

PRO AUDIO DSP

Dynamic Spectrum Mapper

Plug-in Manual
for
Digidesign ProTools, Audio Units and VST.



1. Introduction.

The Dynamic Spectrum Mapper (DSM) application is intended to provide multi-dimensional control over both the spectral response and dynamic characteristics of audio programme, in order to bring a whole new dimension of facility and artistic ability to the sound engineer. The use of novel processing and programme capture algorithms provides the ability to achieve exemplary results in common tasks such as programme compression and enhancement, vocal expression and instrument characterisation very quickly. The ability to capture, modify and re-apply the spectral characteristics of programme is a powerful creative tool in all stages of the production cycle, from recording and mixing to mastering stages.

2. Principle of operation.

The basic architecture of the plug-in is similar in concept to a multi-band compressor, except that the method of band splitting produces far greater frequency resolution in the higher mid and high frequency ranges. This high frequency spectrum resolution allows a high degree of functionality and artistic modification of the programme character. Since the process produces different artefacts to conventional band splitting dynamics applications, a new range of programme character can be achieved, such as tape saturation effects and the softening of high frequency content.

The signal is split into multiple bands using an FFT based process. An adaptive dynamic gain control process operates on the effective power represented within each band, before recombination back into the output signal domain. Manipulation of the threshold curve changes the threshold target level for each band. This allows the frequency content to be modified based on dynamic changes in the signal over both time and rate. Because the dynamic gain control process includes frequency domain information captured from the original programme (as well as user-defined target responses), it provides a large degree of novel artistic functionality, which includes:

Enhance the perceived volume and presence of mixes.

Compress even the most difficult programme with increased accuracy and musicality.

Characterise and personalise output in completely new ways.

Copy and equalise sonic impressions between dissimilar sounding tracks.

Process, enhance and 'tame' vocal tracks with unprecedented accuracy and speed.

Dramatically enhance and increase the intelligibility of dialogue for media post production.

Model and characterize instrument sounds to achieve highly creative artistic impressions.

Create unique dynamic effects in both level and frequency domains to produce exciting new artistic effects.

3. System Requirements.

Pro Tools

- Macintosh running OSX 10.4 (please refer to www.proaudidsp.com for latest operating system compatibility updates)
- Pro Tools LE or HD system.
- iLok USB Key

Audio Units

- Audio Units compatible application.
- Macintosh running OXS 10.4 or later.
- ILoc USB key


VST

- VST compatible application.
- PC running XP SP1 or later and Vista.
- Macintosh running OXS 10.4 or later.
- ILoc USB key

4. GUI Controls.

The GUI controls are set out in four general sections for operational clarity. These are described below.

4.1. Graph and selectors



The graphical display area plots level against frequency and contains two curves.

The **Red** curve displays the **Active Spectrum** of the input signal at all times and is live during programme.

The **Yellow** curve displays the **Threshold Spectrum** of the captured (or manually modified) spectrum. This sets the threshold of compression in each band.

The Yellow active threshold spectrum is obtained by operating the **Capture** button. The button can be operated momentarily to capture a short transient spectrum, or if held down it will average the spectrum over time in order to capture a longer-term character spectrum of the programme.

Note. In order to accurately match the input signal, the Capture function acquires both the level and frequency spectrum of the programme. Therefore the position of the yellow threshold curve on the graph display (and the resulting compression thresholds) will initially depend entirely on the level of the input programme. Quieter and lower level programme will result in lower captured compression threshold levels.

Note. The active threshold spectrum is saved along with all other parameters in the Pro Tools saved settings. Therefore captured responses can be saved and re-used on other tracks and projects.

Compression and gain change can occur anywhere on the graph where the **Red active spectrum curve** is higher than the **Yellow threshold curve**. Where the Red curve remains beneath the Yellow curve the signal at those frequencies will remain unchanged.

The **Link** button ties left and right channels together for stereo operation. The curve is coloured red to indicate stereo linking. If **Link** is deselected, left and right channels will operate independently. During independent operation the red curve is replaced by two blue shaded curves representing the active spectrums of both left and right channels.

