THIMEO
STXTREME AUDIO PROCESSOR
USER MANUAL

User Warnings and Caution

The installation and service instructions in this manual are for use by qualified personnel only. To avoid electric shock, do not perform any servicing other than that contained in the operating instructions unless you are qualified to do so. This instrument has an auto range line voltage input. Ensure the power voltage is within the specified range of 100-240VAC.

CAUTION: HAZARDOUS VOLTAGES The instrument power supply incorporates an internal fuse. Hazardous voltages may still be present on some of the primary parts even when the fuse has blown. If fuse replacement is required, replace fuse only with same type and value for continued protection against fire.

WARNING: The product's power cord is the primary disconnect device. The socket outlet should be located near the device and easily accessible.

The unit should not be located such that access to the power cord is impaired. If the unit is incorporated into an equipment rack, an easily accessible safety disconnect device should be included in the rack design. To reduce the risk of electrical shock, do not expose this product to rain or moisture. This unit is for indoor use only.

This equipment requires the free flow of air for adequate cooling. Do not block the ventilation openings on the rear and front of the unit. The Thimeo logo on the front of the unit also provides ventilation. Failure to allow proper ventilation could damage the unit or create a fire hazard. Do not place the units on a carpet, bedding, or other materials that could interfere with any panel ventilation openings. If the equipment is used in a manner not specified by the manufacturer, the protection provided by the equipment may be impaired.

USA CLASS A COMPUTING DEVICE INFORMATION TO USER

WARNING: This equipment generates, uses, and can radiate radio-frequency energy. If it is not installed and used as directed by this manual, it may cause interference to radio communication.

CE CONFORMANCE INFORMATION: This device complies with the requirements of the EEC council directives: ♦ 93/68/EEC (CE MARKING) ♦ 73/23/EEC (SAFETY – LOW VOLTAGE DIRECTIVE) ♦ 89/336/EEC (ELECTROMAGNETIC COMPATIBILITY) Conformity is declared to those standards: EN50081-1, EN50082-1

Getting started with processing

STXtreme is an extremely powerful and easy to use processor. Our approach is different than most traditional processors because we believe to get the best sounding audio you need to repair any damage before processing. We have created a combination of audio processing methodology combined with the unique set of repair tools Thimeo has created to 'fix' the audio before processing.

Before any processing takes place, first STXtreme repairs the audio. Whatever you put into it, it sounds good when it comes out. So, unlike with other processors, the STXtreme does two jobs, repairing damaged audio and processing it.

The best way to get started with the STXtreme is to select presets. The STXtreme has many settings that are deeply hidden in the user interface, which you probably never want to touch. This is one of the things that makes it as powerful as it is, you can easily replicate the design of many other processors or create your very own, just by configuring all these settings. This does mean that different presets can have very different internal structures and sonic signatures, and because of that it's very important to start with a preset that uses a structure that matches what you want. Starting with a random preset and then trying to change it can be a massive effort if the original preset is very far off from what you want. After selecting a preset that is somewhat close to your goal, you can make a few simple adjustments to get just the sound you want. STXtreme is full of features and parameters. Do not let that scare you. We make them accessible so you can use them, however in most cases you will not need to.

Total Flexibility:

We give you, our customer, unlimited flexibility to create exactly the sound you want. Unlike many other processors, we do not hide controls. Because we do not have a fixed structure, it means you can toggle sections on or off with the icon and move them around in the chain. In the sections that can be moved in chain, you can click on the icon which will give you a drop-down menu where they can be moved.

Where applicable, you can click on the icon to hear the output at this stage of the processing, or the ≠ icon to hear what this processing stage adds or removes.

The fastest way from here to there:

STXtreme has lots of menus but we have lots of short cuts that make navigating much faster.
The fastest way to access any section of processing is simply to click on the label of that section, of the top of the web page, it will then turn orange, if you click on that, it will take you to that section. For example, if you click on multiband one, the screen that will then appear will be the multiband one screen.

Let's get this party started:

The following icons are used:

Home: You will either see this home symbol or the Thimeo logo.
Presets: there are several types of presets. We recommend you select one of each as a starting point.
I/O
Status overview
Network: This is the first page you want to visit on set up to enable web interface access.
Settings: Overall settings for the box including firmware updates.
On/off
≠ Difference: Plays what the processing stage adds or removes.
Monitor: Listen to the output of a stage of processing, eliminating all subsequent stages.
Location: Move processing stage to another place in the chain.
User account
Overview of instances: Go here to switch between instances if you have multiple instances.
Show or hide metering
Advanced settings. You will probably not need to use them.

GETTING STARTED

Installation:

The controls for STXtreme are available via web interface. While some controls are available via the front panel, we recommend setting up from HTML5 web interface once the network configuration is complete.

This is the screen you will see when you boot up the unit.

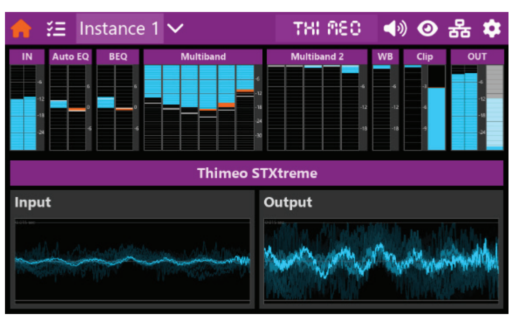

General front screen navigation: All icons are selectable. Once chosen the selected icon will turn orange and screen will change.

Return to home screen by pressing the icon in the upper left-hand corner.

Some buttons will open a pop up screen, which can be closed by pressing the ✕ icon in the top right corner.

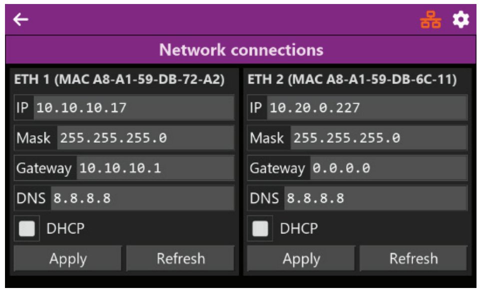

Network connectivity:

To set up network: on the home screen you will see the icon in the upper right hand.
To connect, you can click on the network symbol on the top bar. This takes you to a page that shows the IP addresses for the 2 network connectors. If a network cable is connected and the router dispenses the IP addresses, this page will show your IP address. You can just enter that IP address in a web browser on a device that is connected to the same network to load the web interface.

You can have DHCP checked to automatically obtain an address from the network or uncheck it and manually insert information, which would include IP, mask and gateway.

Web interface:

At this point the web interface should be available, so the next steps can be done via the front panel or the web interface.

The front panel of STXtreme only has a few settings that you can change. To fully use the STXtreme you will need to access the HTML5 web interface. This will enable configuration from remote PC, phone or tablet. The web interface is designed to be used on tablet, and all settings are touch screen friendly.

When you open the web interface, you will see the Quick adjust screen. We recommend you follow the steps below and then fine tune the settings on this page. Most people can get exactly the sound you want with this menu.

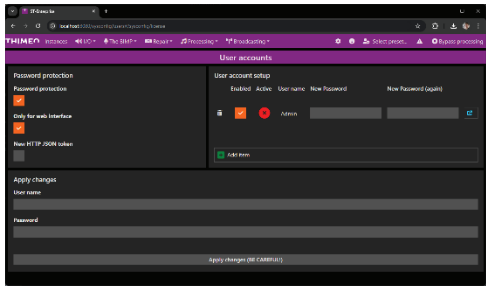

Password

We recommend you use password protection. Select password protection. Then add and or change password as needed. You can set up a password for web interface or entire system. You can also set up various levels of password control via 'user rights.
This allows you to create multiple user accounts and give different users view or change access to different things, such as I/O settings, processing settings, RDS settings, etc.

Basic preset selection: You can not make adjustments from front screen; you can only select a preset and put it on air.

Tap ✓☰ which you will find in upper left-hand corner next to 🏠 button.

There are three categories of presets: repair, processing and broadcast.

Once you select a preset category, only presets that contain presets for that category will be shown. Most people would use the front screen only select a preset to get on air. We highly recommend using the web interface for all adjustments and further preset selections.

I/O

🔈 Audio I/O:

There are a variety of analog and digital inputs and outputs on the unit:

Analog inputs and outputs, both balanced (XLR) and unbalanced (BNC), and digital inputs and outputs. For any left/right input and output, the analog balanced and digital inputs and outputs can be used. The STXtreme supports digital MPX (composite FM signal), so for MPX all outputs can be used. Often you will probably want to use the BNC outputs.

The audio interface can run at one of two frequency modes: 48/192 kHz or 44.1/176.4 kHz. This is mostly relevant if you use digital outputs. Digital outputs can be switched to use 44.1/48 kHz or 176.4/192 kHz in the software. Either way, the inputs and outputs are all perfectly synchronized, to allow for perfect HD/FM synchronization. Diversity delay is also supported, under the FM output settings.

Aside from these inputs and outputs, STXtreme also supports AES67 for inputs, and both streaming and MicroMPX for outputs, all via the network interfaces.

Low Latency:

If you require low latency for headphone outputs, use the low latency output.
For Diversity delay: If you need to synchronize FM and HD, it is under FM output diversity delay settings.   It is possible to set a diversity delay to sync the FM and HD signals. To use this, select by checking the box entitled "diversity delay."

Input 2:

This can be used for backup inputs or other purposes. You will need to tell it which mode to be in. The following modes can be used:

Back up switch to this if input is select.

Combined: which is not normally used

BIMP which is not normally used.

SCA ½ input: not normally used.

For MicroMPX: Go to I/O, then go to the MicroMPX page, there enter the IP address and port number of the decoder. For further information regarding MicroMPX, please see the MicroMPX manual.

Low Latency: If you require low latency for headphone outputs, use the low latency output.

For diversity delay: If you need to synchronize FM and HD, go to More FM output settings, under FM Output, you will find settings for synchronization with HD.

Here is the back panel of the unit to help get you started.

If you are connecting to a transmitter site remotely you may want to use MicroMPX.

The MicroMPX manual can be found at the end of this manual.

Low latency:

You will get the best audio quality in the normal latency modes. Normal latency depends on the settings but is typically around 100 ms. It is possible to reach latencies down to 5 ms, at a cost in audio quality. We do not recommend lowering the latency if you do not need to.

Low latency headphone monitoring: You have the ability to output a less than 5ms output for headphone monitoring., at a cost in audio quality, particularly in the bass section. This is useful for announcers who wish to monitor the on-air signal.

Repair and processing:

3 easy steps to processing nirvana:

The STXtreme helps you get started super fast by setting everything up via preset menus. Repair, processing and broadcast.

You will find the preset menus in the upper left-hand corner of the web interface.

Repair:

Will help you set the unit up for the type of program you have if you have audio that needs some repairing before you process it. This is usually the case with almost all audio. Do not assume your audio does not need to be repaired.

Processing:

Helps you get the sound you want. You can use the processing preset to select the type of sound you want and then make changes from there.   We highly recommend that you use this to start with as a basis for your on-air sound. We do not recommend starting from scratch. Typically for most users, the quick adjust page will suffice for tweaking any audio processing settings.

Broadcast:

This allows you to quickly select the settings you need to get on-air.

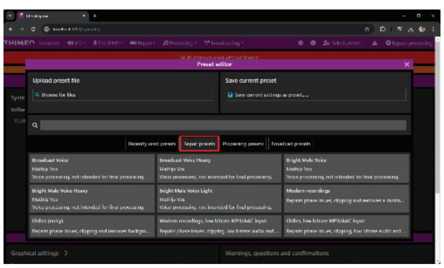
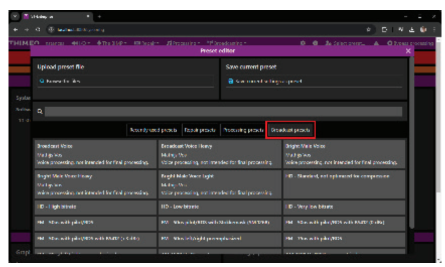
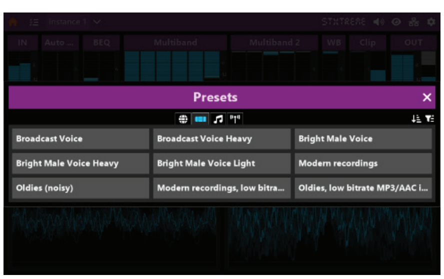

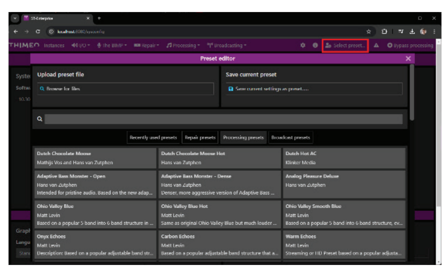

Overall notes:

Certain settings appear in almost all pages for both repair and processing settings.

The title bar typically has the following settings for each filter:

ON/OFF button: Every filter has an on/off button. If it is turned off, it is completely bypassed, so if every filter in the chain is off, it is perfectly identical to the input.

