-zero. The bottom segment of the bit-meter is flashing red to indicate an error condition.




If the internal dither is enabled, the bit-meter indicates that the signal is correctly treated (the bottom segment of the meter is hollow — see above) and the active dither is applied internally (indicated by a green border around the segment).

Again, it is preferable to avoid using the internal trim by lowering the level into the plug-in.

5.5 Overloading Lossless Codecs

The screenshot below illustrates the behaviour of lossless codecs under overloading conditions. A typical audio signal is shown with the display zoomed. The input level is such that the input meters show clipping.

The behaviour of both codecs is the same; as neither has internal headroom, the signal will clip.



Note the **DIFF** signal at well over -80 dB; this overload condition is clearly audible in the **DIFF** signal.

5.6 Reasons to Audition Lossless Codecs

Finally, a few observations on the reasons why you should audition a lossless codec, because a hasty person might initially think that it can serve no purpose. If you have carefully read this manual this far, you will have realised there are several good reasons to audition lossless codecs.

5.6.1 Overloading Input (Sections 4.3 and 5.5)

The input meters will indicate an overloading input, but the **BITSTREAM LEVELS** meters will not show clipping for lossless codecs because they have no internal headroom. During audition, the **DIFF** signal clearly shows the effect of overloading the input to the lossless codecs, and the effect will be very audible.

5.6.2 Lossy Core (Section 5.1)

If it is intended to make use of the lossy core, it is wise to audition the lossy channel that is buried in the lossless codec.

5.6.3 Level change: (Section 3.4)

If a level change is required in the plug-in, and implemented using the **TRIM** tab, again it is wise to verify the audition signal.

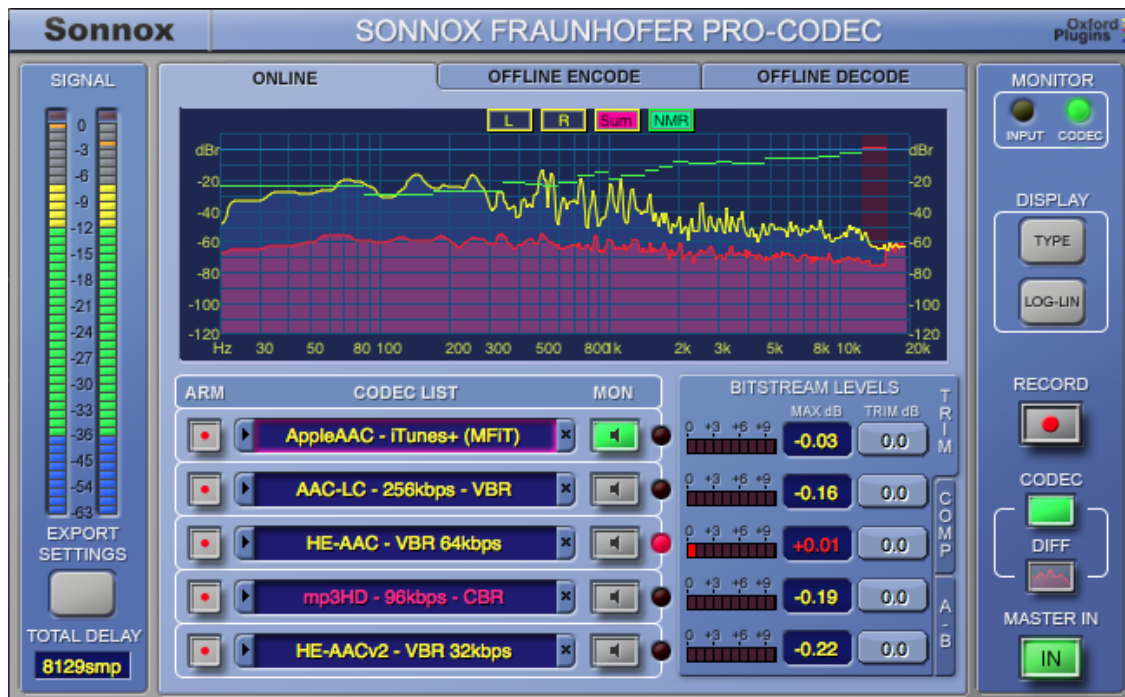
5.6.4 Dither/Truncation (Section 5.2 and 5.3)

It is necessary to correctly dither and truncate the input to the lossless codecs. The audition **DIFF** signal quickly indicates if correct dither and truncation is being applied. Simply put, if any noise signal is seen in the **DIFF** signal, you have not dithered/truncated the input to the plug-in properly. Note that you might not need to dither/truncate if the input files are playing without any processing; again the **DIFF** signal indicates quickly whether this is true or not.

5.6.5 Hearing a difference

Some people might think they can hear a difference between an input and the output of a lossless codec. Using **A-B mode** or the **A-B-X test**, this plug-in allows a user to make up their own mind if they can hear the difference.