The **Freeze Gains** button stops dynamic activity and freezes the response to that which was active at the moment of operation. This arrests any further gain compression and permanently imposes the response of the active spectrum onto the programme – somewhat like a static equaliser.

The **Limit** button selects extra processing, which suppresses sample value overloads. It achieves this by reducing the peak levels of the signal without significant loss of loudness. It also prevents hard clipping and (depending on the programme material) allows higher levels to be developed within the DSM section, which can significantly enhance the loudness and presence of the programme.

The **16 Bit Dither** button toggles from 24bit to 16bit target media when selected

The Process **In** selector is provided for 'click-less' in/out comparison purposes.

4.2. Dynamics control.



Familiar compressor style controls are provided to modify the activity of the dynamic gain control of the programme.

The **Threshold** control changes the **level** of the Active Threshold curve in relation to captured spectrum. The yellow active threshold curve position changes vertically on the graph to reflect this. When the **Threshold** control is set to 0dB the active thresholds correspond to the active levels of the original captured (or manually modified) spectrum. If the Threshold control is adjusted to lower (or higher) levels, the active thresholds in all bands will vary accordingly. An extended positive threshold range beyond 0dB is provided so that Threshold Spectrums captured at lower programme levels can be raised to useable levels.


The **Ratio** controls the ratio of compression for all frequency bands.

The **Gain** control is provided as an independent manual gain make-up function to compensate for level reduction due to compression.

The **Attack** and **Decay** controls modify the rate of gain reduction and recovery common to all frequency bands.

The **Knee** control controls a 'soft' compression ratio function and is expressed by the number of decibels below the active threshold that the soft knee will start pre-compressing, in advance of the active threshold being reached.

4.3. Timing profile.



Normally all time constant parameters within the dynamics section are common to all bands in the spectrum. The **Timing Profile** controls provide a method to modify timing parameters in relation to frequency range, in order to produce more sympathetic action depending on programme type or to produce sound effects.

The **LF ATT** (low frequency attack) progressively **increases the attack time** for lower frequency signals in relation to higher frequency signals. This slows down the attack for low frequency compression events within the spectrum.

The **HF REL** (high frequency release) progressively **decreases the release time** for high frequency signals in relation to lower frequencies. This speeds up the gain recovery for high frequency compression events.

4.4. Parametric Threshold.



The **Parametric Threshold** section allows the user to either modify a captured spectrum manually, or to build a suitable custom spectrum from scratch. It achieves this by changing the active threshold curve in the frequency and level domains.

The familiar EQ-style Frequency Q and Threshold controls impress the appropriate response curves onto the active threshold curve in order to modify the resulting action of the compression.

Note. *The controls neither equalise the signal itself, nor act as side chain EQ.*

The response modifications produced by the parametric threshold control are added non-destructively to the existing active threshold curve and parameters are stored separately, so that a set up can be restored and modified further without loss of the originally captured spectrum.

In buttons are provided for comparison purposes.

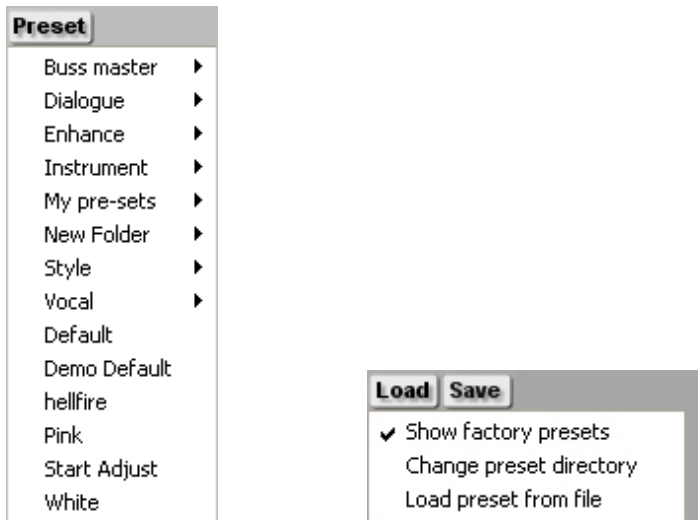
Note. *Manual changes in the response of the parametric threshold will take time to occur in the resulting programme.* This is because the changes applied to the compression sections will only take effect at a rate depending on the time constant settings in the dynamics control section.