≠ **Difference mode icon:** If you want to hear what a filter adds or removes, click on the difference mode icon.

Monitor icon: If you want to hear the output of a filter without any processing, click on the icon. This allows you to better focus on the effect of that specific filter, rather than the effect of the entire processing chain.

Location icon: Some filters can be moved in the processing chain.

For filters that operate in multiple bands, you can listen to the output of each band separately. For this, click the icon for that specific band. Please note that in this mode, subsequent processing is not disabled unless you also click the headphone icon in the top bar.

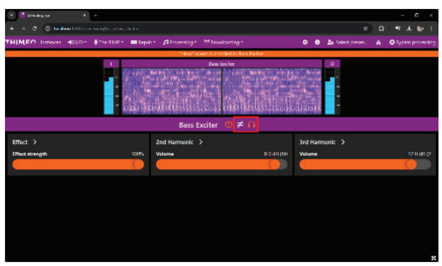

Difference (left) and Monitor (right) buttons of the Bass Exciter, located in the title bar of the filter.

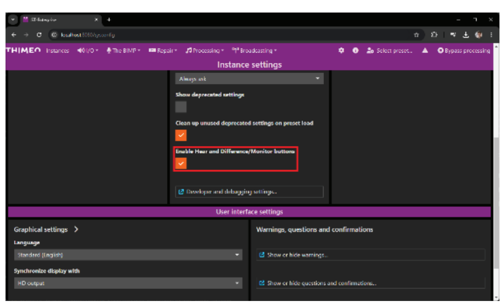

The Difference/Monitor modes (and their buttons) can be disabled outright under "Instance settings."

REPAIR section

The STXtreme is unique in that we repair the audio prior to processing it. This allows for more consistent, not to mention better sounding audio. In Repair you will find the most powerful restoration tool in STXtreme- The Perfect Declipper. Perfect Declipper repairs clipped audio, removes distortion and restores dynamics. Other tools in the repair section to help your audio be pristine includes Dehummer and Delossifier.

Repair overview:

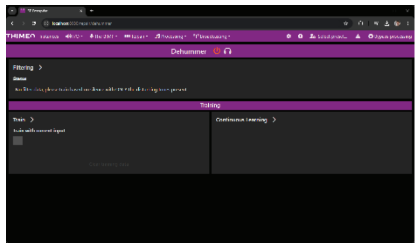

Dehummer:

Enables PNR Noise & Hum removal.

The “Dehummer” can be used to learn and remove any constant noise from a source. For constant noise at specific frequencies, the noise can usually be filtered entirely from the source audio. For other noise (such as hiss or other broadband noise) it is typically possible to achieve several dB of attenuation without significantly affecting the source audio.

Learn (On/Off):

Used to “learn” the noise signature to be removed. With only the noise source applied to the input, turn the “Learn” control “On” for several seconds. This samples the noise such that a filter can be constructed to remove it.

Reset Analysis Data:

Once a particular noise signature has been “Learned” it becomes a single, fixed filter. If the input source changes over time, the noise signature will need to be updated as this filter may no longer apply. This reset control will discard any learned analysis data in preparation for learning a new noise signature.

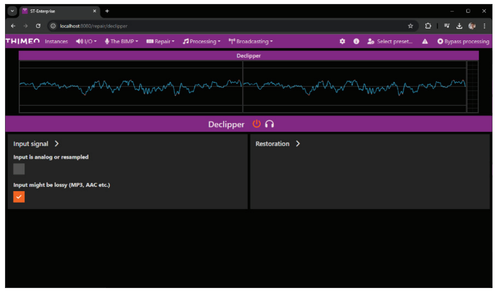

Perfect Declipper:

Perfect Declipper removes audio distortion and restores dynamics caused by clipping. Nearly all music from the last 30 years has digital clipping, which makes it challenging to achieve high-quality sound. Clipping distortion is often amplified during sound processing.

Although perfect reconstruction is not possible because some data is lost, Perfect Declipper is called ‘perfect’ because we think we can mathematically prove that its algorithm approaches the best possible reconstruction, given the available information.

Why should I use it?

In the 1980’s, CDs were expensive and only available in high-end systems - and hence most CDs were recorded at the best possible quality. Since then, due to what is called the ‘loudness war’, music has been released at continuously increasing volume levels. This has come at a cost: Reduction of dynamics and clipping. Clipping means that the loudest spikes in the music are cut off, which causes digital clipping distortion. In the last few years, it has gotten so bad that in some cases you can even clearly hear the distortion on laptop speakers.

Perfect Declipper can restore the clipped parts of the audio, in many cases the result is indistinguishable from the original, not clipped, recording.

Although FM processors typically perform clipping as a final processing stage, many modern recordings are already clipped before they even reach the listener. The Declipper can account for this and repair clipped audio prior to processing.

Input might be lossy (MP3, AAC, etc.)

Enable this setting if the incoming audio has been encoded to MP3 or other bit-reduced format at some point. After lossy encoding such as MP3, sounds are somewhat spread in time. If a sample is clipped, the samples adjacent to it will also receive a part of the distortion. As a result, samples close to clipped samples need to be treated as clipped as well.

The expert settings for this filter have been optimized to work in a wide variety of situations. There is normally no need to adjust them.

Is declipping useful if we are going to clip the audio again anyway at the end of the processing?

Yes! STXtreme’s Advanced Clipper detects whether distortion is noticeable and usually does not cause audible distortion.

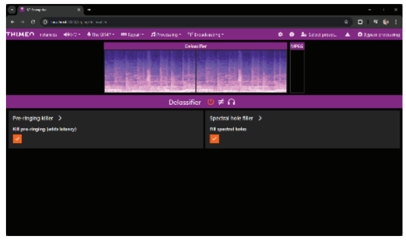

Delossifier:

Lossy audio encoding (such as MP3) reduces the amount of data needed to store or stream audio by using psycho-acoustic models to ‘throw away’ content that the codec deems ‘nearly inaudible’. The trade-off of such data reduction schemes is that they can cause audio artifacts (especially at lower bit rates). In addition to the fact that these artifacts are often already quite audible, any processing that is done afterwards may further accentuate such artifacts, countering the codec’s assumptions of what is or is not audible.

The Delossifier detects and cleans up artifacts introduced due to lossy audio encoding before the audio hits the processing stage.

The expert settings for this filter have been optimized to work in a wide variety of situations. There is normally no need to adjust them.

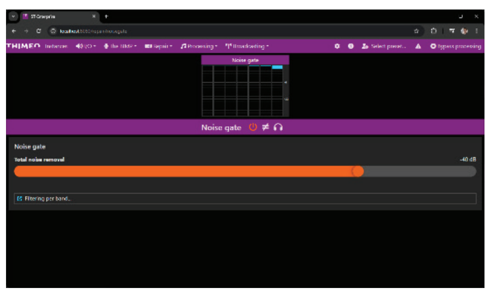

Noise gate:

The noise gate removes background noise, which is mainly useful for older recordings. But even for recent CD’s it helps to let it work very gently, because quantization noise (the number of bits on a CD is not infinite) can be increased a lot by processing which makes music sound slightly harsh. Noise gate can be turned on or off

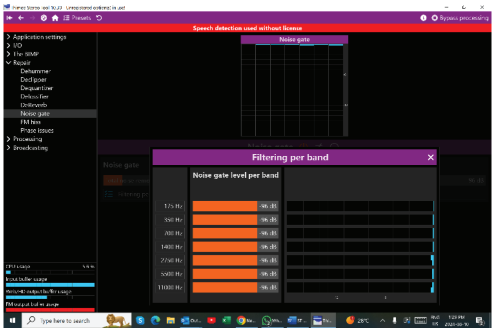

Total noise removal:

The maximum amount of noise to be removed. If you click on this, you will be able to adjust individual frequencies.

Typically, if normal audio is playing, there should be very little noise gate action in the meters. At most 1-2 dB, occasionally. During quiet sections (even if they are very brief), a lot more action is ok, because that is when there's almost only noise.

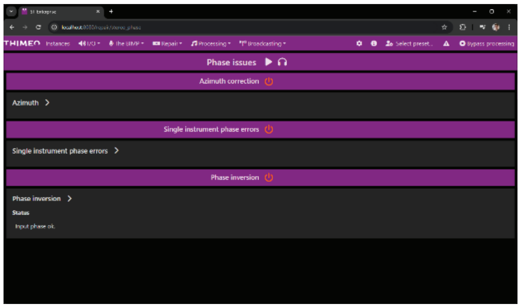

Phase issues:

There can be several phase issues that this section deals with. For example, one channel might be inverted. Another example is music that has been mastered on tape and has phase shifting which Azimuth could correct. Additionally, some songs have been recorded with extreme stereo phase differences which causes problems when the reception goes to mono; this can be remedied as well.

Azimuth correction:

AZIMUTH errors are often present in tape recordings, and on some (typically cheap) CDs. Phasing problems causes playing a recording in mono or through a surround system to result in very ugly artifacts. But even normal stereo playback may sometimes sound a bit unpleasant. The phasing offset is automatically detected and removed by this filter.

This filter only works properly if the sounds at the left and right channel are similar. If this is not the case for a longer period, the azimuth correction will slowly be reduced.

The expert settings for this filter have been optimized to work in a wide variety of situations. There is normally no need to adjust them.

Processing

For processing, we highly recommend selecting a preset that is close to the sound that you want to achieve and then fine tuning from there.

Quick Adjust:

The quick adjust page allows you to fine tune without going very deep. For most users this will suffice.

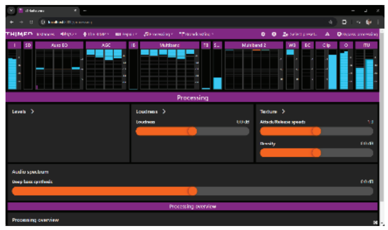

Processing Levels: Input and output levels. We do not recommend making any adjustments to these.

Texture: Simple texture controls.

Attack and release speeds: Affects how fast the levels are adjusted, makes the sound aggressive.

Density: Controls the density, this controls how deep we go into the multiband compressors.

Lows Density: Controls density of low bass frequencies

Mids Density: Controls density of the mid frequencies

Presence Density: Controls density of the presence (mid high frequencies)

High Density: Controls density of the highest frequencies

Loudness: Makes the sound louder or quieter by adjusting the clipper drive. If you drive this too hard it will affect the sound quality.

Deep bass synthesis: Controls how much deep bass is reconstructed if it was removed from recording. For example, with older material the mastering process removed deep bass for practical reasons at the time.

Presets:

The best way to get started with the STXtreme is to select a preset. If you click on any preset, it will give you a description. This allows you to select a preset close to what you want in structure and sonic signature. Once you select your preset then you can use the simple, easy to use adjustments in the processing overview menu.

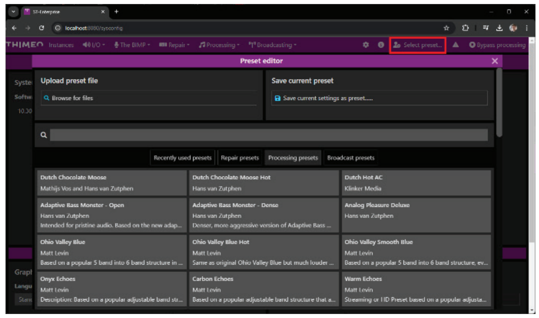

For the most part, you will not need to have a deep dive into the rest of the parameters, but if needed, we have lots for you. Let's get started.

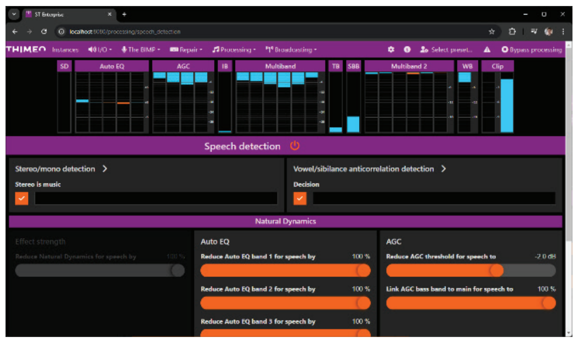

Speech detection:

Detects when voice alone is transmitted and allows you to change certain settings based on whether the content is music or speech.
Typically, you will want different processing during speech (the news for example) than for music. For example, for speech you probably do not want a lot of bass enhancement. The speech detection section detects when speech is being transmitted and allows you to change certain settings based on whether the content is music or speech.

Stereo/Mono detection:

Detects if incoming audio is mono. Stereo sound will never be considered to be speech. Because music beds can be stereo, there are some thresholds that can be adjusted to match when we consider something to be mono vs stereo. The default settings are good for most stations.

Vowel/sibilance anticorrelation detection:

Speech generally consists of vowels and sibilances, which are rarely uttered at the same time. That is very different from music which will have strong tones with for example hi-hats on top of that. This detection uses this to determine when the input appears to be speech. The default settings are good for languages such as English and Dutch. For Spanish, somewhat different settings are needed because many sibilances are much softer.

