



RADIOWORLD

The News Source for Radio Managers and Engineers

WWW.RADIOWORLD.COM

Demystifying the Vorsis 31-Band Limiter

The Subtle Audio Details Being Revealed By The Limiter Have Always Been There

COMMENTARY

BY JEFF KEITH

In 2007 Wheatstone Corp. created our Vorsis development team and asked us to create new audio processing technology that could bridge the gap between loud and clean. One of our team's inventions was the world's first 31-band limiter.

THE AGC COMES FIRST

Our philosophy is that the on-air processor's multiband AGC has one job to do: keep the radio station's signal consistent in volume and spectral balance as source material changes. The Vorsis SST-enabled AGC deftly accomplishes that task with its invisible level and spectral management. With that now out of the way, the following peak limiter has only one thing left to do: control peak energy.

Wondering if it was possible to make a peak limiter just as clean and invisible as the SST-enabled AGC, we decided to find out.

PERCEPTION VS. REALITY

One surprise in human hearing research was that weaknesses of our auditory system can be leveraged to fool the ear into thinking it hears better sound than technically exists (aka, "how can it measure so bad and sound so good?").

In fact, one need only look to the complex technologies behind perceptual codecs like AAC and MP3 for hints on how to hide things from our ears. One of codec technology's better-known secrets is managing the audio spectrum in slices, or frequency bands, that are simply too narrow for our brains to dissect.

Audio researchers tell us that our audi-

tory system can be modeled as a filter bank with 25 overlapping bandpass filters; the 25 "bands" are known as "critical bands." What is important to remember about them is that our ears can't tell that there are separate signals inside a band whenever they occur simultaneously (see Fig. 1).

It gets even stranger than that. When there are audio signals present in different critical bands each signal is heard independently, as long as the signals are loud enough and far enough apart in frequency to stay above and away from an adjacent band's asymmetrical masking threshold (Fig. 2).

If an audio signal is soft enough in level or close enough in frequency to sneak under the masking threshold of an adjacent band, that signal is masked. It is rendered inaudible. Even "Golden Ears" can't hear it (Fig. 3).

The ability of audio signals to mask each other under certain conditions reveals important human psychoacoustic behaviors that we were able to tap into for creating our high-performance audio limiter.

SLICING THE PIZZA

When audio is divided up into numerous frequency bands, the energy within in each band is reduced according to how many bands the audio is divided into. The more bands there are, the less audio

there is in each band.

Probably not intuitive in Fig. 1 is another subtle clue: When there are 25 bands or more, the "sound of processing" within an




Fig. 1




Fig. 2




Fig. 3

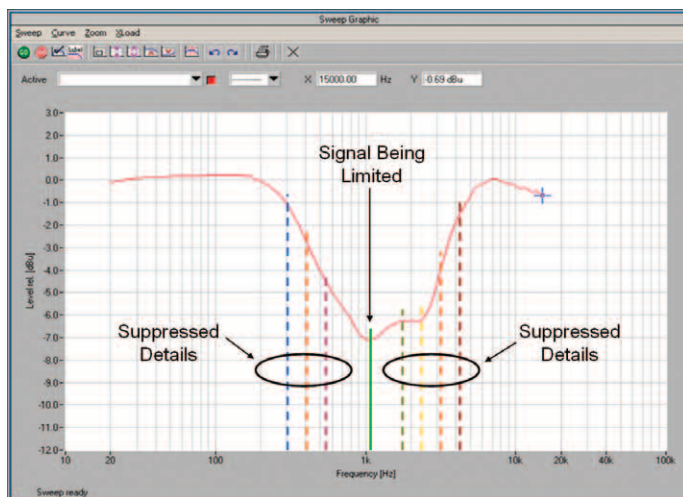


Fig. 4

individual band is also inaudible because it can't break the "in-band, multi-stimulus" masking rule. We decided that 31 limiter bands centered on ISO standard 1/3 octave frequencies would be a good place to start for building our limiter.

MORE LIMITER BANDS: BETTER?

Limiters with only a few bands are often seen with more than 10 dB of limiting depth. With 31 limiter bands, there is no longer a need to drive the limiters deep into gain reduction to create consistency and loudness. In fact, with 31 bands, only light limiting, perhaps 2 dB to 4 dB per band, is all that's required. Such shallow limiting makes it *impossible* for the limiters to create a smashed and fatiguing on-air sound.