4.5. Options menu.



Clicking the **Options Menu** produces additional pop down menus that allow preferences to be set for rotary knob modes and meter over indicator duration times.

4.5.1. Load Menu (VST versions)



For convenient pre-set management, the VST version of the DSM includes a pop down **preset selection menu** and a **load menu** that allows loading and saving of pre-sets to and from specific files and directories.

Preset. Displays a menu of available presets in the current preset directory. Selecting the preset in the menu will load the preset directly into the plug-in.

Load displays a menu with the following options:

Show factory presets. Changes the Preset menu to display the factory presets.

Change preset directory. Changes the current directory from where the Preset menu looks for presets to display.

Load preset from file. Allows the user to load a preset from a different location to the Preset menu, it means a preset can be loaded without changing the current preset directory.

Save. This saves the current plugin settings as a new preset. The location defaults to the current preset directory but can be changed to somewhere else if required.

5. Using the DSM.

With an understanding of the general architecture of the DSM, functionality should be quite straightforward and there are many ways to achieve useful results. The following basic procedures are suggested as operational routines to deal with common situations and illustrate functionality.

5.1. Setting the Capture response.

The DSM differs from conventional dynamic compressors in that it makes use of the original spectral character of the programme. Therefore it is important to establish the threshold spectrum initially.

The best way to hear the capture is to set a gain limiting situation by setting the **Threshold control** to -24dB , the **Gain** to $+24\text{dB}$, **Ratio** to 100:1, put **Attack** to around 10mS and **Release** to 30secs, leaving all other controls at minimum. This will cause the compressor to gain limit the signal to the captured response.

The **'start'** set-up is included for convenience and provides these settings.

Operating the **Capture** function will then sample the spectral parameters of the input programme. It is often best to do the capture at a place in the programme where it is fully under way – most often this equates to the loudest sections.

At this point the programme will be effectively gain limited to the response of the captured spectrum. This can be confirmed by toggling the **IN** button to turn the processing on and off, where the processed sound should closely resemble the original programme over the passage where the capture was taken.

You now have a compression that is removing all dynamic range for signal content above the **Threshold Spectrum** and forcing it to take on the spectrum and level of your captured response.

This is a very good starting point that we will return to throughout operational description.

5.2. Setting Dynamic parameters for compression.

The **Threshold control** can be used to change the active level of the captured threshold spectrum to compress to a higher or lower output target level, or place the threshold spectrum at higher levels in order to impose the captured compressor response and dynamic control only on to louder parts of the programme. For reference, changes to the level of the active threshold spectrum are reflected in the yellow plot in the graph.

The **Ratio control** can be used in a conventional way to restore the required amount of dynamic range for signal within the compression range and the **Knee control** can be used to soften the transition in and out of the compressed level ranges. The **Gain control** can be used to compensate for level changes caused by threshold and ratio adjustments.

The timing controls act in the same way a conventional compressor. Increasing the **Attack time** will slow response to peak events, reducing the **Decay time** will speed up gain recovery after compression events.

The **LF ATT** can be employed to slow down compression of low frequency events to allow greater perceived LF energy in the sound.

The **HF REL** may be used to speed up the recovery of high frequencies to produce the effect of increasing overall HF energy, 'air' and 'fill' to the sound.

Note, since the capture function employs the same processing architecture as the signal control, *compression gain variance due to timing control changes can be compensated for by re-capturing the spectrum with the new settings active*. This is particularly useful when setting timing controls for effects. Very wide time constant settings can be accommodated and experimented with in this way.

Due to the large numbers of effective bands being compressed and the included 'weighting' from the original spectrum of the programme, a very sympathetic and credible compression can be obtained over a very wide dynamic range. In particular, percussive events at higher frequencies modify lower frequency ranges much less than in conventional compressors and the existence of the original spectrum maintains a credible sonic signature even during large amounts of gain compression.

5.3. Character modification.

The frequency spectrum character of a source track or sound can be applied to another simply by capturing the source and playing the target track through the plug-in.