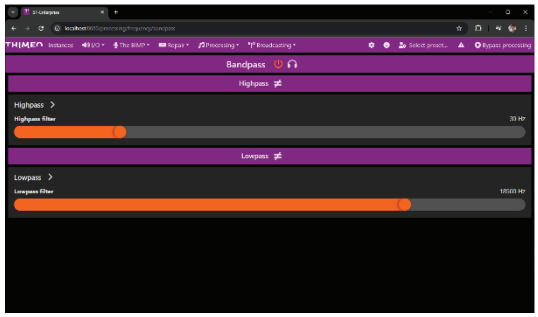

Bandpass:

Configures filters to remove very low or very high frequencies.

Highpass filter:

Removes very deep bass sounds. Controls the highpass frequency. Tones below this frequency are removed.

Lowpass filter:

Remove high frequencies. Controls the lowpass frequency. Tones above this frequency are removed.

This only affects the lowpass frequency for streaming/HD outputs. The FM output lowpass frequency is determined elsewhere.

This filter is very steep. The volume starts to drop a few hundred Hz below the configured frequency, and no frequencies above the set frequency should be coming through. The lowpass filter is always phase linear.

Equalizer:

This controls the equalizer which can be selected on-off. You can then see the parametric EQ and use parametric settings. You can also set it to manual mode, in which case you can draw the response curve graph yourself.

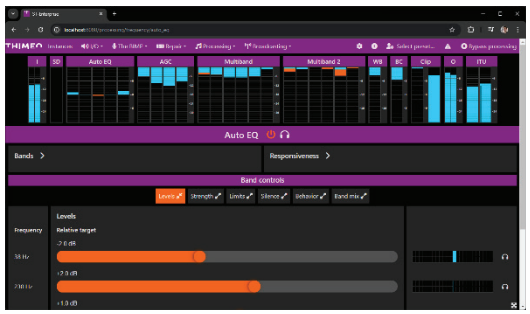

Auto EQ:

Adjusts the spectrum without compression, making it possible to generate a very consistent sound without sounding compressed.

This filter is a dynamic re-equalizer that operates before the AGC in the processing path. Using the familiar Bands user interface found in MB1 and MB2, it can be configured to precisely match bands already in use elsewhere. Subject to user settings, this filter will attempt to keep levels in the bands at a target level with respect to each other. Spectrum Correction is not a compressor and does not affect actual levels, overall. Furthermore, using it does not require any subsequent filters to be altered to accommodate it. If this filter is boosting the level in some band(s), it is reducing some other band(s).

Without spectrum correction, the multiband sections are responsible for two jobs: 1) Controlling audio levels, and 2) Dynamically re-equalizing incoming audio. This is usually not a problem, but it can become one on some material.

Let's say incoming audio usually produces about 10dB of gain reduction in MB1 and all bands are roughly equal in the amount of gain reduction they're doing. Now a song comes along that has very little high end. The top band of the multiband section will release and could end up doing almost no work at all. That is suboptimal, because the multiband now has no margin of error.
If the number of highs drops even further, there is nothing the multiband can do to correct it, and the audio will become muddy and dull. That is where spectrum correction comes in, to fix material that is imbalanced before it hits any further filtering. When audio is playing that is properly

balanced, per the settings, this filter will have virtually no effect on the audio. When doing large amounts of correction, the filter should be unobtrusive. It simply fixes the audio relatively slowly while attempting to be as transparent as possible. The user interface for this filter resembles the multiband compressor UI but has some important distinctions. When the filter is doing no work, you will see nothing in the interface (no colored bars). Bars stretching to the right, in shades of blue indicate a boost in level. Bars toward the left, in shades of red, indicate a cut to the volume in that band.

When audio is playing that is properly balanced, per the settings, this filter will have virtually no effect on the audio. When doing large amounts of correction, the filter should be unobtrusive. It simply fixes the audio relatively slowly while attempting to be as transparent as possible.

The user interface for this filter resembles the multiband compressor UI but has some important distinctions. When the filter is doing no work, you will see nothing in the interface (no colored bars). Bars stretching to the right or top indicate a boost in level. Bars toward the left or bottom indicate a cut to the volume in that band.

The Auto EQ is especially useful when used together with our older compressor designs. If you use the multiband compressor in Adaptive Compressor mode with infinite ratio, Auto EQ is not necessary.

Relative target:

Target level for the Auto EQ.

Auto EQ will keep the level in bands right where desired. This slider determines the appropriate level for each band on a +/- 12dB scale. Where the sliders are in that scale is irrelevant. The only thing that matters is the relationship of the bands to one another. For example, in a two-band scenario if band 1 is at +2dB and band 2 is at 0dB, that identical to band 1 being at 0dB and band 2 being at -2dB.

Important: The best results are obtained if, on average material, the meters are all in the center. The default settings are generally ok, but for different types of content some adjustments may be needed. To determine whether the levels are in the center on average, look at the median display.

PHASE ROTATION:

Makes the peaks of asymmetrical waveforms more symmetrical. In other words, if a peak has a high peak in one direction, this will change the waveform such that the peaks are close to symmetrical, while changing the audio as little as possible.

This protects the compressors, limiters and clipper against certain types of sounds that can be difficult to process and tend to cause distortion. Examples are voices (especially female voices) and instruments like trumpets. These types of sounds are very asymmetrical with large peaks in one direction. If there is also lot of low frequency energy present (such as a song with heavy bass) it can cause additional issues with these waveforms due to intermodulation distortion.

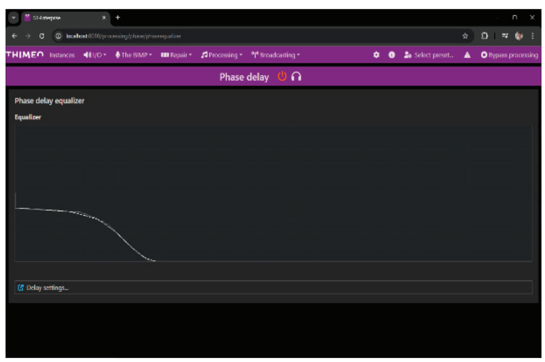

Phase Delay:

Phase Delay allows you to purposely introduce nonlinearities in frequency response. This can for example stretch out a bass sound by playing the initial “kick”, which is usually at a higher frequency, slightly earlier than the “boomy” sound that comes afterwards. This changes the sound of bass and can also make them survive the final clippers better, because the low frequency sound will not push the initial kick down. It can also be used to emulate the nonlinearities found in certain analog processors. The horizontal axis of the graph is frequency, and the vertical axis is the amount of additional phase delay to be applied at that frequency. Curves can be modified by drawing them directly on the graph with the cursor.

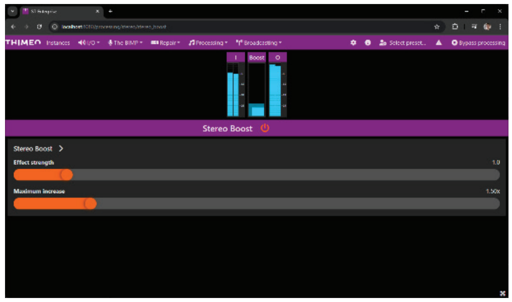

Stereo Boost:

Stereo Image is basically just a multiband compressor which increases the level of the L-R channel. It can be configured to never reduce stereo separation (in tracks with completely separate channels), and it does also not create too much (anti-phase) boost of the L-R channel.

Effect Strength:

Configures how strongly stereo boost works. Higher values sound wider but also increase certain sounds and may cause multipath problems with FM reception.

Maximum increase:

The maximum increase of the L-R channel.
Reduce resulting stereo to multiples the L-R channel with this value. This can be used to reduce extreme stereo separation.

Stereo Image:

Stereo Image gives independent control over phase differences and instrument placement. **But it can easily create artifacts.**

The most interesting feature is to convert stereo to mono without cancellation (loss) of sounds. This creates a sound that is as full as the stereo image, but it is in fact mono. This can be done by setting both **Image phase amplifier** and **Image width amplifier** to 0.00. If you are not located between the speakers, when you press the MONO switch on your audio system you often still hear a big difference, because a lot of the sound disappears. After setting these two sliders to 0, that does not happen anymore. Example uses are radio stations (FM, AM, streaming) that broadcast in mono.

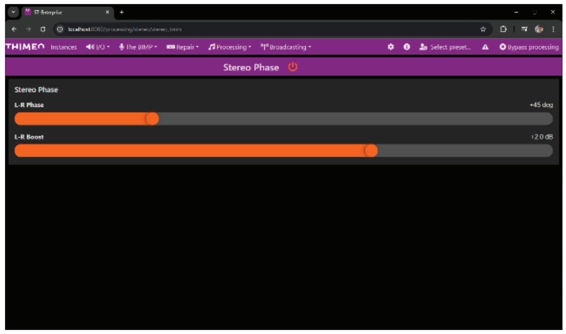

Stereo Phase:

Stereo Phase is a filter that changes the stereo image. It does not really make it more or less strong, just different.

L-R Phase:

How much phase difference it adds.

L-R Boost:

Because changing the phase might slightly reduce perceived stereo separation, a small boost in stereo separation might be in order to compensate for that.

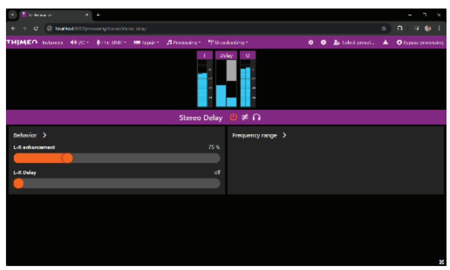

Stereo Delay:

Stereo Delay adds a copy of the L-R (stereo, difference between channels) information to the original stereo signal. This creates a sparkling stereo effect. In cases when the stereo effect in the input signal is already very strong, or almost not present, this filter switches off automatically. On top of these things, the stereo effect is stronger for transients.

This type of stereo widening has very little impact on reverb (which often gets boosted by other types of stereo widening). When used for FM, it also has very little impact on multipath distortion.

L-R enhancement:

Maximum boost of the L-R signal that's added to the original stereo signal.

L-R Delay:

Amount of delay of the L-R signal that is added to the original stereo signal.

When using more delay, the sound appears to 'open up', however it might start to sound like an echo, especially for high frequency transients.

When using 0, there is no delay and the existing stereo information is boosted.

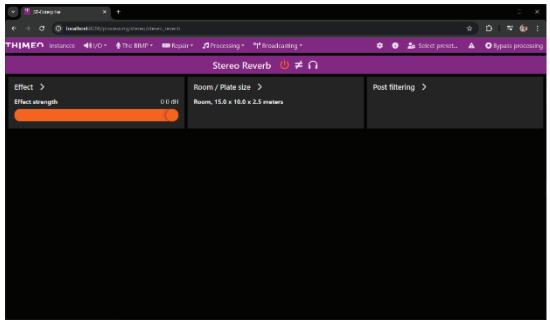

Stereo Reverb:

This adds reverb to the existing L-R difference signal. The benefit of this is that it does not typically add reverb to things like vocals, only to real instruments.

Using a very small amount of this is usually more than enough, otherwise everything will start to sound very reverb-y.

Effect strength:

This controls how much reverb is being added.

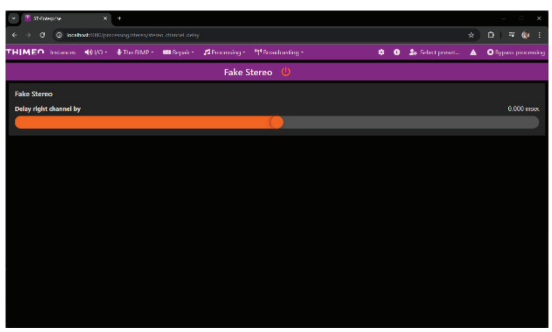

Delay left channel by:

Delay between the left and right channel. This introduces a (very rudimentary) stereo effect. Note that this also changes recordings that are already in stereo, that it also makes sounds like voices stereo (which is generally considered bad), and that the result will sound very bad when played back in mono.

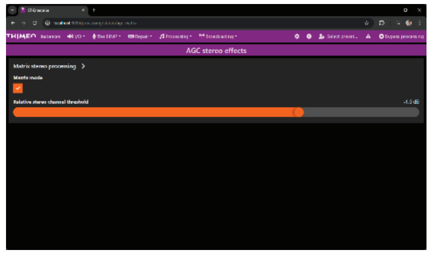

Delay left channel by:

Delay between the left and right channel. This introduces a (very rudimentary) stereo effect. Note that this also changes recordings that are already in stereo, that it also makes sounds like voices stereo (which is generally considered bad), and that the result will sound very bad when played back in mono.