6. Description of Controls



Touch Pad Controls

Several controls in this plug-in use a 'touch pad' associated with a knob or slider. These controls give a clear display of the SI value of a control, and respond to the following actions:

- Left-Click and drag up/down: Increase/decrease the value
- Right-Click and drag up/down: Fine adjustment of value
- Double Click: Directly edit value
- Shift + Left-Click and drag up/down: Fine adjustment of value
- Control + Left-Click: Set to default

Sonnox Menu Button

Provides access to a number of useful options:

Clip Lights — Can be set to hold for 2 seconds, 5 seconds, or indefinitely.

Show Preset Name Path — Shows or hides the Preset name path in the Preset Manager toolbar.

Use Host Channel Configuration — Most hosts indicate channel configuration. For those that don't (such as Sound Forge), the configuration can be forced to mono or stereo.

Clip Trigger — Can be set to one of three states:

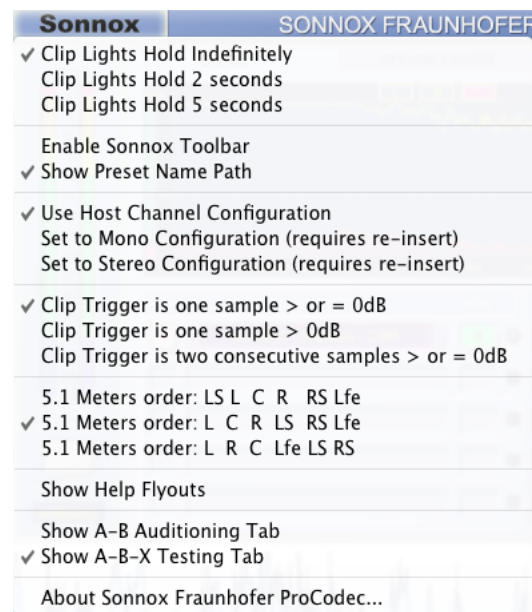
- one sample > or = 0dB
- one sample > 0dB
- two consecutive samples > or = 0dB

5.1 Meters order — One of three popular sequences of metering can be selected.

Show Help Flyouts — This useful tool can be enabled or disabled.

Show A-B Auditioning Tab — Can be set to A-B Auditioning or A-B-X Testing.


About... — Displays the date, version and build number of the plug-in.



Input Section

SIGNAL — The input meter indicates 1dB per segment for the top 36dB of dynamic range, and 3dB per segment thereafter. It includes peak-hold indication. The meter will be mono, stereo or multi-channel depending on context.

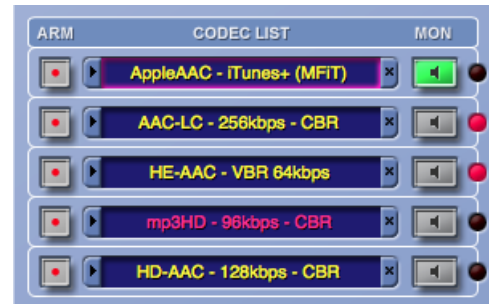
EXPORT SETTINGS — This button calls up a pop-up used to configure the encoded file data.



TOTAL DELAY — This window shows the total delay through the plug-in; click to toggle between samples or milliseconds. The figure given is the maximum of each individual codec delay. This field is updated in the GUI immediately a new codec or other setting is selected. Depending on the host behaviour, the new value may not be available for delay compensation until the plug-in is removed and re-instantiated.

Online Encode Tab

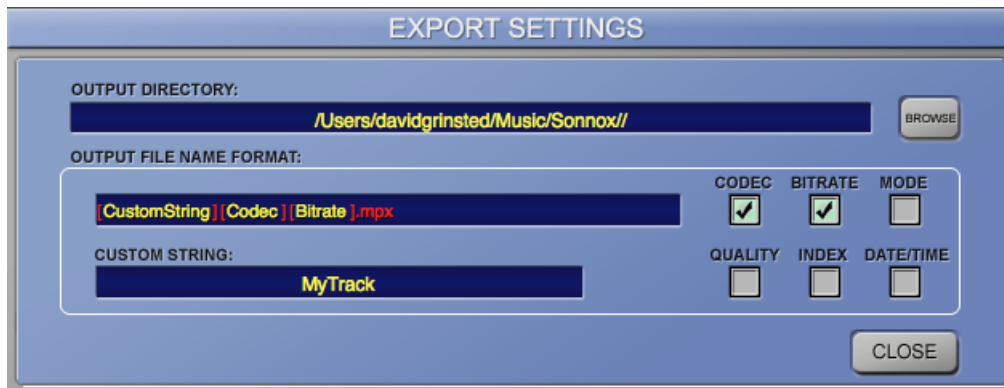
CODEC LIST — Displays a list of codecs chosen for auditioning. The codec type, bit-rate and codec mode are displayed at all times. Any codec pre-loaded in the list can be auditioned without glitches. You can click and drag to re-position codecs in the list (Windows only), or remove codecs by clicking on the little blue 'x' tab at the right hand side of each codec window.



MON (x5) — These select the active codec for routing to the output section. All monitor paths are delay compensated for glitch-free switching.

ARM (x5) — Prepares an activated codec in the codec list for writing to disk.

Export Settings Pop-up



OUTPUT DIRECTORY — Selects the destination directory for output files.