A BETTER LISTENING EXPERIENCE, AND A SURPRISE ...

When implemented correctly, a 31-band limiter can manage the energy of electrical signals without our ears noticing that any work has been done. But there is another secret to reveal, and it's the one everyone notices first: how the limiter uncloaks fragile audio details often turned into incomprehensible mush by other processors' limiters.

In a conventional multiband limiter, the broadness of each limiter band (see Fig. 4) allows the act of limiting to affect a large portion of the audio spectrum, reducing signals that likely have no need to be reduced. In those limiters, quieter details coexisting in the band but below its limit threshold are pulled down along with the stronger signals, often to below audibility. Many users find themselves driving the multiband limiter

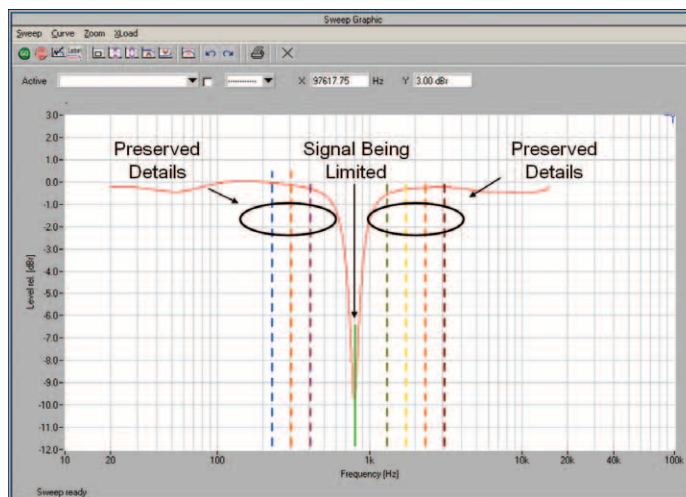


Fig. 5

even harder trying to get the lost detail back.

Conversely, as Fig. 5 shows, each band of a 31-band limiter is quite narrow. Note how little audio spectrum is affected by one limiter band. There is something even more important to notice: how the high selectivity of the limiter's crossover allows nearby audio detail to be completely untouched. It is all still there. This is completely different behavior from the way multiband broadcast limiters with only a few bands work. It sounds different too.

HOW DOES A 31-BAND LIMITER SOUND?

When a band of a 31-band limiter reduces its gain, two things happen — one expected, and one a surprise.

The expected thing is that the signal being limited's level has been restricted to the band's limit threshold, just as in any limiter.

The surprise is more interesting. The act of limiting a signal in one narrow band psychoacoustically (not electrically) raises the level of subtle audio details residing near to, but not inside of the band in limiting and its masking threshold.

The mechanism for this effect may need an explanation.

Even though the audio signals in the bands adjacent to the one in limiting have not undergone any modification, our brain has decoded it very differently and hears subtle details in the program material not often heard through other broadcast audio processors. Exactly why this effect is heard could be open to speculation. However we've come to the conclusion that the mechanism for it is probably quite simple:

The audio in the band undergoing limiting and the audio in nearby limiter bands not being limited have undergone a change in their relative gains. Our brain doesn't notice the effect of limiting because it is constrained to such a narrow band. But there's a perceived increase in the level of the signals in the "non-limited" bands, even though their electrical amplitudes have not been changed. This is entirely opposite behavior from what limiters with only a few bands do to the audio when they carve up huge chunks of the audio spectrum just to limit a single isolated signal.

The magic is that there is no magic, only science. The subtle audio details being revealed by the 31-band limiter have always been there; they are "part of the song," if you will. The first time people hear a Voris processor with a 31-band limiter, the openness of the audio and its vivid detail catch them off guard, just as the effortless and clean acoustical loudness gained by the 31-band limiter does, without having to drive either it, or the final clipper, "hard."

SUMMARY

In the beginning our goal was to return the lowly peak limiter, a device intended to manage peak energy and then get the heck out of the way, back to its roots. Believing that an audio limiter should be seen and not heard, we went about creating the world's first 31-band limiter. People can't believe that they won't hear it until they don't hear it.

Jeff Keith, CPBE NCE, is Voris senior product development engineer for Wheatstone Corp.