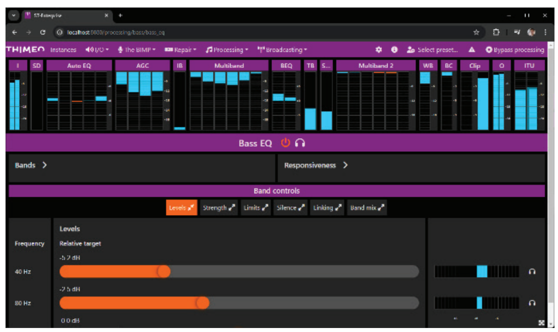

BASS EQ:

Bass EQ is similar to auto EQ. See Auto EQ for description. The purpose of it is to add bass if the material is lacking bass. The default mode is 2 bands, but the number of bands can be changed. The topmost band can not be changed. This is a reference band. The bands below it can all be changed. Band 1 and 2 can be changed and center frequencies can be moved as well.

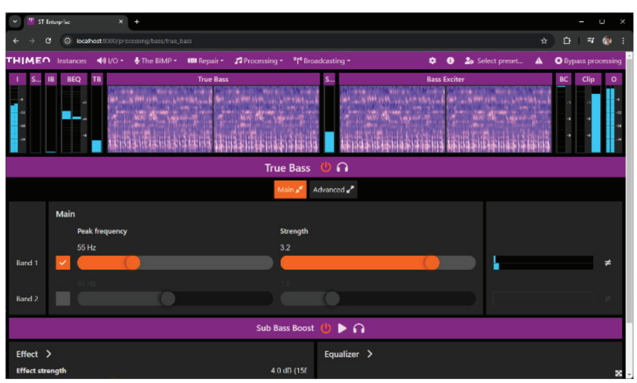

True Bass:

Lost subharmonics generator. It is fixed at two bands. True Bass attempts to generate lower harmonics (deep bass tones) that either were not recorded properly (due to for example a highpass filter, a bad microphone) or appear to be missing.

It can generate bass at a specific frequency, which matches the rest of the signal. True Bass was designed to only generate bass that sounds like it should always have been there. It typically will have no effect on vocals, only on actual instruments that play at very low frequencies.

Some examples: Input: Square wave, 50 Hz, with the base frequency (50 Hz) filtered out - only 150, 250, 350, 450 Hz etc. remain. True Bass will recreate a 50 Hz tone in this situation:

Input:

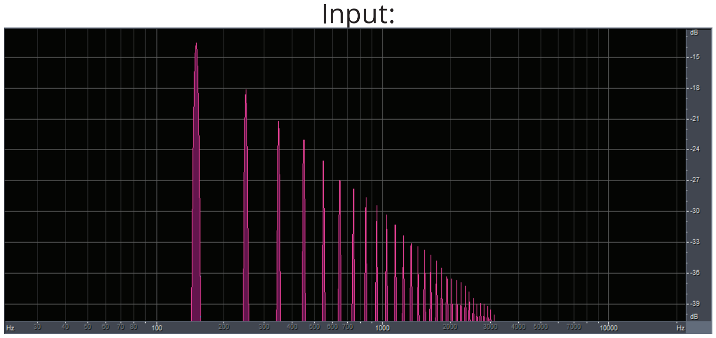

Output:

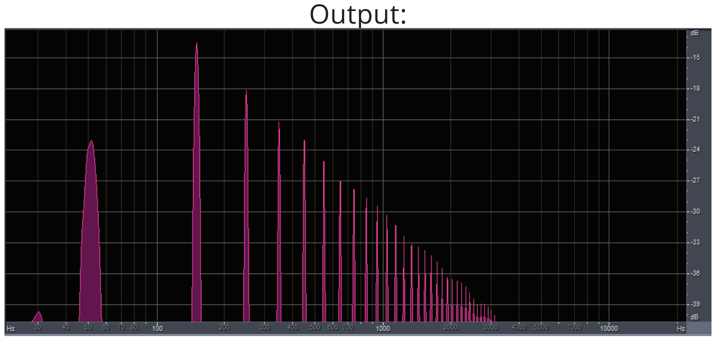

However, if instead of a square wave, we feed it a sine wave at 150 Hz, it will not generate a subharmonic:

Input:

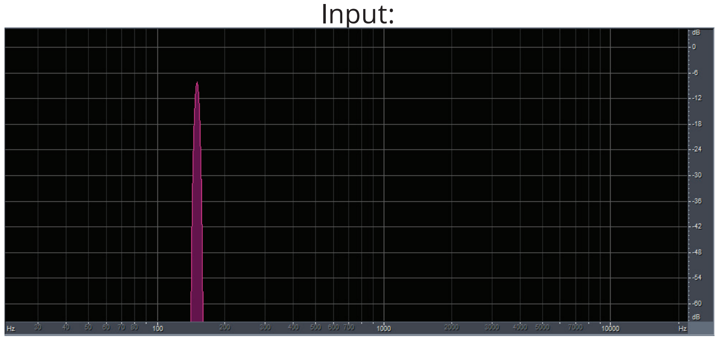

Output:

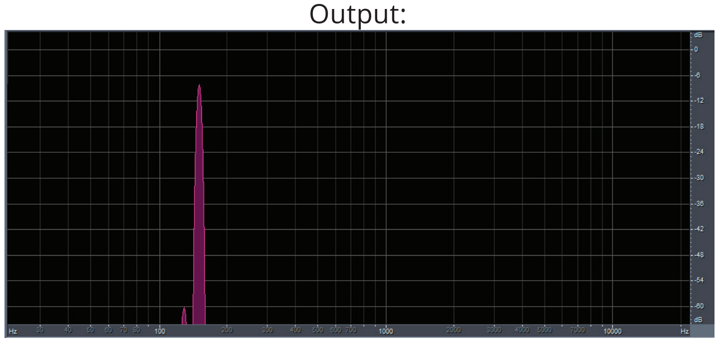

Peak frequency

The frequency around which this True Bass filter works.

The filter drops off pretty steeply to higher frequencies, but also to lower ones.

Strength:

The amount of effect that True Bass has. Use caution as using more can give unnatural effects.

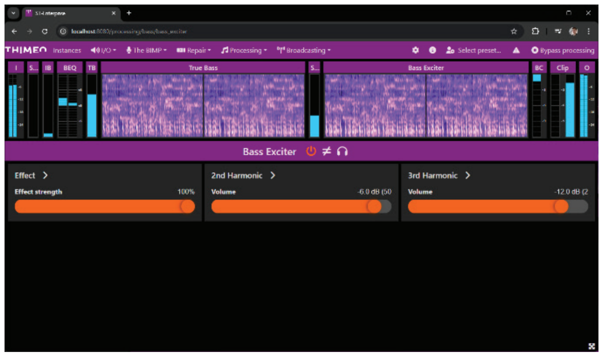

Bass exciter:

Creates upper harmonics, which generates “warm”, “car radio” bass. This is useful for small speakers on which you cannot hear the original lower bass. It can also give a bit more warmth on larger speakers.

2nd Harmonic:

The second harmonic adds warmth to the sound. This is useful for bass, but can also be used for vocal, to add some warm sparkle. But you should use very little of it on non-bass frequencies, clearly less than –30 dB.

3rd Harmonic:

The 3rd harmonic should typically only be used for bass, it doesn’t sound too nice on higher frequencies.

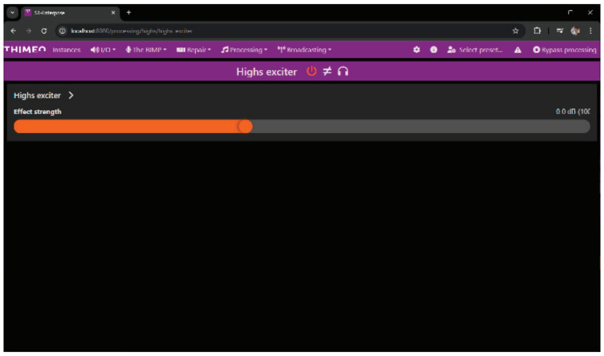

Highs exciter:

This creates high frequencies. It is useful for old recordings. Using too much can make things sound overly bright and annoying, so only use very small amounts.

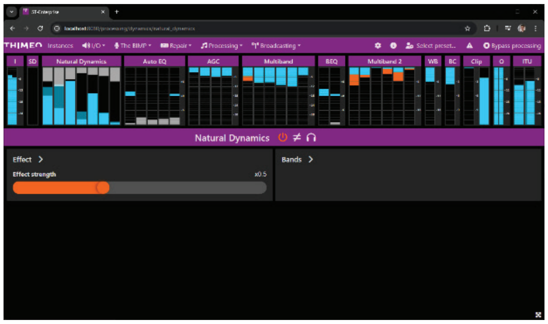

Dynamics

Natural dynamics:

In addition to being heavily clipped, modern recordings often lack any dynamic range. It is much easier for an audio processor to work with music that has some dynamic range rather than trying to further compress already heavily compressed recordings. “Natural Dynamics” restores dynamic range and “punch” to compressed recordings whilst not altering audio containing sufficient dynamic range.

Effect strength:

Increases or decreases the overall effect of “Natural Dynamics”.

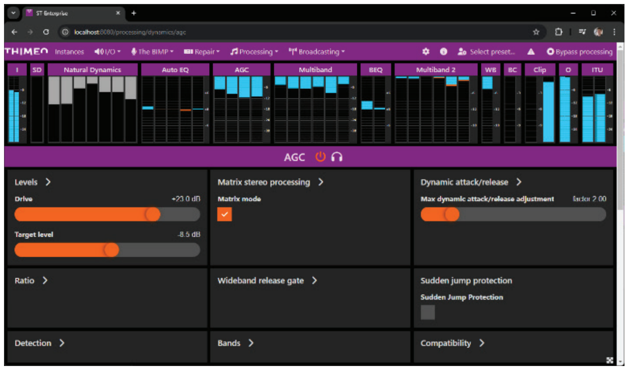

AGC (Automatic Gain Control):

With the Multiband compressor running in Adaptive Compressor mode with infinite ratio, the AGC really is not needed anymore; the Adaptive Compressor can basically handle anything you throw at it.

The AGC attempts to make the audio level constant by gradually increasing and decreasing the gain. If sudden level jumps are present in the audio, it can remove these transients without lowering the output level too much. This avoids pumping and similar annoying effects.

Drive:

Sets the AGC input level.

Most of the included presets were built assuming that the input level for music typically peaks to about 0 dB. Use the AGC drive to drive it harder that that.

Target level:

Determines the target maximum output level.

Attack and Release: these settings control how fast/aggressively the AGC responds to increases or decreases in input loudness.

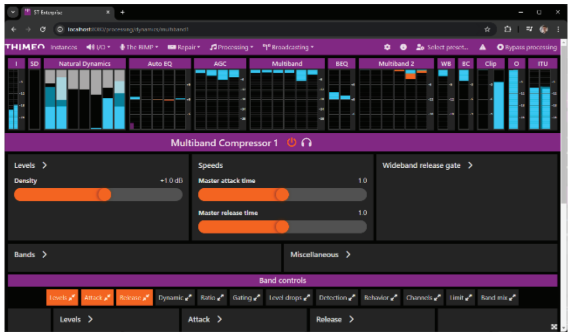

Multiband compressor 1 and 2:

STXtreme supports multiple compressor designs. One of those, the Adaptive compressor is very different from traditional compressors. It can operate on its own without an AGC. With our compressor, you tell it what you want the end result to be as opposed to telling it what to do which is what older traditional compressors do. This section describes the rationale behind this compressor. The most important settings are described below the description, you can skip to there if you just wish to know want to know what these settings do.

Attack behavior:

In adaptive mode the compressor constantly determines the optimal attack speed. This handles extreme content much better. When presented with extreme content the behavior is not very sensitive to the settings. It will move down rapidly almost regardless of the settings. But if the level difference is small, it moves very slowly - which protects already dense material.

With this design you can get away with a wide range of attack speeds.

Release behavior:

The release behavior mostly ignores the input level, and acts on very different parameters. Recovery after a brief loud S sound is quick almost regardless of the settings, and there are no notable gaps after hi-hats either. However, for already dense content, release can be configured to be very slow.

Too quiet dense audio:

Occasionally very dense audio comes out quieter than dynamic audio. This is easy to explain: During dynamic audio we constantly release, and if after releasing a new loud sound starts it will play at a higher level because we still need to attack to push the level down. For dense audio, the compressor does not release much, so it stays all the way down.

If you actually measure the level of the output audio over some amount of time, it is often not even that far apart. But to our ears, it sounds much louder because we focus on the spikes in the audio, not on the average level; we kind of ignore moments of silence.

A workaround for this is to use lower ratio's; if you use a lower ratio, the loudest (typically densest) parts of the audio will come out louder compared to not so loud input. So, with an AGC set to infinite ratio, followed by a compressor with a lower ratio, this problem is mostly gone. But this is definitely not ideal, and it hides the issue instead of solving it.

There is a setting “Make dense audio louder”, to handle this. It also makes the audio sound better, because even in dynamic audio there can be moments of dense audio (such as loud vocals), which come out louder as well, so vocals are stronger with this setting, and cause less “ducking” in the background music.

Windowing:

With our design you can use a 12 dB window on the final compressor without causing any major issues; all it does is to make the audio more dynamic. So, this works great in combination with very fast attack and release times.