BROWSE — Opens a file browse directory, prompting you to select an output directory.

OUTPUT FILE NAME FORMAT — Displays the complete string that will be written to an output directory. You can re-order any of the fields in red brackets for a completely custom filename. The filename fields are:

CUSTOM STRING Displays the identifying name string of the filename, eg. a song name.

CODEC Add the codec type. Optional.

BITRATE Add the encoder bit-rate. Optional.

MODE Add the encoder mode; either CBR or VBR. Optional.

DATE/TIME Add the current time and date when the file is created. Optional.

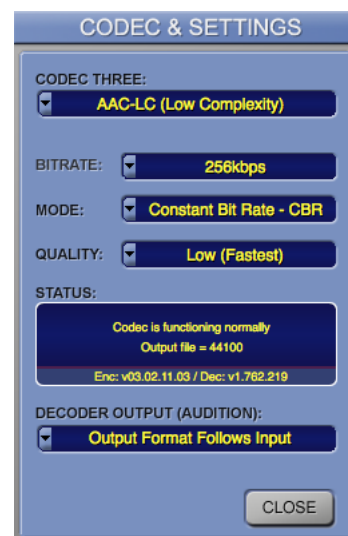
INDEX Add a codec specific number index. Optional.

QUALITY Add the encoder quality information. Optional.

Codec & Settings Pop-up

CODEC — Use the drop-down menu to choose from a list of available codecs. Some codecs will not be available, depending on the host sample rate and the number of channels. For example, HE-AACv2 is stereo only.

BITRATE — Choose the bit-rate of the selected codec. Some bit-rates may not be available, depending on the sample rate of the host DAW and the number of channels. If the bitrate is accompanied by an asterisk, the codec is resampling the audio stream.



MODE — Select from the available list either CBR or VBR encoder algorithms for the selected codec.

QUALITY — Select the quality by which the mode is calculated. The highest quality setting often causes codecs to re-sample the signal (in which case the bitrate and the quality setting will be displayed in red). If the codec is resampling, the auditioned signal will *not* reflect this, but the encoded signal written to disk will be at the highest quality setting, which is higher than the auditioned signal. This is the one case where the auditioned signal will not reflect the written signal.

STATUS — Shows the status and versions of the selected codec.

DECODER OUTPUT (AUDITION) — Output format follows input.

Dither at 24-Bit Level — Only visible if the selected codec is HD-AAC. Select if input signal is not correctly dithered and truncated to 24 bits.

Dither at 16-Bit Level — Only visible if the selected codec is mp3-HD or HD-AAC 16-bit mode. Select if input signal is not correctly dithered and truncated to 16 bits.

Output Section

MONITOR — Indicates the plug-in's audio output; either **INPUT** (ie. the plug-in is bypassed) or **CODEC**. You can switch between input and codec output by clicking on the LEDs. A special case for the mp3-HD codec, whereby the **MONITOR** signal is synthesized, is denoted by an **INPUT 16** indication.

DISPLAY TYPE — Toggles between a Lissajous phase display and the FFT display.

DISPLAY LOG/LIN — Toggles between logarithmic and linear scale FFT displays.

RECORD — Writes the active codecs pre-selected for record-arm to disk.

CODEC — When selected, will route the encoded signal to the output.

DIFF — When selected will route the difference signal to the output. The difference signal is the output encoded signal subtracted from the unprocessed input signal.

MASTER IN — With the **IN** button selected, all processing is active. If not selected, the outputs are disabled. However, much of the processing remains active, for example, the FFT display and the overload detection. The **MASTER IN** control is automatable.

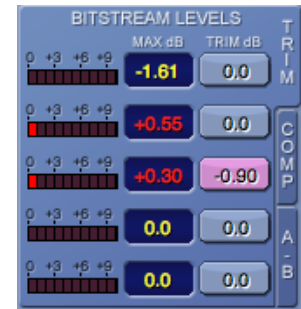


The five **MON** buttons and the **MASTER IN**, **CODEC** and **DIFF** buttons all provide smoothly interpolated and glitch-free audio switching. It is advisable to use these for auditioning, rather than using the host application's bypass function (which is commonly not interpolated).

Bitstream Levels / Trim Tab

METERS (x5) — Indicate when a decoded signal is above 0dB. Includes peak hold.

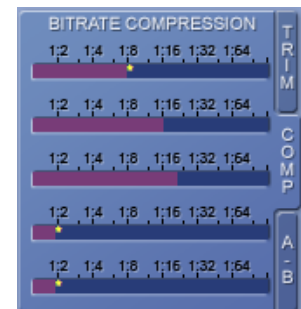
MAX dB (x5) — Indicates the maximum level detected by the decoder. This value is yellow if negative, and red if positive. Clicking in this window resets the peak hold value.



TRIM dB (x5) — Indicates the input trim to the individual codecs. This is a Touch Pad control, allowing a fader style movement to determine an attenuation of up to a maximum of 12dB. If this value is non-zero, the Touch Pad background is pink. The **TRIM dB** controls are automatable.

