



1- The spectral range is divided into N equal filter bands. If this is equal to FFT-size then the bandwidth of each filter is 1 FFT bin. If Band is FFT/2 then the bandwidth of each filter is 2 FFT bins. And so on...

2- This defines the lower and upper dynamic thresholds. All filter Bands whose amplitudes (ranging between 0.0 and 1.0) fall below the lower threshold are multiplied by the controller “low” (see number 3). All that fall in-between are multiplied by “mid” and those above are multiplied by “high”.

Thus effectively the spectrum is divided into three layers depending on the energy or amplitude contained in the spectral data. And the amplitude of each layer can be independently controlled.

This is the same process that the GRM contrast plug-in applies to the sound.

4- This is very important and to a large degree determines the outcome of this instrument. In short, the spectral data of the consecutive analysis windows contained within N seconds

of sound (in the above picture 0.33 seconds) are averaged, prior to going through the threshold analysis. Note that this averaged spectrum is not itself synthesized but passed unto the threshold analysis. In effect this eliminates very fast changes in the spectrum of the individual filter bands, thus smoothing any erratic jumps in the amplitude analysis result. Ultimately this helps minimize the bubbly fft artifacts that normally arise as the result of window-to-window changes in the spectral data (often associated with noise rather than musically desirable sounds). With higher average values and a fine tuning of the threshold values we can achieve a rather good noise reduction method: setting the "low" slider to 0 and the other two sliders to 1, so that only the very constant (due to high average setting above 0.5) low amplitude spectral elements are eradicated (such as low level hiss).

Note however that the larger the average value the more CPU is required. So it is wise to avoid very large values (more than 0.5) in real-time performance. This is not a problem with bouncing to disk (it will only elongate the processing time).

5- This parameter will further smooth the averaged analysis data using a spectral smoothing algorithm. This is not as important a controller as number 4 but it does improve the quality when set between 0.1 and 0.5. Any lower than that and the analysis data is more or less blurred in terms of morphology. 0 indicates complete freeze of the data.

6- This parameter defines the amount/depth of multiplication applied to the filter-band's amplitudes. 1 denotes a full processing (wet signal) and 0 is equivalent to dry signal only. For a subtle effects (such as noise-reduction) a little bit of the dry signal added is a good idea.

Further notes

Also worth noting is that since this instrument does not alter the frequency data of the input sound there is no loss of phase information. So the blurry fft effects can be completely avoided here, as long as the FFT, window and overlap sizes are tuned properly in relation to the types of morphologies contained in the sound-files being processed.

Batch processing

I have used this instrument for noise reduction before. If one has for instance 40 sound-files of the same source recorded at the same level (such as a long sound-file chopped into smaller chunks) then it would make sense to set the parameters for one sound only and batch process an entire folder with all 40 sounds contained in it.

This is possible with the non-real-time bounce to disk option (the red button). Instead of dropping a file drop a folder in the drag-and-drop area. Now pressing the red button will batch process all the files contained within that folder and automatically rename them (keeping the original files).