Typically, first set up the compressor to sound good overall. Then add some windowing to create more punch.

Gating:

Casper the Friendly Gate

Gates are scary. They can get you stuck at some level. So, we have Casper the Friendly Gate, which isn't scary at all. It acts indirectly on internal controls, which makes its behavior very smooth. If you use high release speedup values for dynamic audio, using the Casper gate will tame it down a lot, but without loosing the speedup.
So always set this up first. And then use the other 2 gates if needed.

Input gate:

This is an input gate as it's also present in AGC's. If the input level drops below some threshold we slow down or stop.

You can set a target attenuation level to move towards (“return to home”) during gating. That only responds to the input gate.

Dynamic gate:

The Dynamic gate responds to the average level over some amount of time, which is a lot safer. Still, use this one wit care - this is the only gate that can get you stuck at a weird level.

Level drops:

Use this to release faster if the level drops a lot. Use gating to control when the speedup stops.

Detection:

The setting that controls if you are using RMS or peak mode.

Important: Even in peak mode, the RMS size is still being used! So, make sure to set that up even in peak mode.

Channel and band linking:

Changed band linking uses differences in dB's.

If you are using multiband compressor 1 as an AGC as well, you'll probably need to set up band linking to make it sound good; without it the bands can get very far apart, which tends to sound bad.

Limiting:

We have 2 different types of limiters in the Adaptive compressors.

Intellilimiter:

The Intellilimiter works together with the main compressor in an intelligent way. If we limit a lot, that also affects he total level of the compressor. Which in turn responds by releasing faster, so we build a bit more density, but without constantly limiting a lot because we go over the threshold. With how the Intellilimiter works together with the normal compressor, the end result sound completely natural, and not limited at all.

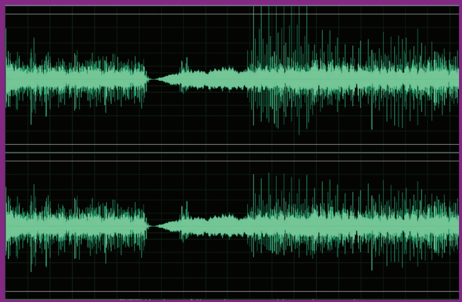

Here is the same audio, processed without Intellilimiter on top, and with Intellilimiter at the bottom. What is clearly visible is that not just the highest peaks got reduced (as a direct effect of the limiting), but lower peaks are also reduced. The end result is a more dynamic, less restrained sound. Also, short loud sounds may not have much effect on the measured level, but they still sound loud to our ears, so reducing them and everything around them sounds more natural.

While it is not visible in this waveform, the quieter parts are still quieter. In that regard the Intellilimiter behaves a lot like the Casper gate; it tames down the audio. But where the Casper gate looks at how quiet the quiet sections get, Intellilimiter instead looks at how loud the peaks are after compression.

Using the Intellilimiter has more impact on overall loudness than using a normal limiter, since it also affects the audio between loud sounds.

The Intellilimiter can run in peak or RMS mode. Since the compressor responds to Intellilimiter, if the latter 's running in peak mode, some peak effects will be introduced in the compressor behavior.

Limiter:

The (normal) limiter is mainly used to protect against very loud sounds, which would sound too loud or cause problems for the final clipper. We recommend using this in tandem with the Intellilimiter, to precisely control how much "real" limiting you want to perform. Where the Intellilimiter mainly affects loudness, the real limiter has much less impact on loudness but will hurt the audio quality more. Using the two together gives you very precise control over how much you want to have of each.

The limiter always runs in peak mode, the Intellilimiter can run in both peak or RMS mode. If you are planning to combine the two, it is probably a good idea to run the Intellilimiter in peak mode and use the same attack speed for both. If you do that, the maximum amount of real limiting should be close to the threshold difference between the two. So, if the Intellilimiter is set to 2 ms, peak, 10 dB, and the real limiter is set to 2 ms, 12 dB, typically the real limiter should rarely have more than 2 dB effect on the output.

Density:

Changes the input and output levels of the compressor such that the total output level is not affected but compressor starts to work at much lower input levels. Works on all bands simultaneously.

Master attack time:

Value with which to multiply all attack speeds. Higher means slower attack.

Master release time:

Value with which to multiple all release speeds. Higher means slower release.

Threshold:

The input level above which the compressor becomes active. You should adjust these settings have more or less of each band in the output. We recommend configuring these settings such that each band is attenuated by the same amount. The average attenuation is shown by a white line; right click on the display to reset the measurement.

Attack time:

This has to do with the amount of time for the signal to be compressed. Slower times will give you more dynamics, but it is very important that attack and release are working together in harmony.

Release time:

This controls how fast the volume is raised when the input level goes down. Faster times will give you a more aggressive but potentially more consistent sound. Again, it is very important that this is working in harmony with attack.

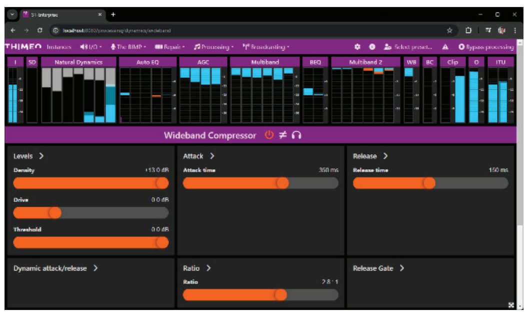

Wideband compressor:

This is not used often. This is a broadband compressor. Reduces dynamic range of the audio and limits it. Volume compression reduces the dynamic range of a sound. Limiting limits the maximum audio level below a certain threshold.   Because it is wideband it operates on all audio at once, so for example loud bass can push everything else down. This can be used to create a "hot" pumping sound.

The settings are the same as for the multiband compressor.

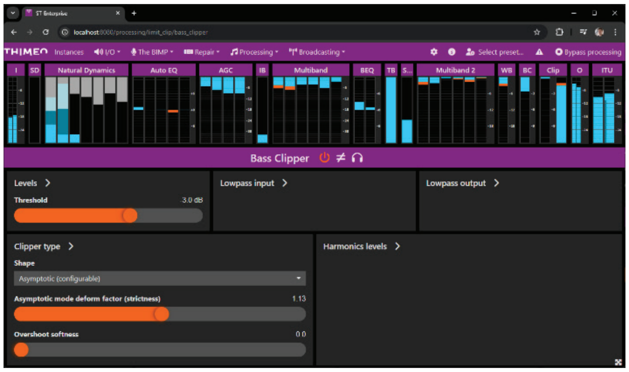

Bass Clipper:

Clips the bass before the final clipper. The purpose is fourfold. Protects the final clipper against excessive bass, control the bass tightly so it is constant in level, generate a small amount of harmonics to make the bass more audible on small speakers, and replicate specific bass flavors from other processors or create your own. This will also allow you to make adjustments in your audio that could make your processors replicate the sound of other processors on the market should you wish to do so.

It is possible to meticulously control for each separate harmonic how loud you allow it to get. This allows you to precisely control the sound.

Levels:

Threshold: The threshold for the bass clipper. If a bass sound waveform peaks above this level, is it clipped.

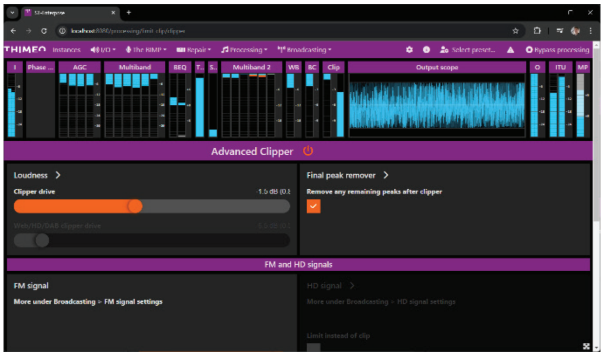

Advanced Clipper: The advanced clipper is different than most older traditional processors in that it does not generate harmonics distortion unless you drive it very hard. Because of that, you can achieve more loudness on FM and safely use it for streaming/HD outputs.

On top of that the composite clipper in STXtreme is very different from older traditional clippers and can give you 2 to 3dB extra loudness on air at the same quality. Additionally, it can optimize for better reception in fringe areas and reduce multipath issues.

We recommend before you tweak your audio you determine whether you can send audio via composite to the transmitter site. That is a prerequisite for using the composite clipper.

The following is an in-depth description of the clipper and the advantages of composite clipping. The clipper settings are described below this section, and you can skip to those if you just want to know what the settings do.

The advantages of composite clipping

Compared to traditional left/right clipping, composite clipping can give you more than 2 dBs of extra loudness without sacrificing audio quality or dynamics.

The pilot and RDS take up space.

If you have a signal that is clipped in left/right, and after that is stereo modulated, you have a signal with a fixed peak level. But after that, you still need to add the pilot and RDS signals, which together require about 13.5% of the total signal (assuming 9% pilot + 4.5% RDS injection). The image below shows what happens: You have the full stereo modulated audio signal (the gray block), and you have to add the pilot and RDS (the green line) to that. The green area surrounding the gray block is an area that has been unnecessarily cleared by the left/right clipping - had you known in advance that the pilot and RDS combined were moving in the opposite direction, you could have raised the clipping thresholds at that point.

To clarify: Say your total signal must stay below 1.0. If at some point, the value of the pilot+RDS is -0.1, then your audio could have gone up to 1.1 at that point. Instead, because the pilot+RDS can go up to 0.135, the audio was clipped at 0.8865. That is a difference of more than 0.2, which is 2 dB. Note that this is not possible all the time, there are (rare) locations where you actually do need to keep the audio below 0.8865. But it is never less than that - and almost always more.

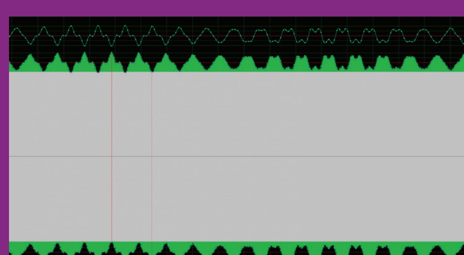

Room needed for pilot and RDS.

Making use of the stereo modulation

It gets worse though. Consider a waveform that peaks much higher on one channel than the other. See the 2 images on the left top. The thin line is the center of the waveform, the thick blue line is the clipping threshold. If you were to clip it like that, you would get the results that you see in the bottom right.

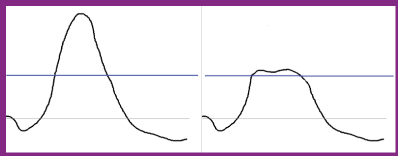

In the 2nd row, you can see what can happen to the L-R after stereo modulation: If it is short enough, the peak has a 50% chance of going in the opposite direction. On the right you can see what this means after modulation (without clipping): There is almost no overshoot left anymore, so clipping does not have much impact yet.

In the 3rd row, you can see the resulting signals. Left shows the demodulated audio after composite clipping. The peaks are way louder than the threshold, and the levels are different, so the stereo effect is maintained well. On the right you can see what it looks like after traditional L/R clipping: A lot more audio has been removed and the result is almost mono.

It gets worse though. Consider a waveform that peaks much higher on one channel than the other. See the 2 images on the left top. The thin line is the center of the waveform, the thick blue line is the clipping threshold.

If you were to clip it like that, you would get this result. A lot of the signal is gone, so you'll get a very clipped sound, and on top of that, the result is nearly mono - the fact that one channel was much louder that the other isn't really visible anymore in the end result.

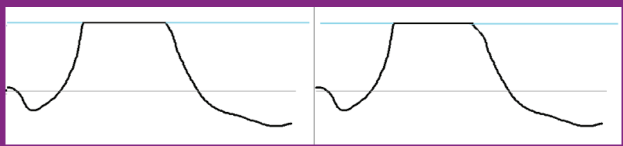

So, what do we do when we create a composite signal? Instead of working with the left and right channel, we are going to work with the L+R (just the mono sound) and L-R (the stereo sound, the difference between the two channels) signals. You can see what those signals look like on the right.

(Typically, we use (L+R)/2 and (L-R)/2).

So, let us first check what happens in stereo modulation. Basically, the L-R signal is multiplied by a 38 kHz sine wave. Which means that there is a good chance that the resulting waveform is less loud (the average sample value is 3 dB's lower), and if it is short enough, there is a 50% chance that the resulting waveform points in the opposite direction.

If you add the L+R and modulated L-R signals back together again, the result is always the same or lower as the peaks of the left and right signal, and often by a substantial amount. Clipping this resulting signal will have far less impact on the audio than clipping the separate left and right channels has.

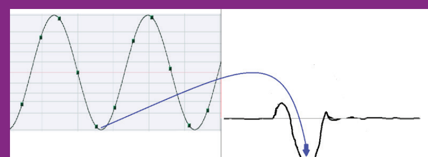

And indeed, this is what the demodulated composite clipped signal looks like in the above example. As you can see, it is very close to the original signal, and the stereo difference has not been reduced as was the case for the left/right clipped signal.