Comp(ression) Tab

BITRATE COMPRESSION (x5) – Shows an indication of the amount of data compression in the output data stream or file. The compression is expressed as a ratio, from 1:1 down to around 128:1. A yellow asterisk denotes an estimated value.



A-B Auditioning Tab*


A (x5) — Selects the codec to be routed through 'A'.

B (x5) — Selects the codec or input to be routed through 'B'.

A-B — Glitchlessly switches output routing between active codecs selected in 'A' or 'B'.

HIDE — Cloaks the AB matrix state and FFT scope.

BOTH — Displays both 'A' and 'B' difference signals.



A-B-X Test Tab*

A (x5) — Selects the codec to be routed through ‘A’.

B (x5) — Selects the codec or input to be routed through ‘B’.

A-B-X — Glitchlessly switches output routing between active codecs selected in ‘A’ or ‘B’, or the anonymous codec in ‘X’.

RESET — Sets the ABX machine back to a start condition. Resets counters to zero and sets X to either A or B.

CHOOSE — Increments cycle counter. Examines whether choice of A or B correctly matches X. If correct guess, increments the success counter. Randomly re-allocates A or B to X.

REVEAL — Calculates the ‘success’ counter divided by the cycle counter and displays the result.

DISPLAY — Shows number of cycles completed during test; when **REVEAL** is pressed, it will show the success ratio.

***NOTE:** Select the **A-B Auditioning Tab** or the **A-B-X Test Tab** via the **Sonnox** menu button



Offline Encode Tab

BROWSE — Opens a file browse directory, prompting you to select a wav or AIFF file for offline encoding.

ENCODE — Encode all queued files in the list.

TRIM BEHAVIOUR — Four possible settings are available:

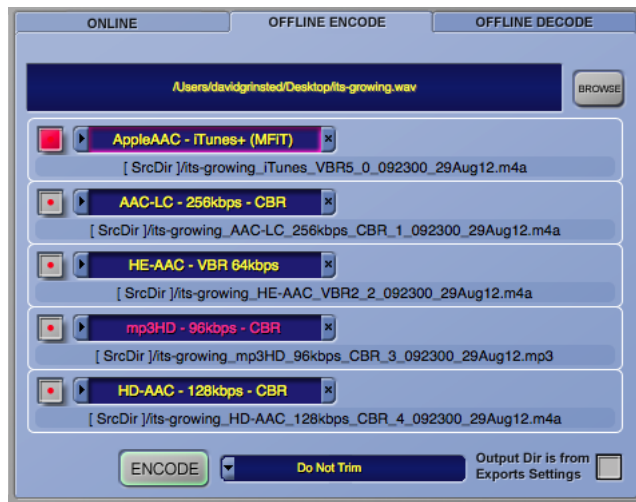
Do not trim: No gain adjustments are made.

Use online trim values: The trim values from the online page are used.

Prevent Decoded Levels > 0dB: The source file is encoded, then decoded and the maximum sample value is found. If it is greater than 0dB, the source file is re-encoded with the minimum gain adjustment necessary to bring the decoded max level down to 0dB.

Normalise Decoded Levels: The source file is encoded, then decoded and the maximum sample value is found. The source file is then re-encoded with a gain adjustment to force the decoded max value up or down to 0dB.

Output Dir is from Export Settings — The output file(s) will be written to the directory specified in the **EXPORT SETTINGS** window, instead of the same directory as the source file.



Offline Decode Tab

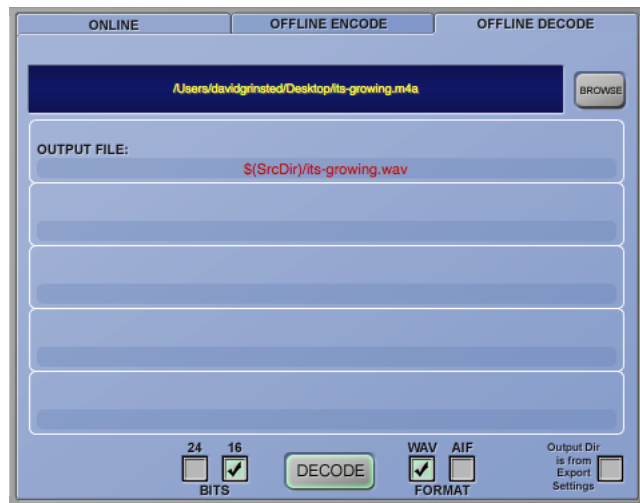
BROWSE — Opens a file browse directory, prompting you to select a compatible compressed file format for offline decoding.

DECODE — Decodes all files in the list.

BITS — The output file format can be set to 16 or 24-bit depth. Files created with a bit depth of 24 bits will be approximately 50% larger than files created with a depth of 16 bits.

FORMAT — The output file format can be set to wav or AIFF.

Output Dir is from Export Settings — The output file(s) will be written to the directory specified in the **EXPORT SETTINGS** window, instead of the same directory as the source file.



