The Vorsis Bass Management System®

December 2007 – Jeff Keith

600 Industrial Drive New Bern, NC 28562 USA 252-638-7000
www.vorsis.com
techsupport@wheatstone.com

Background

The artistic handling of bass, particularly the deep and percussive bass present in modern program material presents significant challenges for a broadcast audio processor and this is especially true when competitive on-air loudness is desired. One reason for this is that there is simply a lot more energy at bass frequencies in today's contemporary music recordings.

Compounding the issue is the non-flat frequency response of the human ear which happens to be less sensitive to frequencies at the extremes of the audio spectrum. In order for bass frequencies to be 'prominent' on the air after the average level of midrange frequencies that the ear is more sensitive to have been raised by processing, some bass boost is required in the audio processor in order to restore a spectrally balanced sound.

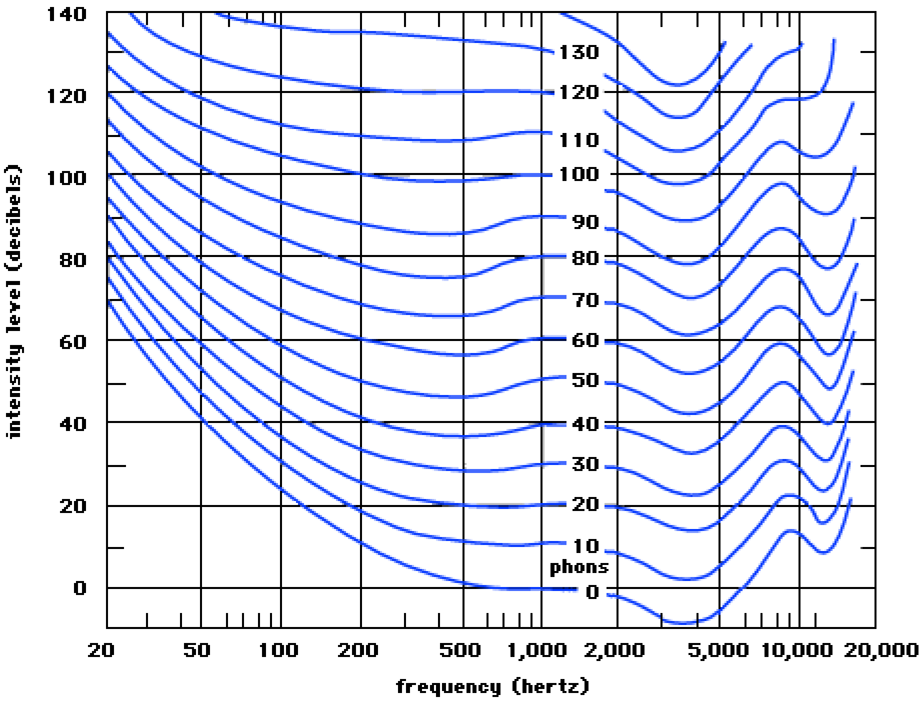

The well-known Fletcher-Munson curves are depicted above and illustrate the dramatically non-flat frequency response of the human ear at various audio frequencies and loudness levels. As the graph shows, almost regardless of the listening volume, higher acoustical power is required at low and high frequencies compared to those at mid frequencies in order for the ear to perceive equal loudness.

Getting a Handle on Bass

Conventional automatic gain control (AGC/compressors) can be used to effectively manage medium and long term bass program energy, but such dynamics control is wholly unsuitable for controlling short term percussive dynamics.

For instance, if sufficiently fast attack and release times are utilized in order to adequately control the electrical peak levels, the impact of percussive energy is reduced, leaving the bass ‘flat’ and ‘lifeless’. Intermodulation distortion can also occur on sustained bass when the necessarily rapid time constants are used. The resulting processing artifacts alter the character of on-air bass response in subjectively negative, unnatural, and certainly non-artistic ways.

Conventional Audio Processor Bass Control

One of the methods commonly used to manage percussive bass energy is the well known “Bass Clipper”. This strategy usually combines some form of dynamics control for managing sustained bass energy and leaves the peaky, percussive portions of the bass waveform to be constrained by a simple clipper. The first “bass clipper” that this author is aware of dates back to the custom-built Gregg Labs audio processors.