Going back to our starting limiting set up in [4.1], operating a capture on the source track will force the output to follow the captured response. If the target track is played through this (without re-capturing) it will take on many of the spectral characteristics of the source sound.

The dynamic controls may be used like a conventional compressor to achieve the desired dynamic range characteristics

If dynamic compression activity is unwanted, the **Freeze Gains** function can be used to suspend compression and the plug-in will operate like an EQ that has the source character pre-loaded. Slower attack and release settings are best for this function, as the slower response time will facilitate a more representative frequency response at the moment the gains are frozen.

Novel **Spectrum Thresholds** can be constructed by modifying an existing one or building completely from scratch using the **Parametric Threshold** section. Set-ups with pink and white frequency spectrums are provided as useful starting points if building spectrums from scratch.

5.4. De-essing.

This is a process of simply capturing the vocal response during a period where 'essing' does not occur and re-applying it to the vocal track at the appropriate threshold levels. Some included set-ups illustrate this.

Going back to the original limiting situation we had set up in [4.1], setting both the **Threshold** and **Gain** to 0dB will ensure that any capture will be calibrated to unity at the incoming signal level.

If a capture is performed during a period where the singer is not 'essing' and this is applied to the whole track, louder essing events will breach the **Threshold Spectrum** within the frequencies they occur and those parts of the spectrum will be forced back to the captured response.

Since the captured response is a valid spectrum of the singer's original voice, the quality of the de-essing can be very high.

If desired, the **Ratio** and **Knee** functions can be adjusted to achieve a progressive application of the target response and the **Threshold** can be adjusted so that the effect commences at the appropriate levels.

Alternatively (or in combination with the above) one can visually identify essing events in the graph area and a custom threshold spectrum can be constructed using the **Parametric Threshold** section to build an appropriate response with a 'dip' in the offending region. This can then be applied at the appropriate levels to reduce these events.

Faster **Attack** and **Release** times can often be used effectively because the action is restricted only to high frequency events.

Slower **Attack** times can be effective at retaining the sibilance of the voice whilst still suppressing unwanted longer term 'essing'.

5.5. Vocal enhancement.

Vocal enhancement and intelligibility improvement are conceptually similar to the above; except that the character and dynamics of the voice are deliberately modified under a range of levels, programme and timing conditions to achieve the desired result.

The whole voice may be compressed to a captured spectrum of itself, to achieve maximum consistency and penetration whilst retaining vocal character. Or it may be compressed with a capture of some other vocal parts or tracks in order to achieve a higher degree of consistency between takes or to emulate aspects of other tracks.

It can be made to ‘morph’ from one character into another by either applying a previously captured spectrum or creating a custom **Threshold Spectrum** in the **Parametric Threshold** section, setting the threshold at some intermediate level and applying the appropriate compression ratio and timing parameters.

Setting lower ratio values will cause the sound to take on the character of the **Threshold Spectrum** progressively as the voice gets louder.

By manipulating the **Attack** and **Release** controls it is possible to cause the voice to blend towards the **Threshold Spectrum** over a period of time without crushing the sibilance. This can be useful for preferentially moderating sustained vocal phrases, particularly when they often have a tendency to become harsh in many vocal styles.

Because the compression is aiming for a credible spectrum of the original voice, explosives and other unwanted vocal sounds are effectively tamed without a damaging loss of character to the overall vocal performance.

5.6. Loudness and presence enhancement.

Using the **Limit** function in combination with the **DSM** level control and compression can result in significantly greater loudness and presence for both single tracks and entire mixed programme. Because the **Limit** function allows higher modulation levels without clipping, it is possible to pass signals that would otherwise cause sample value overloads.

This can be achieved after setting the appropriate **Threshold Spectrum**, **Attack** and **Release**, by increasing the overall threshold or gain a few dB in order to progressively overdrive the **Limit** processing. Tight control of the levels by the DSM can avoid gross distortion from the limiting section.

Whilst it is possible to use a **Threshold Spectrum** captured from the programme itself, it is also possible to achieve interesting results by producing a modified threshold spectrum that delivers greater gain in the mid ranges (where we are most sensitive) and allowing the DSM to impose this response onto the louder or percussive parts of the programme.