7. Presets and Project/Session Data Handling

The Sonnox Fraunhofer Pro-Codec plug-in contains internal data not normally found in a plug-in, such as filenames and directory paths. This has prompted the adoption of a slightly different preset model in order to allow the creation and utilization of independent libraries of data presets.

7.1 Internal Data Rules

The following special data have special rules applied to them:

- a) The **Export Settings data** (the **Custom String**, filename **Attributes** and **Output Directory**)
- b) The five **Trim dB** values
- c) The five **Rec Arms**
- d) The **Offline Encode Headroom** selection

Unless detailed otherwise, all other data are always returned in any form of preset load.

7.1.1 Rule 1: Project / Session loads (including host preset loads)

Will load all data as it was saved, including all the **Special data** above.

7.1.2 Rule 2: Onboard Preset Manager (PM) preset loads (including PM copy/pastes, + and -)

Will *not* load any of the special data listed above.

To be more specific:

The **Output Directory** and filename **Attributes** remain as they have been set

The **Custom String** is reset to the default 'MyTrack'

The five **Trim dB** values are reset to 0dB

The five **Rec Arms** are reset to off

The **Offline Encode Headroom** selection remains as previously set

Also, note that PM preset loads ***must*** match the session sample rate and the session channel configuration to avoid incompatible codec settings (the codecs are sensitive to session sample rate and channel configuration).

See section 8 for more on the Preset Manager Toolbar

7.1.3 Rule 3: Persistence for plug-in insertion and deletion:

Only the **Output Directory**, filename **Attributes** and **Offline Encode Headroom** selection are remembered across a delete/insert. All other controls (whether special or otherwise) are not remembered, and so are defaulted on plug-in insert (notably the **Custom String** is defaulted to ‘MyTrack’).

7.2 Other Data

There is limited provision for ID3 metadata. The plug-in inserts the following metadata into an mp3 stream:

- Date and time
- Codec type, bitrate and other settings
- Encoded by manufacturer message (‘Sonnox Fraunhofer Pro-Codec Plug-in’)

These metadata are inserted into an output stream only; they are not stored internally in the plug-in, and are not applicable to Presets or session data.

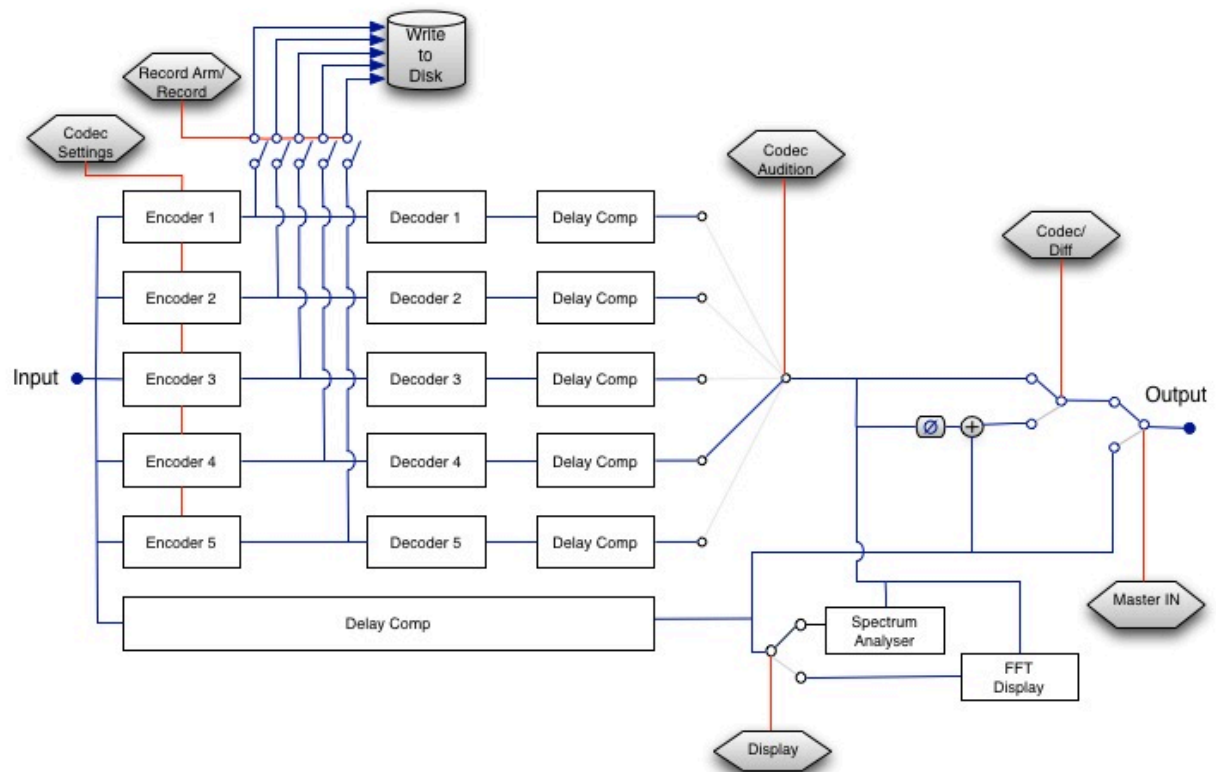
Currently, AAC metadata (eg. for iTunes) is not supported.