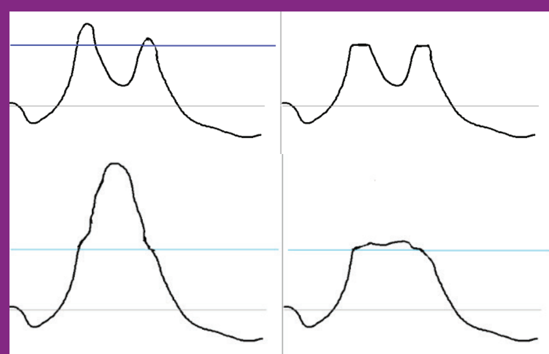

Making use of lower and upper sideband differences

We are still not done! There is another trick that we can use to make the end result even louder, with even less clipping applied. Normally, the area around 38 kHz is perfectly symmetrical: The upper and lower sidebands are exactly the same, just mirrored. It is possible to transmit only one of the two sidebands at twice the level. Most receivers (there are some exceptions) will work perfectly fine if you do that.

So, the next trick that we can use is to analyze both sidebands separately, and when one of the two causes a bigger spike than the other, quickly put in a bit less of the sideband that causes the spike and a bit more of the other. The image below shows the difference between the two sidebands, if you can dynamically - and very quickly - switch between the two, that will greatly reduce the amount of clipping that is being done and hence increase the output volume and dynamics further. The arrows show which is the optimal sideband to choose at those points in time.

Because some receivers really do not like single sideband transmissions and will start to blend to mono, we have to limit this effect. We typically keep it below at most 10% of change to the resulting sample values. At levels below 20%, no receivers seem to suffer from any negative effects. And you still gain close to 1 dB of extra loudness compared to transmitting a perfectly symmetrical signal.

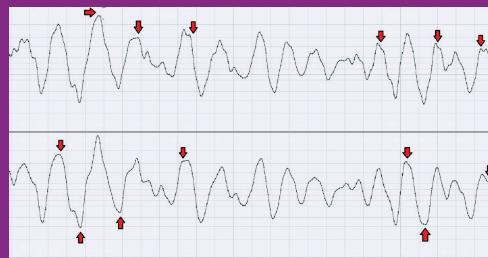

Making use of lower and upper sideband differences

A very small amount of asymmetry at the top part of the L-R area

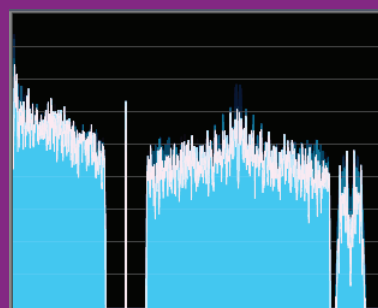

Here, the whole L-R area is clearly asymmetrical.

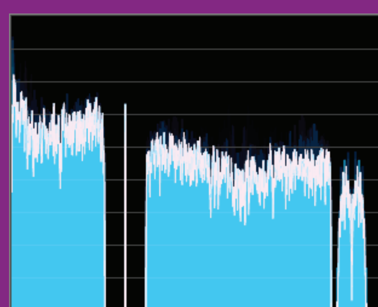

Improved reception

If the RF signal that is being transmitted gets wider, this tends to cause reception issues, especially in fringe areas an when there are multipath reception issues, and potentially interference to stations at nearby frequencies. The ITU has published a recommendation (ITU-R SM.1268) that the RF signal should ideally stay inside of. For mono content without overmodulation, this is guaranteed, but when the stereo pilot, modulated stereo signal and RDS subcarriers are added, that is no longer the case. If you broadcast outside of this mask, you are broadcasting in an area where most receivers will not be listening - they will filter it out, which causes distortion in the audio.

A composite clipper has access to all these subcarriers. That is why, inside STXtreme, we have added a software-based FM exciter and a spectrum analyzer. These analyze the output of the composite clipper and adjust the clipper parameters on the fly to force it to stay inside the ITU mask. That does add a small amount of extra clipping, but only in the brief moments where the signal would otherwise exceed the ITU mask.

This filter can be enabled or disabled. In the Thimeo STXtreme it is called “RF bandwidth limiter”.

The images below show the resulting RF bandwidth without (left) and with (right) this protection enabled.

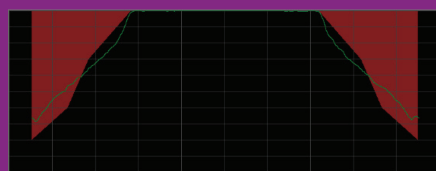

Without RF bandwidth protection. Each horizontal line is 10 dB, so the overshoots are close to 15 dB.

With RF bandwidth protection. The signal never exceeds the mask.

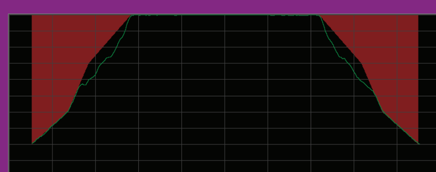

The total effect on the audio

We started this with the claim that composite clipping is advantageous to the audio. So how advantageous is it really? The biggest effect that it has is that it allows far more high frequencies to get through. In tests with extreme amounts of clipping on white noise, we measured that the high frequencies were more than 5 dB louder with composite clipping than with left/right clipping. While the lows and mids do not directly benefit from composite clipping, they do benefit from the fact that the highs get “out of the way”, so they are a bit louder as well. For music and with “sane” levels of clipping, the effect is less extreme, but there is still a very clear difference in the audio level, the audio quality and the number of dynamics in the audio.

The images below show - for several songs - the demodulated output of left/right clipping, standard composite clipping, and composite clipping with asymmetry between the upper and lower sidebands, processed in STXtreme. It is clearly visible that the peak levels are much higher - and hence the amount of clipping is much lower - when using composite clipping, and the effect gets even bigger when asymmetry is added.

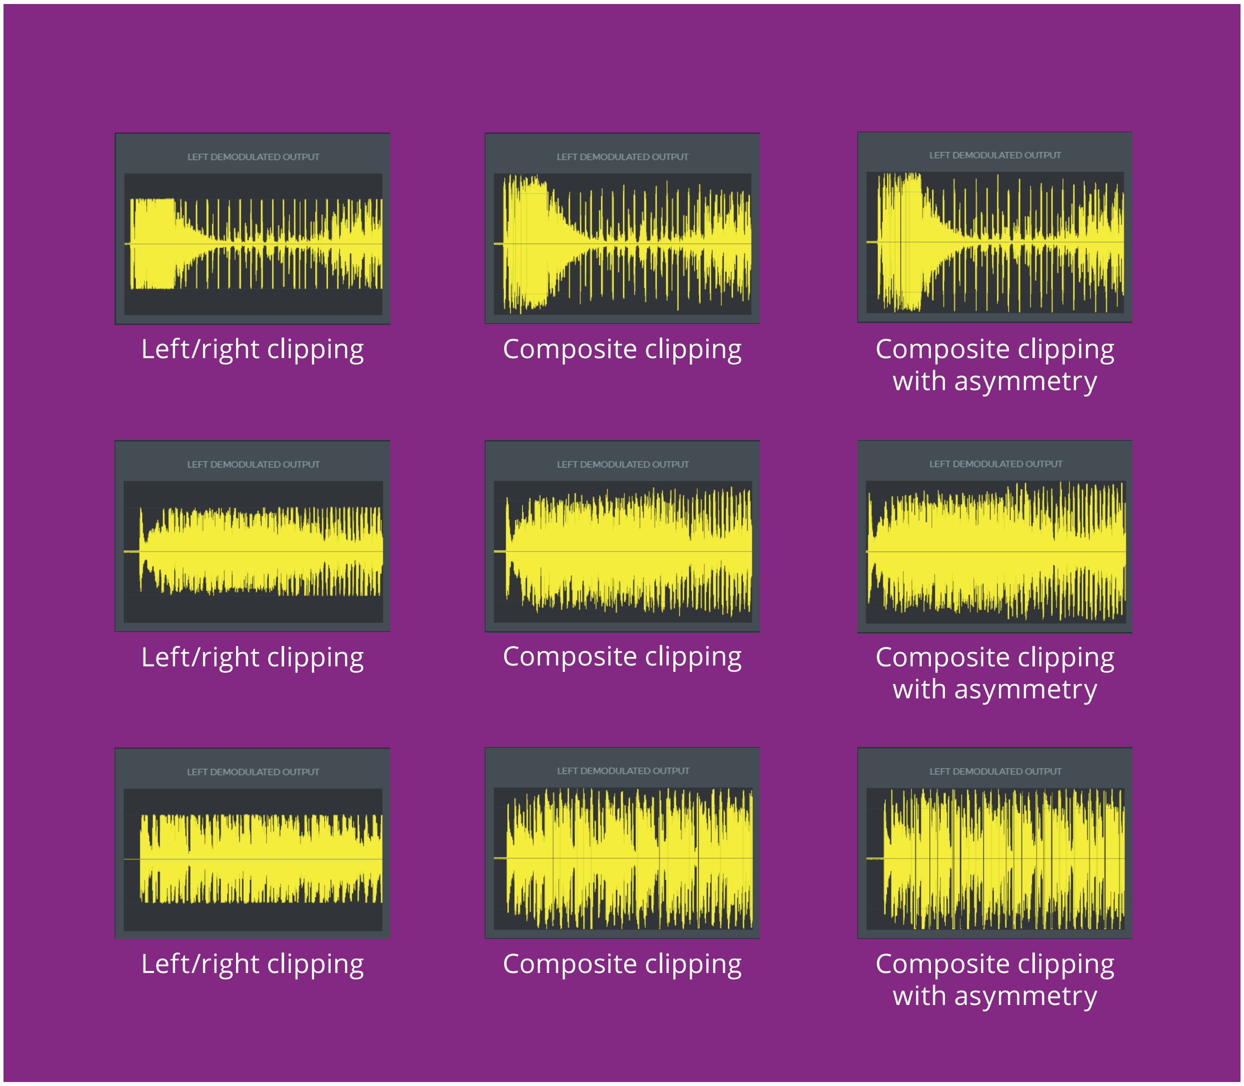

Left/right clipping | Composite clipping | Composite clipping with asymmetry

Left/right clipping | Composite clipping | Composite clipping with asymmetry

Left/right clipping | Composite clipping | Composite clipping with asymmetry

Loudness:

Clipper Drive: Sets the amplification before going into the clipper. Clipper always occurs at 0dB, this slider adjusts the input level to match.

Web/HD/DAB clipper Drive: Allows setting the clipper level for streaming separately.

Final peak remover: Removes small peaks above 0dB in the audio. This is a way to remove any small remaining peaks after clipping. It should always be on.

FM and HD signals

HD signal:

Limit instead of clip: Uses a limiter instead of a clipper for streaming output. Clipping is our preferred method; however, some people prefer limiting. Clipping typically gives you a louder, more punchy sound, whilst limiting if used can lower the level of transients. You can safely use clipping even with lossy codecs as long as you do not drive the clipper too hard which is different than most other processors.

Sound

Even harmonics: Allows even harmonics to be generated for low frequency clipping. Even harmonics typically sound nice and warm and enabling this allows the output to sound louder.

Bass Shape: Deforms the bass to achieve louder levels and sound better on some speakers. Deforms low bass sounds if they are being clipped by making the tops flatter to make them sound louder at same peak level.

Strength: Controls how much deep bass sounds are made louder.

Allow harmonics up to: Controls up to which frequency the bass may be deformed to sound louder.

Artistic effects

Sparkling highs: Enables sparkling highs sound sparking and less clean.

Dirty Bass: Allows more distortion in bass frequencies, caused by bass frequencies.

Dirty Mids: Allows more distortion in the mid frequencies.

Dirty Highs: Allows more distortion in the high frequencies.

Open Sound: By allowing a bit of extra distortion in some situations, the highs sound far more open.

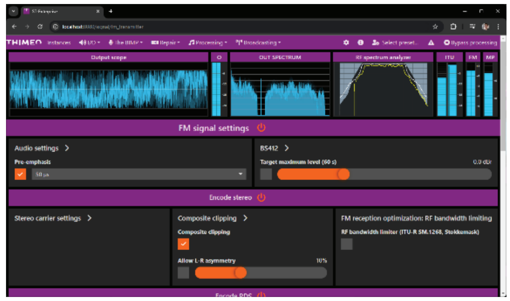

FM SIGNAL SETTINGS:

FM signal settings section

All the settings to generate a compliant FM signal.

All the settings that are related to the FM signal can be found here, such as pre-emphasis, stereo and RDS encoding, RDS texts, Stokkemask (ITU-R SM.1268) and ITU-R BS.412 compliance and multipath distortion protection.

FM processing is more difficult that normal processing. Because of Pre-emphasis you often lose highs when enabling the FM processing settings (Composite Clipping helps a lot against this), Stokkemask and multipath distortion protection can have some impact on the amount of stereo separation, and BS412 lowers the level a lot.