To achieve this, higher overall threshold settings can be used to apply a modified (artificial) threshold spectrum only to the top few dB of the dynamic range.

Set-ups named as ‘louder’ are included as an illustration of this concept. The starting point for most of these was the ‘pink’ set up, from which the spectrums were further modified using the **Parametric Threshold** section.

Note. For percussive programme, greater overall punch may be achieved by setting slower attack times and allowing the limiter to handle the high level transients.

Note. For maximum loudness, Knee function settings should be minimised to avoid premature compression.

5.7. Sound effects.

The novel architecture of the plug-in produces a different set of artefacts from conventional compression and this can be used for a range of sound effects. Because the DSM can capture and aim for specific responses, novel effects may be generated if used in conjunction to other processes such as equalisation. Some examples are provided below.

5.7.1. High frequency content softening.

Because the compression occurs within many bands simultaneously, very fast attack and release times can be used over relatively small level ranges, particularly in the higher frequency regions.

By applying a suitable threshold spectrum that preferentially compresses the HF content, a useful softening of the high frequency content can be achieved which is reminiscent of tape saturation effects. Some included set-ups illustrate this effect.

5.7.2. High frequency content enhancement.

By applying significant compression at lower thresholds and ratio ranges, then preferentially speeding up the recovery of the HF compression using the **HF REL** control, a pleasing increase in HF detail and presence can be produced, which is reminiscent of a style processing known as ‘aural enhancement’.

Unlike the application of simple equalisation, because the HF levels remain under control by the compression, no extra harshness is produced.

This can be particularly effective when applied to acoustic stringed instruments, woodwinds and flutes to accentuate the high frequency overtones.

5.7.3. Low frequency enhancements.

Increasing the **LF ATT** control will preferentially slow down the attack speed for low frequency programme content. This is useful as a way of increasing the perceived warmth of the sound and preventing short-term percussive resonance to be cut short excessively.

By producing a custom **Threshold Spectrum** that rolls off in the lower frequencies and placing it at a suitably high level, it is possible to cause the LF to be moderated at high levels. This is useful for reducing LF over loading when aiming for the maximum possible loudness in programme with significant dynamic range.

By applying Equalisation prior to the **DSM**, it is possible to pre-boost certain parts of the programme and allow the DSM to re-normalise it at higher levels, to effectively ‘undo’ the equalisation. Effects such as low-level sibilance enhancement, low level warming and other dynamic frequency effects can be produced this way.

6. Additional information.

6.1. Peak levels and Limit function.

Because the DSM modifies the frequency response, and therefore changes the waveform shape of complex signals, it is possible that larger peak levels will result. As in the case of simple equalisation, this is true even when reduction is applied at certain frequencies. Therefore if maximum peak modulation is required, the **Limit** function should be employed to suppress sample value overloads in the output programme.

From V1.3 onwards the limiter acts only on the top levels above –1dBfs, leaving programme below that unchanged. So under most normal conditions the limiter may be left in place for safety, without adversely affecting the sound of the programme.

From version V1.3 onwards the input and output meters display levels up to +6dBfs above notional maximum level. The output level meter may therefore be used to monitor level **before the Limiter section**, where the yellow section corresponds to the amount of limiting being applied.

6.2. Relative frequency distribution.

The DSM uses FFT methods to produce the active frequency bands. This means that the density of effective bands naturally increases towards the higher frequency ranges, producing the greatest frequency resolving ability at high frequencies.

Whilst this provides a high degree of control in the higher mid and HF frequency ranges most identified with overall programme character, the natural reduction of active bands in the LF region may make it unsuitable for some classes of sound effect that focus primarily on low frequencies within the programme.

6.3. Latency.

Total processing delay Pro Tools.

44.1KHz & 48Khz:	48 samples
88.2KHz & 96KHz:	96 samples
176.4KHz & 192KHz:	192 samples

Total processing delay Audio Units.

44.1KHz & 48Khz:	52 samples
88.2KHz & 96KHz:	104 samples
176.4KHz & 192KHz:	208 samples