8. Preset Manager Toolbar

The Sonnox Fraunhofer Pro-Codec plug-in comes equipped with its own onboard Preset Manager, which is displayed at the top of the plug-in window as if the host created it. The reasoning behind this is to allow increased portability of your Presets across all the host applications, while also providing a consistent and versatile interface. While most host platforms allow creation and loading of Presets, those host-created Preset files are not portable between different host applications. With the Oxford plug-ins' Preset Manager, you can create a named Preset in one host application and load it when using an alternative application.

The Sonnox Preset Manager is fully described in a companion document — 'Sonnox Toolbar and Preset Manager Operation Manual' — available for download on the Support Documentation page of the **Sonnox** website: www.sonnox.com

9. Pro-Codec Signal and Control Flow Diagram



10. Specifications

10.1 Codec Specifications

Available bitrates are dependent upon the selected sample rate.

Encoder	mp3 / mp3surround
Channel Configurations	mp3: Mono, Stereo mp3surround: 5.1
Sample Rates (kHz)	mp3: 8, 11.025, 12, 16, 22.05, 24, 32, 44.1 & 48 mp3surround: 44.1 & 48
Bitrates	mp3 Mono: 8-320 kbps in variable and constant rate mp3 Stereo: 16-320 kbps in variable and constant rate mp3surround 5.1: 128-320 kbps in variable and constant rate

Encoder	mp3-HD
Channel Configurations	Mono, Stereo
Sample Rates (kHz)	Mono & Stereo: 32, 44.1 & 48
Bitrates	Variable (lossless)
Limitations:	Only 16-bit integer input values supported

Encoder	AAC (AAC-LC, HE-AAC)
Channel Configurations	Mono, Stereo, 5.0, 5.1
Sample Rates (kHz)	AAC-LC: 11.025-96 kHz HE-AAC: 22.05-96 kHz HE-AAC v2: 22.05-48 kHz
Bitrates (Stereo, 44.1kHz)	AAC-LC: 96-256 kbps in variable and constant rate HE-AAC: 16-128 kbps in variable and constant rate HE-AAC v2: 18-56 kbps in variable and constant rate

Encoder	AAC (MPS)
Channel Configurations	5.1
Sample Rates (kHz)	AAC-LC: 44.1 & 48 kHz HE-AAC: 44.1 & 48 kHz
Bitrates (Stereo, 44.1kHz)	AAC-LC: 192-320 kbps HE-AAC: 48-128 kbps

Encoder	HD-AAC
Channel Configurations	Mono, Stereo, 5.1 (7.1 not currently supported)
Sample Rates (kHz)	44.1, 48, 88.2, 96, 176.4, 192
Bitrates	Variable (lossless)
Limitations:	Supported oversampling modes: 1 for input sample rate 44.1 or 48kHz 2 for input sample rate 88.2 or 96kHz 4 for input sample rate 176.4 or 192kHz

10.2 Sample Rates

The supported session sample rates of the plug-in are from 32kHz min to 192kHz max (where the host DAW allows). The supported codec sample rates are from 8kHz min to a maximum of 48kHz (MP3), 96kHz (AAC family) or 192kHz (HD-AAC only, see above).

Under certain settings at higher session sample rates, downsampling occurs between the session and the audition path, or between the session and the encoder. This downsampling is lossless to 24 bits.

The following tables indicate the preferred sample rates of the codecs (see section 4.7). Only supported sample rates from 32kHz and higher are shown.

mp3

Bitrate Range (kbps)	Supported Sample Rates (kHz)	Preferred Sample Rate (kHz)	Number of Channels
40-47	32	22.050	1
48-55	32	24	1
56-63	32 / 44.1 / 48	32	1
64-320	32 / 44.1 / 48	44.1	1
96-111	32 / 44.1	32	2
112-319	32 / 44.1 / 48	44.1	2
320	32 / 44.1 / 48	44.1 / 48	2
128-191	44.1 / 48	44.1	5.1
192-320	44.1 / 48	44.1 / 48	5.1

AAC-LC

Bitrate Range (kbps)	Supported Sample Rates (kHz)	Preferred Sample Rate (kHz)	Number of Channels
32-55	32	32	1
56-127	32 / 44.1 / 48	44.1	1
128-191	32 / 48 / 96	48	1
192-264	44.1 / 48 / 96	96	1
264-288	48 / 88.2 / 96	96	1
40-95	32	32	2
96-111	32 / 44.1 / 48	32	2
112-255	32 / 44.1 / 48	44.1	2
256-384	32 / 48 / 96	48	2
384 – 529	44.1 / 48 / 96	96	2
529-576	48 / 88.2 / 96	96	2
160-239	32	32	5 , 5.1
240-279	32 / 44.1 / 48	32	5 , 5.1
280-639	32 / 44.1 / 48	44.1	5 , 5.1
640-800	32 / 48 / 96	48	5 , 5.1