Audio Settings:

Pre-emphasis: The pre-emphasis settings for your region. 50 Us or 75Us is available.

Lowpass filter: FM lowpass frequency. This overrides the main lowpass frequency setting for FM output only.

Mono reception compatibility removes: Extreme phase differences to reduce multipath and improve mono compatibility. If the left and right channel are in anti-phase, the L-R signal is very strong, which makes the RF bandwidth on FM wider, which in turn increases the chance of getting multipath reception issues. This filter makes it possible to determine the maximum phase difference between the left and the right channel.

Center Bass: makes bass mono. Bass frequencies typically have a very small effect on experienced stereo but a big effect on FM reception. So, its usually a good idea to broadcast the bass in mono.

BS412:

Target Maximum level: Determines the maximum BS412 output level. In most countries where BS412 is mandatory, this must be set to 0dB, but in some it is at +3dB.

Stereo carrier settings:

Pilot injection level: The stereo pilot volume, in % of the maximum level.

Composite clipping: Enables the composite clipper.

Allow L-R asymmetry: The maximum amount of asymmetry in the L-R part of the MPX spectrum. To add even more loudness than traditional composite clipping, this allows making the L-R part of the spectrum slightly asymmetrical when that allows louder, less clipped audio to get through.

FM Reception optimization: Enables Stokkemask clipping. This measures and dynamically controls the RF bandwidth which is mandatory in certain countries and will improve reception in fringe areas. It may come at a small cost in stereo separation.

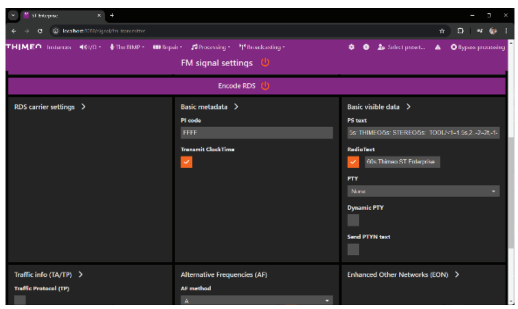

Encode RDS:

We have a full featured RDS encoder inside. In the RDS settings you can enter in all the information you would put into any RDS encoder including PI code, PS text, radio text, PTY, traffic information, AF, and EON data.

For dynamic data, can be sent to the eternal RDS encoder via UECP remote menu. If you click on that menu, you will see options for UECP, or ASCII, via TCP or UDP and you can also specify the port number.

RDS carrier settings:

RDS injection level: The RDS signal volume, in % of the total level. Higher values may improve RDS reception but may also increase multipath distortion.

Use quadrature mode: Encodes in quadrature mode. This allows the total of the stereo pilot and RDS signal from about 13.5 to 12% which allows for about 0.15dB extra loudness when not using the composite clipping. With the composite clipper the difference is about 0.02dB.

Basic meta data:

PI code

The PI code of your radio station.

The PI code controls things like automatic switching to Alternative Frequencies, it is important to set this up correctly. Here is why:

- If you use the same code as another station in your area, you might 'steal' listeners from them (they automatically switch to your station if its reception is better), or you might lose listeners to them in areas where their signal is better. Beside this, this is *extremely* annoying to listeners.

- If you change your code at a later time, some radios where your station has been programmed on a preset button may turn the sound off because they 'think' it is the wrong station, even if reception is perfect. Only reprogramming the preset button helps in that case - but no all listeners will understand this.

PI codes are regulated differently in different countries, please check with your local authorities what you need to do.

The PI code needs to be entered as a hexadecimal number. The default is FFFF, which disables RDS on some receivers so never leave it at that value.

The first 2 digits of the PI code are a country code, but again, many receivers ignore it. Whatever you do, make sure that the last 2 digits are unique in your area, and that they are not FF.

Basic visible data:

PS text

The Program Service text (8 characters)

According to the original RDS specification, the PS text should be static, and its intended use is to display the station name. However, most stations use this text to broadcast lots of things beside the stations name, such as traffic information, phone numbers, even advertisements. Please be aware that certain receivers - mainly in cars - do not display quickly changing texts because it is seen as a danger because it could distract the driver.

STXtreme can broadcast a constant (not changing) text, as is officially mandated by the RDS specification, but it also supports a lot of extra features such as alternating texts, scrolling texts, and dynamically inserted texts such as the name of the current song.

Note that the RDS transmission speed is very low (at most 22 characters per second) so building up the texts can take some time. Because the RDS signal is very weak, bad reception, especially for (moving) car radios, often leads to missing parts, which means that - even when the reception is good - texts need to be transmitted at least two times to be displayed. Many radios even ignore texts (or portions thereof) that are not transmitted at least two times. This lowers the effective transmission speed to 11 characters per second. So, it can take a lot of time before all the texts are displayed properly. The transfer speed is also shared between different types of data **(PS text, RT text', Transmit ClockTime, ...)**. STXtreme' s RDS encoder contains an automatic priority mechanism, which lets texts that have just been changed take more bandwidth.

The descriptions below are valid for both types of texts, although some (such as scrolling) are not useful for RadioText due to the slow transmission speed.

Many radios can only display a limited set of characters: Basically the letters A-Z (capitals/upper case only!), the numbers 0-9 and some special characters such as /, \, -, < and >. This is mainly an issue for PS texts, because they are also displayed on cheap LCD displays on car radios.
So especially for PS texts, only use the characters listed above. (Lower case characters are automatically converted to upper case by all the receivers that we tested).

Constant texts

To transmit a constant text, just type that text in the appropriate field (PS Text or RadioText) in the interface.

The PS Text must be at most 8 characters, the RadioText can be up to 64 characters. Any extra characters are ignored.

Example:
POWER-FM

Alternating texts To show different alternating texts, write all the texts you want to be displayed in the appropriate field in the interface. Split the different parts with a / character.

Example:
POWER-FM/88.5 FM
(Note: The. character will be replaced by a space on radios that cannot display a .)

By default, every part of the text will be displayed for 5 seconds.

Changing the timing of alternating texts To change the amount of time a text is displayed, put the requested time in front of the text, as follows:

- 1.5s: POWER-FM
  Displays the text POWER-FM for 1.5 seconds. When specifying a time in seconds, STXtreme will automatically make sure that the text is transmitted at least twice so radios can receive it properly, even if that takes longer than the specified time.
- 3t:POWER-FM
  Transmits this text 3 times. Stereo Tool will transmit this text 3 times, then switch to the next text. Note that it is possible to use 1t here, but it will lead to bad reception on many radios.

Example:
3s:POWER-FM/1.5s:88.5 FM
Displays POWER-FM for 3 seconds, then 88.5 FM for 1.5 seconds.

Scrolling texts

*Due to the slow transmission times, scrolling is only useful for PS texts. Transmitting a PS text takes at least (2 transmissions * 8 characters / 22 characters per second) = 0.72 seconds. (For RadioText that's 2 * 64 / 22 = 5.8 seconds, which makes scrolling useless.)*

For left scrolling, add a < character at the front of your text. To scroll at higher speeds, add multiple < characters. Right scrolling is also possible, use > for that.

Examples:

```
<THE BEST HITS IN TOWN
<<THE BEST HITS IN TOWN AT HIGH SPEED
3s:POWER-FM/1.5s:88.5 FM/<THE BEST HITS IN TOWN
```

The scrolling speed can also be adjusted using timing settings:

```
3t:<<THE BEST HITS IN TOWN
```

In the previous example, the first and last characters are displayed only very briefly. There are two possibilities to solve this.

Adding spaces:

```
3t:<<      THE BEST HITS IN TOWN (ends here)
```

But it is also possible to specify different timings for different moments of the scrolling. That makes it possible to display the start and end longer than intermediate parts.

Instead of using just 2s: which sets the display time for every part of the scrolling, different values can be set for different positions.

During scrolling, the first position is called '1', the second '2' etc. The last position is also known as '-1', the one before that as '-2', etc.

So what you want is to display positions '1' (start) and '-1' (end) for 1.5 seconds, and everything in between (2 up to -2) should be broadcast 2 times (the minimum for good reception). That can be done as follows: 1=1.5s,2..-2=2t,-1=1.5s:

Examples:

```
1=1.5s,2..-2=2t,-1=1.5s:<THE BEST HITS IN TOWN
3s:POWER-FM/1.5s:88.5 FM/1=1.5s,2..-2=2t,-1=1.5s:<<THE BEST HITS IN TOWN
```

Word wrapping texts

Instead of using scrolling, you can also use word wrapping. That way you can type a longer text, and the words will all be displayed separately.

Word-wrapping is enabled by adding two pipe || characters.

Example:

```
||THIS TEXT IS WORD WRAPPED
```

Again, it's possible to adjust the time that a word is displayed:
1.5s:||THIS TEXT IS WORD WRAPPED

If a word does not fit on the display (for PS texts if it's longer than 8 characters), the remaining characters are truncated. To avoid that, you can combine word wrapping with scrolling.

Combining word-wrapping and scrolling

If you combine word wrapping with scrolling, every word in a text is displayed separately. If it does not fit on the display, scrolling is used to display the whole word.

Example:
||<TRANSMISSION IS A LONG WORD

Now the scrolling problem from before re-appears: The start and end of the word are displayed only very briefly. This can be solved using the timing settings described for scrolling.

So, to display the text, every word 1.5 seconds, but if a word scrolls its start (1) and end (-1) should be displayed 1 second, and the parts in between (2 up to -2) should be transmitted 2 times, that leads to: 0=1.5s,1=1s,2..-2=2t,-1=1s:

Example:
1.5s,1=1s,2..-2=2t,-1=1s:||<TRANSMISSION IS A LONG WORD **Including dynamic texts**
Dynamic texts can be used to include all kinds of changing information in your RDS texts. For example, the name of the song that is currently playing, the current temperature, the weather forecast, you name it.

Special characters

As you can see above, some characters have special meanings in STXtreme. More specifically, the characters <, >, |, :, / and \ cannot be used normally in a text.

If you need to use any of these characters, put a \ in front of it. So: \<, \>, \|, \:, \/ and \\.

Example:
\<POWER!\>

Windows-1252 input

Texts are entered using the standard Windows ASCII table.

Converts special characters, assuming that they are entered using the standard (Windows-1252) character mapping.

If this is disabled, characters are transmitted unchanged - so then it is assumed that the RDS character set is used.

Long PS text: On modern receivers, it is possible to enter a long PS text of up to 32 UTF-8 characters. This supports characters from all languages. Certain modern receivers will display this instead of the standard PS text if both are transmitted.

RadioText

Configures the RadioText text that is displayed on some radios.

RadioText is a longer (up to 64 characters) text. Not all radios that can receive and display RDS also display this text.

RadioText can be configured using the same formatting as PS text. Because it takes a long time to transmit the (much longer) RadioText, up to about 5 seconds for a single transmission of the whole text, options such as scrolling are not very useful, and because the text is so long and reception errors frequently occur, it's better to avoid dynamic texts - unless there's a lot of time for each text to be transmitted (at least 20-30 seconds).

STXtreme also support RadioText Plus (RT+). With RT+, you can tell the receiver what the meaning of parts of the RT text is. For example, you can indicate that a part of the text is the artist name or title of a song, or a link to a website. Some receivers can use this info to - for example - display album art, or even direct links to buy a song. Receivers that do not support this will just ignore the extra information.

An RT text can contain at most 2 RT+ fields. To indicate the start or end of an RT+ field, use \+XX to indicate the start of a field, and \- to indicate the end. XX can either be the decimal value of the RT+ field (00-63), or one of our predefined names (AR for artist, TI for title, WW (case insensitive) for a link. The \+ and \- fields are not part of the length of the RadioText string!

Example:
  Now playing \+tiRadar Love\- by \+arGolden Earring\-
  For more info go to \+WWwww.thimeo.com\-!

PTY

The type of program that you are broadcasting.

Many receivers can automatically search for some type of broadcast, such as a news broadcast. If you don't always broadcast the same type of program and you want to update this setting, make sure that **Dynamic PTY** is enabled.

The list of PTY codes is very different for RDS (EBU, used in Europe) and RDBS which is used in the US. The PTY codes are changed when you change the **Pre-emphasis**.

Dynamic PTY

The PTY changes regularly.

Send PTYN text

Text field to specify the PTY in more detail.

This can be used to transmit an 8-character PTYN text that will be displayed on receivers, usually instead of the PTY code. For example, the PTY code, which needs to be selected from a (limited and sometimes outdated) list of 31 entries, might be SPORT. And the PTYN text could then be FOOTBALL.

Receivers will still filter based on the **PTY** value.

Traffic info (TA/TP)

Traffic Protocol (TP)

Traffic Protocol: This radio station broadcasts Traffic Announcements ***(Traffic Announcement (TA))****.*

Traffic Announcement (TA)

Traffic Announcement: Currently traffic information is being broadcast.

Receivers will switch to this station, even overriding other audio sources such as CD's, USB sticks, memory cards and cassette tapes, and adjust the audio volume. Make sure to turn it off when the announcement is finished!

Turn off TA after

If for some reason TA is not turned off, this will turn it off automatically after some time.

