VORSIS Application Note

AN2007-02

The Vorsis Bass Management System™

January 2008 - J. 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.

Further compounding the issue is the human ear’s non-flat frequency response making it 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.

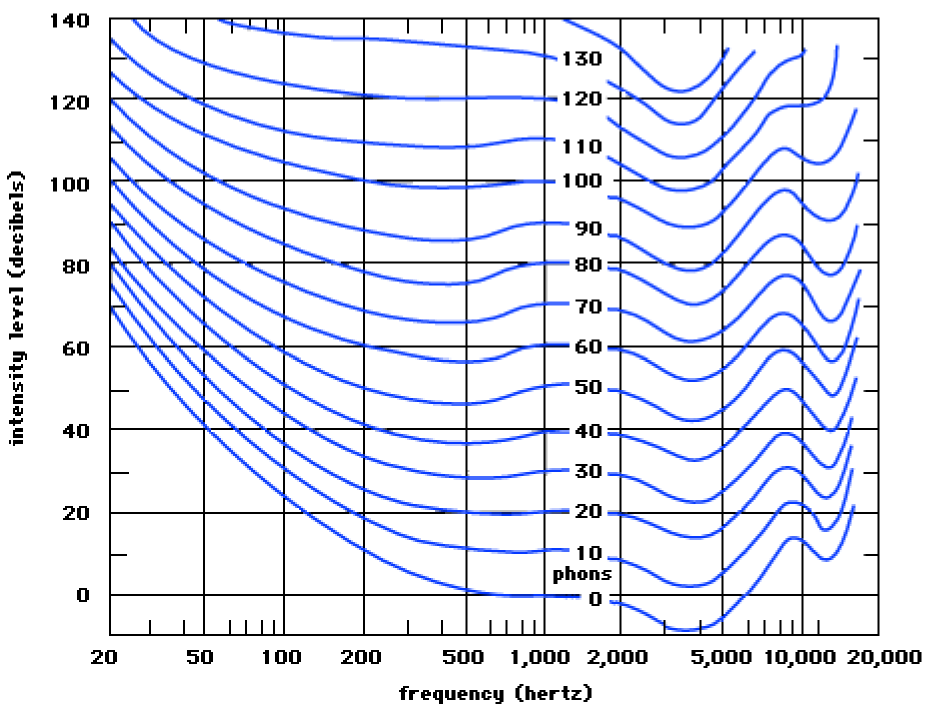

The well-known Fletcher-Munson curves are depicted above and illustrate the dramatically non-flat frequency response of the human ear at various audio frequencies and loudness levels. The graph shows that almost regardless of the listening volume higher acoustical powers are required at low and high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

Getting a Handle on Bass

Conventional automatic gain control such as AGC's and compressors can be used to manage medium and long term bass program energy, but such forms of dynamics control are wholly unsuitable for controlling the short term percussive bass dynamics.

For instance, if sufficiently fast attack and release times are utilized in order to 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 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 conventional 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.

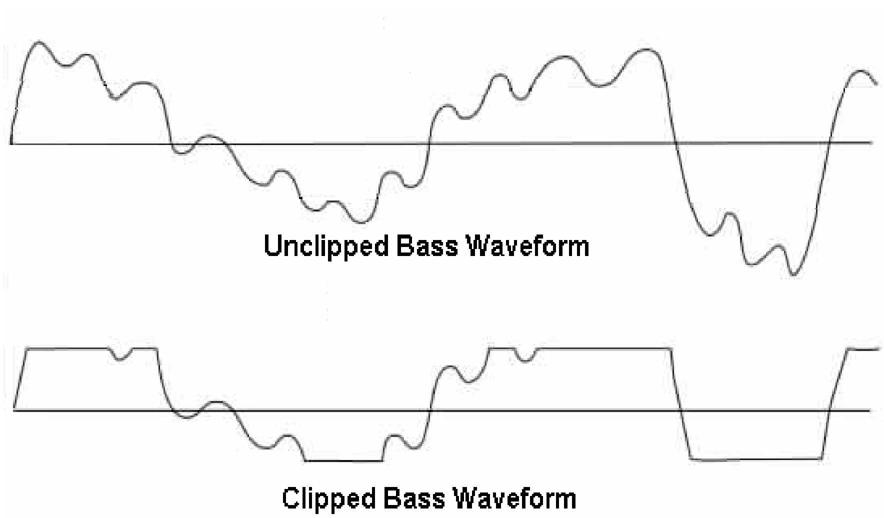

But…Clipping Generates Distortion

When performed in a 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 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 those simplistic approaches.

The Vorsis Bass Management System™

Given our understanding of advanced signal processing and the shortcomings of "bass clippers" in other audio processors, it was clear 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 the 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 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. Unfortunately this isn't all that easy because the goals tend to be mutually exclusive

The foregoing discussion brings up a weakness 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 character of the wave shape of the bass signal. This behavior 'disassociates' the fundamental bass signal from the harmonics that distinguish it as belonging to a certain instrument and it also raises the peak level because the harmonics have been removed. This is exactly the opposite of what we were trying to accomplish!

While the above can create bass with a definite "thud", that bass will lack the associated detail that identifies what type of instrument generated the low frequency transient in the first place. This just isn't musical.

Because Vorsis™ audio processors preserve important harmonic relationships better they will sound more natural on a wider variety of program material - even when tuned to be quite competitively loud.