HE-AAC

Bitrate Range (kbps)	Supported Sample Rates (kHz)	Preferred Sample Rate (kHz)	Number of Channels
12-17	32	32	1
18-39	32 / 44.1 / 48	44.1	1
40-56	32 / 44.1 / 48	48	1
16-27	32 / 44.1 / 48	32	2
28-95	32 / 44.1 / 48	44.1	2
96-128	32 / 48 / 96	48	2
64-69	32 / 44.1 / 48	32	5 , 5.1
70-159	32 / 44.1 / 48	44.1	5 , 5.1
160-319	32 / 44.1 / 48	48	5 , 5.1
320-640	64 / 88.2 / 96	96	5 , 5.1

HE-AAC v2

Bitrate Range (kbps)	Supported Sample Rates (kHz)	Preferred Sample Rate (kHz)	Number of Channels
12-17	32	32	2
18-39	32 / 44.1 / 48	44.1	2
40-56	32 / 44.1 / 48	48	2

MPS (AAC-LC)

Bitrate Range (kbps)	Supported Sample Rates (kHz)	Preferred Sample Rate (kHz)	Number of Channels
192	44.1 / 48	44.1	5.1
256	44.1 / 48	44.1	5.1
320	44.1 / 48	44.1	5.1

MPS (HE-AAC)

Bitrate Range (kbps)	Supported Sample Rates (kHz)	Preferred Sample Rate (kHz)	Number of Channels
48	44.1 / 48	44.1	5.1
64	44.1 / 48	44.1	5.1
96	44.1 / 48	44.1	5.1
128	44.1 / 48	44.1	5.1

10.3 Delays and CPU Usage

10.3.1 Sample Delays

The sample delay (latency) of the plug-in depends upon the particular codecs selected. The monitor outputs are delay compensated, so if more than one codec is selected in the **CODEC LIST**, then the reported delay is the maximum of the individual codec delays. This field is updated in the GUI immediately a new codec or other setting is selected. Depending on the host behaviour, the new value might not be available for delay compensation until the plug-in is removed and re-instantiated.

Individual codec sample delays are given in the following tables:

		44.1kHz	48kHz
mp3	CBR max/min	13489 / 7345	13489 / 7345
	VBR	37041	39127
mp3-HD		3072	3072
AAC-LC	CBR/VBR	3648	3648
HE-AAC	CBR max/min	7398 / 7300	7398 / 7300
	VBR	7300	7300
HE-AAC v2	CBR/VBR	9153	9153
HD-AAC		4096	4096
Apple AAC	VBR	5120	8192

		88.2kHz	96kHz
mp3	CBR max/min	26978 / 14690	26978 / 14690
	VBR	74082	78254
mp3-HD		6144	6144
AAC-LC	CBR/VBR	7296 / 3648	7296 / 3648
HE-AAC	CBR max/min	14796 / 6302	14832 / 6312
	VBR	14600	14604
HE-AAC v2	CBR/VBR	18306	18306
HD-AAC		6144	6144
Apple AAC	VBR	12288	12288

10.3.2 CPU Usage

All codecs are mathematically intense, and some have significant latency as a result.

Separately from the latency, some of the codecs are also CPU-intensive. The user can take advantage of advances in CPU power in modern computers if appropriate, but can also use the plug-in extensively with older and possibly lower-powered computers by selecting fewer codecs simultaneously.

In particular, the HD-AAC lossless codec creates a heavy CPU load. High sample rates exacerbate this situation; for example, without a powerful CPU, the HD-AAC codec might cause a prohibitively high CPU load at 192kHz.

All codecs that are added to the **CODEC LIST** are always active, in order to report accurately the NMR signal and potential overload state. This is irrespective of which codec is selected for monitoring. Remember that this requires both encoding and decoding of the signal in real time. Thus it is possible to have dropouts with an underpowered CPU if several highly intensive codecs are active simultaneously.

Online recording will add to the CPU activity (and file-handling activity).

Offline encoding and decoding are CPU-bound, not real time bound.

Graphical Screen: The graphical area is not CPU-intensive, and should not cause performance problems even on older CPUs.

10.3.3 Internal Buffer Sizes

Additionally, for most codecs, the plug-in uses internal buffer sizes of 1024 for up to 48kHz operation; 2048 for greater than 48kHz up to 96kHz operation; and 4096 for sample rates higher than 96kHz. We recommend that the buffer sizes of the sound card, and thus of the host application, are set to match these figures.

Buffer sizes for the HE-AAC and HE-AACv2 codecs are twice these figures (so 2048 for up to 48kHz operation; 4096 for greater than 48kHz up to 96kHz operation; and 8192 for sample rates higher than 96kHz).

If the buffer sizes are not matched, there will be an inefficiency in the data transfer, and in some applications this leads to uneven CPU behaviour.