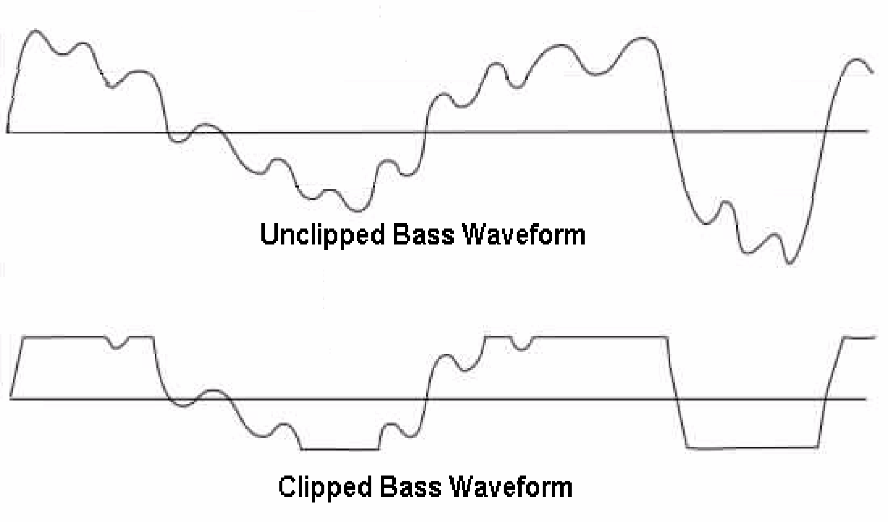

But… Clipping Generates Distortion

When performed in a perfectly symmetrical manner, audio clipping creates distortion containing only odd harmonics and filtering may be used to reduce those harmonics to an acceptable level. In fact most bass clipper topologies including the Gregg clipper mentioned previously have some form of post clipper filter which is intended to reduce the level of undesirable harmonic (distortion) energy created by clipping.

The clipper/filter topologies in some audio processors may be quite simplistic and/or can behave in undesirable ways when tasked to handle a wide variety of program material. Undesirable audible side effects such as strong intermodulation between bass and higher frequencies are quite common in such simplistic approaches.

The Vorsis Bass Management System®

Given our deep understanding of advanced signal processing and the shortcomings of bass clippers in other audio processors, it was clear to us that a completely different approach for handling bass in an on-air audio processor was both required and warranted.

Our design team was intent on tackling bass processing from a new perspective and the result of our work is the new Vorsis Bass Management System® (VBMS®) algorithm. In simplistic terms it is a set of high-performance DSP algorithms designed to overcome the shortcomings of previous bass amplitude control topologies while also outperforming them.

Making Bass Play Nice

The human ear is quite sensitive to the relationship between a complex signals' fundamental frequencies and its harmonics as it is the harmonics that give musical instruments their unique timbre. We knew we had to come up with a method to preserve harmonics in order for instruments to sound natural to the listener – and do it without adding noticeable and objectionable distortion.

The foregoing discussion brings up a failing of common bass clipper topologies: bass energy is brute force clipped and then filtered afterwards to remove most of the distortion products. While such a scheme does reduce distortion, it is a static process that always does the same thing regardless of the wave shape of the bass signal.

Unfortunately, such simplistic behavior also ‘disassociates’ the fundamental bass signal from the harmonics that distinguish it as belonging to a certain instrument. While it can create some mighty impressive bass “thud” on the air, the resulting bass lacks the associated detail that identifies what type of instrument generated the signal. This just isn’t musical.

A Vorsis® audio processor preserves bass harmonic relationships better and this is why musical instruments sound more natural processed by a Vorsis® processor - even when it has been tuned to be quite loud. Other audio processors with their conventional bass clippers tend to sound ‘clinical’ compared to one equipped with VBMS®.

The VBMS® algorithm uses various forms of masking to conceal the undesired dynamic and distortion products inevitably created. The difference is that it does so by manipulating the amplitude and phase of harmonics rather than simply removing them.

A Look Under The Hood

The Vorsis Bass Management System® contains a number of mathematically interrelated subsystems. In the most simplistic terms it is a set of rules-based algorithms making up a relatively small, DSP-based three-dimensional neural network that interoperates with the AGC/Compressors and multiband limiters. A portion of the computational load is handled by the multiband limiter’s Motorola DSP fixed point algorithms with the remainder by a Texas Instruments floating point DSP.