Alternative frequencies (AF)

AF lists are used to tell a receiver which other frequencies carry the same content as the current frequency. Receivers can use this to automatically choose the frequency with the strongest signal. It can also tell receivers (together with some other data) which other frequencies usually carry the same content, but not always - and protect against switching to a different version of the same programming during times when the broadcasts are different. For example, a station might have multiple transmitters with regional ads or regional news.

AF method

Selection of the AF method: A, B or B for retransmission on multiple frequencies.

AF method A is simple: You can select up to 25 frequencies that carry the exact same content. A receiver can always switch between these frequencies.

AF method B makes it possible to specify which frequencies carry the same content, and which frequencies carry localized (sometimes different) versions of the same content. A receiver should not switch (depending on settings that the user can sometimes overrule) to a frequency that carries a localized version of the same station, at moments when the content is different from the frequency that the user is listening on.

AF method B for retransmission is a special case: If the content of a station is rebroadcast elsewhere where there is no separate RDS encoder available, then the first station needs to transmit the AF method B lists for multiple stations.

The settings below depend on the A/B setting selection. For method B it is possible to enter multiple lists if you open the advanced settings.

RDS Alt Freq 0-25 (AF method A)

List of up to 25 alternative frequencies.

Some radios use this list to check which other frequencies to scan for a stronger signal. However, most radios seem to rely solely on the **PI code.**

Main (AF method B)

The AF method B tuning frequency.

For AF method B encoding, the RDS encoder needs to know the frequency that the signal will be transmitted on. So, fill in your own frequency here.

Same (AF method B)

The list of AF method B frequencies that carry the same content as ***Tuning frequency****.*
Example: 90.1 91.5 93.2

Regional (AF method B)

The list of AF method B frequencies that carry a regional variant of ***Tuning frequency****.*
Example: 95.5 96.6

Enhanced Other Networks (EON)

Most of the settings here are self explanatory. The following two need an explanation:

AF

List of EON AF frequencies.

Several EON AF methods can be used and combined here.

Method 4 is used for AF method A data and looks like this: “4 91.1 92.2 93.3”. The number of frequencies can be anything up to 25.

Methods 5, 6, 7 and 8 are used for AF method B pairs, and are entered as the method number plus 2 frequencies in the correct AF method B order. So, if you want to transmit 2 pairs using method 5 and 6, you might type “5 91.1 92.2 6 94.4 93.3”.

14B burst

Enables 14B burst transmissions when the **TA** value changes.

Group Sequence

Enable Group Sequence

Enables the RDS Group Sequence.

The Group Sequence can be used to specify exactly which RDS groups must be transmitted in what order, and what should be transmitted if a specific group is not available.

Group sequence

The group sequence to be transmitted.

Example: 0A 2A 0A 8A

This will transmit a lot of 0A (PS text) groups, and only a small number of 2A (RT text) groups. 8A is normally used for TMC (traffic information for navigation systems), but it is only available if 8A group data is sent. See also **Extended group sequence**.

Extended group sequence

Extended Group Sequence information.

Example: 8A: 3A 1A

If in **Group sequence**, group 8A is encountered but no 8A data is available, then this list means that it needs to try to transmit a 3A group. If there is also no 3A data, then it should transmit a 1A group.

Multiple such fall-through lists can be separated with a /. (Example: 8A: 3A / 9A: 1A).

RDS remote control and monitoring:

UECP (Remote)

UECP settings:

UECP is a standardized protocol that is used to send information to RDS encoders. It is supported by many software and hardware RDS encoders, and by many playout systems.

Originally, UECP was used over an RS232 connection, but many products are now supporting it via TCP/IP. STXtreme only supports the latter.

Enable UECP

Enables remote RDS control via UECP.

If this is enabled but no UECP server is connected, then the behavior depends on **Use standard RDS settings when connection closed.**

Packets

Packet counter (display only).

This can be used to verify that UECP data is coming in and passing through the Site/Encoder filtering.

Mode

Switches between UECP (recommended!) and ASCII mode control.

UECP is a standard created by the same group that created the RDS standard, and we recommend using that when available.

However, some playout systems and other programs don't support UECP mode, and they support different dialects of an ASCII mode.

ASCII mode has been tested with Arctic Palm Center Stage Live RDS. At the moment it does not work with Magic RDS (we will fix that).

If you want to use ASCII mode, here are some example commands:
Set a static 8-character PS text:
PS=HELLO
PS HELLO
PS:HELLO

Set a dynamic PS text, automatically split into sub-texts if longer than 8 characters:
DPS=Hello there, this is a long text!
Alternative names: PS_TEXT, DYNPS, DPS1.

Set a dynamic PS text, automatically split into sub-texts if longer than 8 characters:
DPS=Hello there, this is a long text!
Alternative names: PS_TEXT, DYNPS, DPS1.

Set a Thimeo-style PS text, with instructions for scrolling, timing etc.:
THIMEOPS=1s:Hello/2t<This text scrolls

Set an RT text:
RT=This is my radio station
Alternative names: TEXT, RT_TEXT, DYNRT, RT1, NEWS.

Send RT+ data:
RT+=a,b,c,d,e,f
Alternative names: RTP, RTPLUS.

Delete RT+ data:
RTPRUN

Other commands:
PI=1234
PTY=11
DI=0
MS=1

TCP/UDP

Listen to UECP in TCP or UDP mode.

Port

The UECP HTTP port address.

Use standard RDS settings when connection closed

Determines behavior when the connection to the UECP server gets lost.

If this setting is enabled, whenever the connection is lost, the RDS encoder will return to the settings that are configured inside STXtreme. If it is disabled, it will keep using the last received UECP data.

If you are using UECP to update things like Now Playing information, you should probably enable this and configure some standard text inside STXtreme.

Block access to text fields:

This is a security feature. Some stations need external data, for example for traffic information. Enabling this setting blocks access to all RDS text fields.

Specific address

Listen only to UECP commands targeted to this address.

UECP can send commands via multicast to multiple RDS encoders. If it does that, it can use addressing information to target a subset of encoders.

Specific address

UECP "Specific Address".

See UECP specification and UECP encoder settings.

Allow traffic for site

UECP "Site Address".

See UECP specification and UECP encoder settings.

Allow traffic for Encoder Address

UECP "Encoder Address".

See UECP specification and UECP encoder settings.

DSN

UECP "DSN".

PSN

UECP "PSN".

Monitoring section:

Overview of transmitted RDS groups. You can use this to monitor how many packets of each RDS group type have been sent.

Encode RDS2:

RDS2 logo file: You can upload logo PNG or JPEG file. This will be transmitted in the RDS2 data. Try to make the logo file as small as possible in order to transmit it quickly. RDS2 logos are currently only supported by a limited number of receivers.

AM Signal settings:

Lowpass filter: Lowpass filter for AM

NRSC: Enables NRSC AM pre-emphasis.

AM Asymmetry strength: How asymmetrical the sound may be. Typically, this is 125%.

AM downmix to mono: For mono AM transmissions, this converts the audio to mono without any phase cancellation problems.

AM stereo (C-QUAM): Enable this for C-QUAM stereo transmissions.

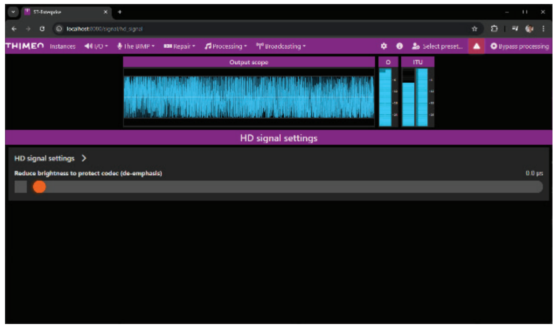

HD Signal settings:

Reduce brightness (de-emphasis): Because FM reduces brightness due to pre-emphasis, presets are often brighter on non-FM outputs. This setting allows you to reduce that brightness for the HD output.

Pre-emphasized clipping /limiting to reduce codec overshoots: This reduces overshoots in high frequency content caused by lossy audio codecs.

Optimize audio for lossy compression (clean): This puts the clipper in a mode that does not cause harmonic distortion even if this comes at a small cost in loudness. This makes it possible to use this clipper with lossy codecs without increasing codec artifacts.

Reduce parts of stereo to help low bitrate codecs: For very low bitrate streams, i.e. DAB or HD, this reduces stereo content in frequencies that our ears are less sensitive to in order to simplify the signal that goes into the codec to reduce codec artifacts.

Watermarking

Nielsen PPM watermarking in Thimeo products

Getting started.

Support for Nielsen PPM watermarking is currently limited to the United States and Denmark, and it is only available in the Windows 64-bit version of ST-Enterprise. To use it, you first need to obtain a CSID from Nielsen by contacting your Nielsen Account Representative. Once you have that, please send an email to support@thimeo.com to confirm to us that you have such a license, and we will send you an unlock code to use it in our software.

Contact Nielsen Encoder Support at encoders@nielsen.com or call 800-537-4872, option 2 if any assistance is required contacting your Nielsen Account Representative.

Kantar watermarking is available in ST-Enterprise and STXtreme:

Kantar embedding currently only works if the **Sample rate** is set to 48 kHz or a multiple thereof (96, 192), and **Don’t process high frequencies above lowpass filter** is disabled.

To get a watermarking embedding license, please contact Kantar Media support at www.kantarmedia.com/watermarkinghelpdesk with following information:

- Product name and version
- Customer name
- Country
- If different, country of broadcast
- Channel(s) to be watermarked.
- Customer internal name for the hardware platform
- Authorization Code for each hardware or login contact for online solution.

The AutorisationCode.txt file can be found in the Stereo Tool directory after installation and can be regenerated if needed by running AuthorizationCodeCL.exe in that same directory.

Kantar watermarking log files are created in C:\Users\username\ST-Enterprise.log. If you rename the ST-Enterprise executable, ST-Enterprise is replaced by that name, so for example for ST-Enterprise10.exe it will be in C:\Users\\ST-Enterprise10.log.

Running Kantar watermarking causes an extra delay in the audio. The total delay is equal to the setting under **Processing latency** plus the processing delay plus the Kantar delay, which is 2048 samples. So, the **extra** delay when enabling Kantar at 48 kHz, compared to the set value, is 128.00 ms for Latency 4096, 85.33 for 2048, 64.00 for 1024, 53.33 for 512, 48.00 for 256, 45.33 for 128. Note that turning Kantar on or off causes a jump in time of 2048 samples - when turning it on, it inserts a bit of silence, when turning it off, some audio is discarded.

Kantar watermarking resynchronizes its time stamp with the Windows system clock at least once a day (when a new day starts), plus whenever the timestamp has an offset to the system clock that exceeds 10 seconds.

Kantar watermarking for HD panel

- **FM Channel**

  The Kantar Channel Name for FM output.

  A license file can contain licenses for multiple channels, and you can select different channel names for different signals, so make sure to select the channel that you want to encode for.

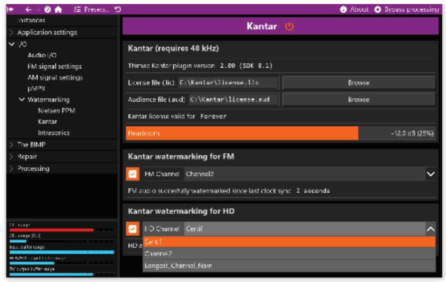

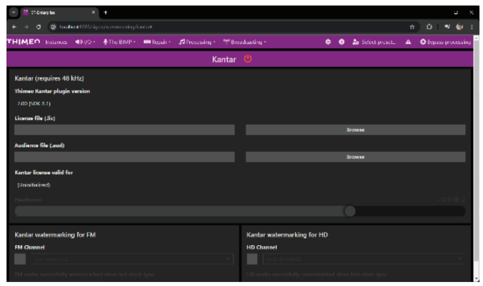

- **HD Channel**
  The Kantar Channel Name for HD output.

  A license file can contain licenses for multiple channels, and you can select different channel names for different signals, so make sure to select the channel that you want to encode for.

- **FM audio successfully watermarked since last clock sync**
  Shows how many seconds the watermarking has worked since the last clock time reset.

  This is useful to verify that watermarking is active, and to see when the watermarking clock time has been updated. If the time does not move while this watermark is enabled, watermarking is not working - you should have received an error message at an earlier point if that is the case.

  Always verify that the signal that you think you are watermarking is indeed running.

- **HD audio successfully watermarked since last clock sync**
  Shows how many seconds the watermarking has worked since the last clock time reset.

  This is useful to verify that watermarking is active, and to see when the watermarking clock time has been updated. If the time does not move while this watermark is enabled, watermarking is not working - you should have received an error message at an earlier point if that is the case.

  Always verify that the signal that you think you are watermarking is indeed running. For example, if you are feeding a left/right demodulated FM signal to a streaming encoder, it may sound ok, but it would have the FM watermark. In that case, no HD signal would be watermarked at all. So, always verify that the output(s) that you expect to be watermarking are all enabled, and no others.