10.4 GUI Controls

Input Meter: 0dB to -63dB

Graphical Display Frequency: 20Hz to 20kHz (44.1 and 48kHz)

Graphical Display Gain: +10dB to -120dB (to -180dB with zoom)

Bitstream Level Trim: -12dB to 0dB

Trim Meter: 0dB to +10dB

Trim Display: -144dB to +10dB

10.5 Summary of Codec Features and Applications

Codec	Features	Typical Applications	Typical Bitrate
mp3	General purpose low bitrate	<ul style="list-style-type: none">• Apple iPod• iTunes	128 kbps (stereo)
mp3-HD (mp3 High Definition)	Lossless audio codec with 16-bit depth and up to 48kHz	<ul style="list-style-type: none">• Home networks• Music distribution and production• Archiving	Roughly half the bitrate of a 16-bit uncompressed wav file.
AAC-LC (AAC Low Complexity)	High performance audio codec for excellent audio quality at low bitrates.	<ul style="list-style-type: none">• Apple iPod• iTunes• ISDB television broadcast (Japan)	128 kbps (stereo)
HE-AAC (High Efficiency AAC)	High performance audio codec for good quality at bit-rates of 32-48 kbps per channel.	<ul style="list-style-type: none">• XM Radio• Mobile music download• Digital Radio Mondiale	64 kbps (Stereo)
HE-AAC v2	High performance audio codec for good quality at bit-rates of 16-24 kbps per channel.	<ul style="list-style-type: none">• 3GPP music download• Digital Radio DAB+• Internet radio streaming to mobile devices eg. iPhone	48 kbps (Stereo)
HD-AAC (High Definition AAC)	Lossless audio codec with 24-bit depth and up to 192kHz	<ul style="list-style-type: none">• Home networks• Music distribution and production• Archiving	Roughly half the bit-rate of a 24-bit uncompressed wav file.
Apple-AAC	High performance audio codec for excellent audio quality at low bitrates.	<ul style="list-style-type: none">• Apple iPod• iTunes	256 kbps VBR (stereo)

11. Copyright and Acknowledgements

The Sonnox Fraunhofer Pro-Codec plug-in was developed by the design team at Sonnox, in collaboration with Fraunhofer IIS.

The Audio Coding software contained in the Sonnox Fraunhofer Pro-Codec plug-in was developed and provided by Fraunhofer IIS of Erlangen, Germany.

Trademarks and content copyright © 2007-Present Sonnox Ltd. All rights reserved.

Fraunhofer IIS and related logos are registered trademarks.

Apple AAC, iTunes, iTunes Plus and Mastered for iTunes are trademarks or registered trademarks of Apple Inc.

DIGIDESIGN, AVID, and PRO TOOLS are trademarks or registered trademarks of Digidesign and / or Avid Technology, Inc.

VST is a Trademark of Steinberg AG.

All other product and company names are trademarks or registered trademarks of their respective holders.

These products are protected by several patents either active or pending in the UK, the EU and the USA.

This product is manufactured and supplied by Sonnox Ltd.

11.1 Commercial Use of mp3 Content

It is important to distinguish between the use of codecs and the use of content produced by the codecs. It is therefore important to note that any mp3 royalty fees are independent of which encoder is used to produce the mp3 file content. If royalty fees are due for a particular project or enterprise, they will be due irrespective of whether the Pro-Codec is used or some alternative encoder is used. There is no circumstance where using the Sonnox Fraunhofer Pro-Codec will increase licence fees or require additional fees of any sort.

Both the mp3 and AAC compression routines are based on technology that is patented. At present no royalties are required for AAC codec usage, and private or non-commercial use of mp3 content does not require royalty fees.

Under some circumstances the commercial (ie. revenue-generating) use of mp3 content will require a separate licence irrespective of how that content is generated.

To summarise:

There are no licence fees for private or non-commercial use of mp3 content.

Enterprises with a turnover of less than \$100,000 per year are exempt from mp3 royalty payments.

There are currently no royalties collected for AAC commercial content distribution.

If you are in doubt whether fees are applicable for your project, please contact

<http://www.mp3licensing.com>

12. Manual Revision History

Version 1.0 date 11th May 2011 — Initial release version.

Version 1.1 date 5th September 2012 — Updated to reflect new software version 2.

Platform Specific Supplement

S1. Supported Platforms

Avid Pro Tools (LE, RTAS, M-Powered, Pro Tools HD) and Pro Tools 10 (HDX & Native)

VST Native

Audio Units Native

S2. System Requirements

These requirements are current at this revision of manual. For latest system requirements, please see the website www.sonnox.com

Pro Tools

Pro Tools 7, 8, 9 & 10

Approved Pro Tools CPU, OS and hardware configuration (see www.avid.com)

Mac OSX 10.4 or higher

Windows XP / Vista 32 & 64 bit / Windows 7 32 & 64 bit

Ram 1GB minimum

iLok USB key with latest drivers

Audio Units

Audio Units compliant host application (e.g. Logic, Digital Performer etc.)

Mac OSX 10.4 or higher

Ram 1GB minimum

iLok USB key with latest drivers

2nd generation iLok required for 64-bit plug-ins

VST

VST compliant host application (e.g. Nuendo, Cubase, WaveLab etc.)

Mac OSX 10.4 or higher

Windows XP / Vista 32 & 64 bit / Windows 7 32 & 64 bit

Ram 1GB minimum

iLok USB key with latest drivers

2nd generation iLok required for 64-bit plug-ins