The explicit details of the VBMS® design will not be disclosed. However a rudimentary explanation is this: FFT spectral analyses are performed on the audio exiting the AGC/Compressor and multiband limiter and these analyses are used in various ways to tune the algorithm in real time and in a manner that satisfactorily manages certain forms of bass frequency spectra.

It should be no surprise that there are many time constants associated with the VBMS® algorithm and most are extremely critical to the unobtrusive operation of the algorithm. Those were hard-coded as fixed coefficients in DSP firmware after being very carefully tuned for predictable and transparent operation with an *extremely* wide variety of program material.

Certain less critical parameters were assigned to adjustable user controls located on the limiter control screens and enable tailoring of the behavior of the algorithm according to individual taste. Each of the adjustments has more than sufficient control range to enable deep exploration of *all* aspects of the algorithm that would be useful to an end user.

In summary, the VBMS® is tasked with ensuring that bass energy is always well controlled and bass harmonics are preserved in a way that sounds as artistic and natural as possible, regardless of competitive loudness requirements.

Bass Primer – More Than Just Fundamentals

A sound is said to have a missing fundamental (or phantom fundamental) when its harmonic overtones suggest to the brain that there is a presence of a fundamental frequency but the sound itself actually <u>lacks</u> that frequency.

The brain perceives pitch not just by the fundamental frequency, but also by the ratios of higher harmonics to the fundamental. In fact, we may still perceive a 'phantom' pitch (and even with a different timbre) even if the fundamental frequency is completely missing. This physiological effect is created entirely within the ear/brain interface and this oddity of the human auditory system can be tapped to create an illusion of deep bass in circumstances where it might not be otherwise audible.

Researchers once thought that the phantom fundamental effect occurred because the missing fundamental frequency was being replaced by distortion introduced by the various well known mechanical nonlinearities of the human hearing system. However, experiments later showed that even when a wideband noise that *should* have masked the natural nonlinear distortion of the ear was added to the signal, listeners *still* heard a pitch corresponding to the missing fundamental!

It is now widely believed that the brain processes the information present in overtones, and even the ones that the ear itself creates, to 'calculate' absent fundamental frequencies. The precise way in which this happens is not well understood but seems to be based on a form of autocorrelation-like operation involving the timing of neural impulses along the auditory nerve.

Beyond Physiological

The concept of manipulating bass overtones in the absence of the fundamental is not new, and in fact it is sometimes used to create the illusion of deeper bass in consumer-grade electronics gear. By selectively processing certain bass overtones a rich bass effect can be created even with relatively small speakers that simply cannot produce the fundamental frequencies. When correctly done the manipulation of bass overtones compels the brain to replace the missing low bass fundamentals allowing us to "hear" what we perceive as deep bass response from speakers of virtually any size.

The VBMS® algorithm contains mathematical methods to extract and reveal the bass fundamentals and their overtones if they exist, in order to create the illusion of deep bass on speakers large and small. The result is full, rich, natural sounding bass regardless of the speakers used to reproduce it. Note that this part of the algorithm doesn't actually *create* overtones, but 'looks for them' via FFT where it 'extracts' and amplifies them when present (conditional overtone creation occurs later).

Deep bass fundamentals are measured by the VBMS® algorithm and if present within a specified and automatically controlled bandwidth, certain harmonics are extracted and used as a reinforcement signal. The resulting bass is tight, natural, and never harsh, even when it is auditioned on very high quality monitoring systems or when program material has been mastered with far too much low end.

VBMS® Operation

A Drive control is available to the user which adjusts the relative position of the VBMS® operating window within the dynamic range exiting the AGC/Compressor. Higher settings of the control move the window 'down', causing more energy within the bandwidth set by the "Freq" control to be above the processing threshold resulting in higher bass loudness.

Lower settings of the control move the operating window 'up' and reduce bass processing. Such a setting might be utilized when the AGC/Compressor is being operated with excessively long attack times in combination with high thresholds in the lower frequency bands.

There are three "VBMS®" operating styles available to the user called Hard, Soft, and OFF.

When switched off the algorithm is completely out of circuit although there is little benefit to doing this because it is smart enough to 'do nothing' when that is required.

The hard setting causes the controller to operate with a 0.25dB transition knee and because it is moderately abrupt generates low and high order harmonics. In a process that borrows from the technology behind the Timbral mode in our Vorsis AP1000 we mathematically manipulate certain harmonic energy occurring at and *above* the second harmonic.

At the soft setting the controller has a transfer function of 1.0dB. This is less abrupt than the hard setting and generates few high order harmonics. Like the hard mode, we again manipulate certain high order harmonic energy at and above the second harmonic.

Most facets of the VBMS® algorithm are completely automatic. Observed variations in behavior will always be appropriate for the frequencies and waveforms present. Bass energy is favorably managed in an artistic way in contrast to the dull, disassociated ‘thud’ when heard through the bass control mechanisms of other processors. It is important to remember that certain harmonics are never entirely eliminated but are instead manipulated in amplitude and phase to favorably alter both the sound as well as the waveform <u>shape</u> so that it remains useful for good amplitude control.

The operating bandwidth is adjustable over the range of 60Hz to 300Hz and only audio <u>below</u> the Frequency setting is processed by the algorithm. A lower setting produces ‘heavier’ and ‘tighter’ bass while higher settings produce ‘broader’, ‘warmer’, bass. Regardless of the setting bass energy of any magnitude and frequency is favorably controlled.

The Frequency tuning control may be adjusted to personal taste and program format, although there seems to be a general ratiometric relationship between the AGC/Compressor Band 1 to Band 2 crossover frequency and the setting of the Frequency control that sounds best to *our* ears. The numeric range of this ratio seems to be between **1.26** and **1.40**.

Example:
Suppose the AGC/Compressor Band 1 to Band 2 crossover is set at 80Hz. In that case the “ratiometrically optimum” VBMS® Frequency would be somewhere between 108Hz and 112Hz because:

<u>80Hz * 1.26 = 108Hz</u> and <u>80Hz * 1.40 = 112Hz</u>

Time Alignment

The design process presented us with some challenging time alignment requirements because certain signals within the algorithm need to be fed forward or backward to other sections of the same algorithm or to other parts of the main signal processing chain. Rest assured that the following considerations have been adequately addressed:

- The signals from the various outputs of the algorithm always sum correctly into multiple summing nodes within the main processing chain.

- The lower frequencies handled by the various VBMS® signal paths have longer time delays than the non-VBMS® paths and these delays have been compensated for.

- Combining of the many dissimilar latency signals within the signal processing chain will occur correctly *regardless* of the setting of the operating controls.

Great care was exercised during the design cycle to ensure that signals not processed by the various parts of the algorithm always undergo sample accurate time alignment as operating control settings are altered. This ensures proper phase and time coincidence with program energy that *is* processed by the signal paths within the VBMS® algorithm.

Safe Area Operation

There are no particular cautions associated with the setting of the VBMS® Output and Drive controls other than some care should be exercised when setting the Output control higher than approximately +3dB *when* the Drive is *also* set higher than approximately -2.0dB. Under those conditions certain program material *might* be able to 'unhide' spectra that the algorithm is charged with keeping hidden below the masking threshold of its adjacent critical bands.

The graphic below reveals the operating relationship between the Drive and Output controls. A careful analysis of the graph shows that virtually all adjustment combinations falling within the shaded area are "safe" and should prove satisfactory with a wide variety of program material.

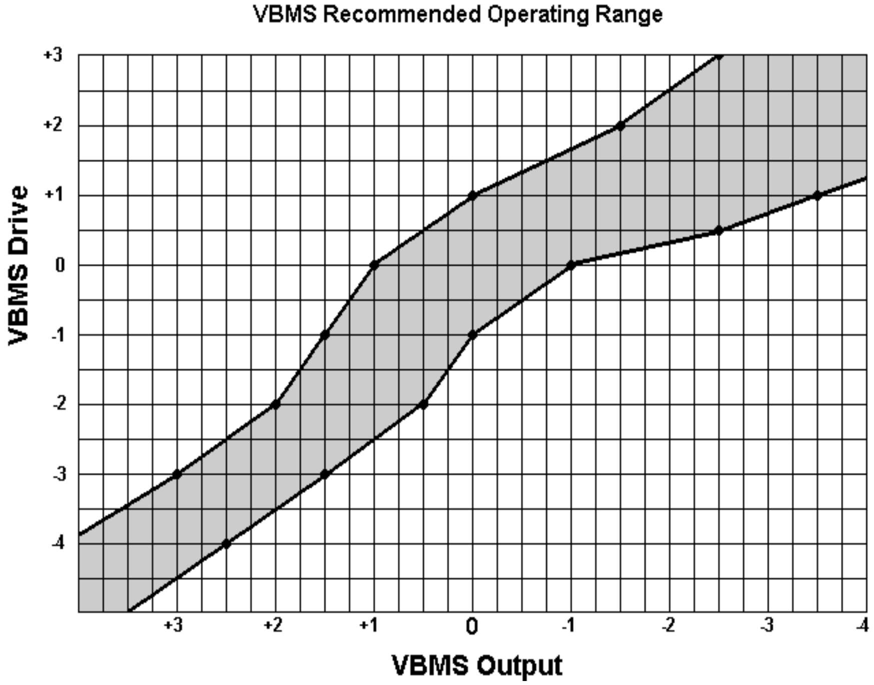

The higher the Frequency *and* Drive settings are, the more program energy that will be managed by the algorithm. Generally speaking there is no right or wrong settings. Whatever sounds best for the format, market, and competitive situation are correct.

Intentionally overdriving the algorithm input with excessively high Drive settings will not cause the obnoxious forms of distortion normally associated with “bass clippers”. This is because the algorithm knows about certain program energies in its incoming and outgoing signal paths and automatically adapts to *minimize* unintentional intermodulation spectra that might otherwise result from aggressive settings or extremely bass-heavy program material.

VBMS missteps…

The algorithm has a predictable and repeatable dynamic error that varies with the setting of the Drive control and the amount of low frequency dynamic range present at the output of the five-band AGC/Compressor.

This error is somewhat parabolic and while being measurable using special test signals is rarely audible with program material. This is because departures from ideal behavior are usually inaudible in the presence of program material because the error is *masked* by that material.

The graphic below demonstrates how the operation of the algorithm deviates from the perfect mathematical model.

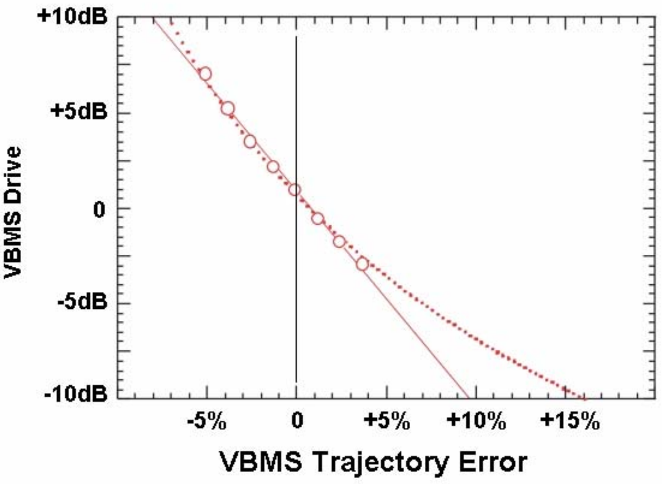

If it were possible, theoretically perfect behavior would be represented by the vertical line extending above the “0” on the horizontal axis.

The straight diagonal line depicts the algorithm’s deviation from ideal behavior at both extremes of the drive control settings.

Slightly parabolic is shape is the line representing the algorithm’s departure from the ideal when excited by certain laboratory test stimuli having dynamic characteristics similar to typical program material.

Circles located along the diagonal axis depict the almost insignificant difference in behavior between the hard and soft modes.

Departure of the algorithm from perfect behavior causes a very slight tendency to over control bass rather than under control it and it is caused by intentional mathematical rounding in the module of the algorithm responsible for predicting bass waveform crest factors.

The same algorithm is also tasked with computing waveform trajectories at dynamically selected frequencies outside the VBMS® operating bandwidth in response to certain program-related factors that gives the algorithm hints about program content in the remainder of the audio spectrum.

Although it results in less than perfect algorithm operation, the rounding off of certain dynamic coefficients is necessary in order to constrain the amount of DSP resources required to operate it. This is because nearly infinite DSP resources (as well as processing time!) would be required to operate the model “perfectly”.

**Wheatstone Corporation
600 Industrial Drive, New Bern, North Carolina 28562
252-638-7000
www.vorsis.com
techsupport@wheatstone.com**