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

(In early AP1000 and FM5 firmware we referred to the VBMS™ function as a "Bass Clipper" and its operating modes as "Clipper styles" though it differs quite significantly from known forms of "bass clippers". We felt that similar operating terms might make VBMS™ easier to operate for those with experience with conventional bass clippers. Because the VBMS™ is intelligent and many aspects of its operation 'automatic', even less experienced users can achieve excellent on-air results.)

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 small, three-dimensional neural network that interoperates with the AGC/Compressors and multiband limiters. A portion of the required computational load is handled by the multiband limiter's Motorola DSP fixed point algorithms with the remainder by Texas Instruments floating point DSP.

The explicit details of VBMS™ operation will not be disclosed nor will we explain why both fixed and floating point DSP is involved in the algorithm. 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 in a manner that satisfactorily manages certain forms of bass frequency spectra.

It should be no surprise that there are many time constants inside the VBMS™ and most are *extremely* critical to the unobtrusive operation of the algorithm. The critical constants 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.

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

The VBMS™ operates in real time to ensure that bass energy is always well controlled and that the bass harmonics are preserved in a way that sounds artistic and natural. This adds quite a bit to the DSP processing load however the reason for this extra step will soon be become clear.

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 in fact lacks that frequency.

The brain perceives tone pitch not just by the fundamental frequency, but also by the ratios of any higher harmonics to the fundamental. In fact, we may 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 may be utilized to create the illusion of deep bass in circumstances where it would not be otherwise audible.

Researchers once thought that the phantom fundamental effect occurred because the missing fundamental 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 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 somehow 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 may be based on a form of autocorrelation involving the timing of neural impulses along the auditory nerve.

Beyond Physiological

The concept of bass overtones being manipulated 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 reinforce 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 certain harmonics are reinforced within a specified and automatically controlled dynamic bandwidth. The resulting bass is tight, natural, and never harsh, even when auditioned on very high quality monitoring systems or when program material has been mastered with far too much low end.

VBMS™ Operating Hints

VBMS™ Drive

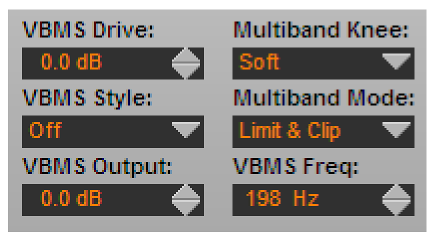

The Drive control adjusts the relative position of the 3dB-wide operating window within the available dynamic range exiting the five band AGC/Compressor and is adjustable over a +/- 10dB range.

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

In typical operation and with reasonable time constants and thresholds in the five band AGC/Compressor, the "normal" setting of the Drive control will be found between -4.5dB and +1.0dB.

Important! The Drive Control is purposely labeled ‘upside down’ in order to make operation more intuitive - higher settings *increase* VBMS™ processing while lower numbers *decrease* it!

VBMS™ Style

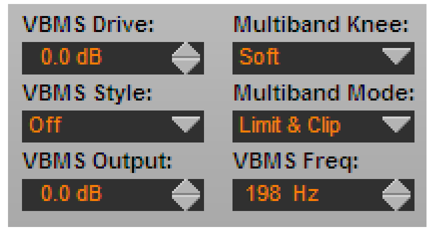

There are three “VBMS™” operating styles; Hard, Soft, and OFF.

OFF - the VBMS™ is completely out of circuit. There is little benefit to switching the algorithm off because it is smart enough to ‘do nothing’ when ‘doing nothing’ is required.

Hard – The Hard setting has a transfer function with approximately a 0.25dB transition knee and because this knee is moderately abrupt it generates both low and high order harmonics. Afterwards, in a process that borrows from the underlying technology behind the Timbral mode in our Vorsis AP1000 audio processor, we mathematically manipulate certain harmonic energy occurring at and *above* the second harmonic.

Soft – The Soft setting operates with a transfer function with approximately a 1.0dB transition knee. This transfer function is less abrupt than the Hard setting and therefore generates fewer high order harmonics. In an operation that has similarities to that used in the Hard mode we manipulate certain high order harmonic energy at and *above* the second harmonic.

(Note that certain harmonics are never entirely eliminated but are instead manipulated in amplitude and phase to favorably alter the sound of the Hard setting in a manner that also retains a waveform <u>shape</u> that is important to amplitude control.)

Most facets of the VBMS™ algorithm are completely automatic and observed variations in its behavior between different program elements will always be appropriate for the frequencies and waveforms present. Bass peak energy is favorably managed in an artistic way in contrast to the dull, disassociated ‘thud’ accompanying bass drum hits when heard through the bass control mechanisms in other audio processors.

Due to the differences in the Hard and Soft transfer functions - how and where they depart from linear operation, and how, where, and how much the harmonic spectra is manipulated the Hard and Soft styles can sound *significantly* different on certain program material.

VBMS™ Frequency

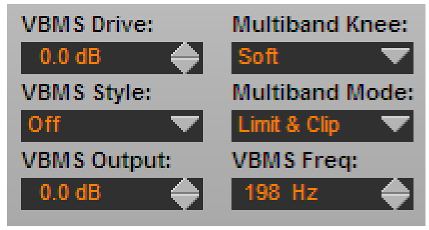

The VBMS™ Frequency is adjustable over the range of 60Hz to 300Hz with audio information that is present <u>below</u> the Frequency setting being processed by the algorithm.

Generally speaking a lower Frequency setting produces ‘heavier’ and ‘tighter’ bass while higher settings produce ‘broader’, ‘warmer’, bass.

Regardless of the Frequency setting bass energy of any magnitude and frequency is favorably controlled.

The Frequency 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:

80Hz * 1.26 = 108Hz, and 80Hz * 1.40 = 112Hz

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.

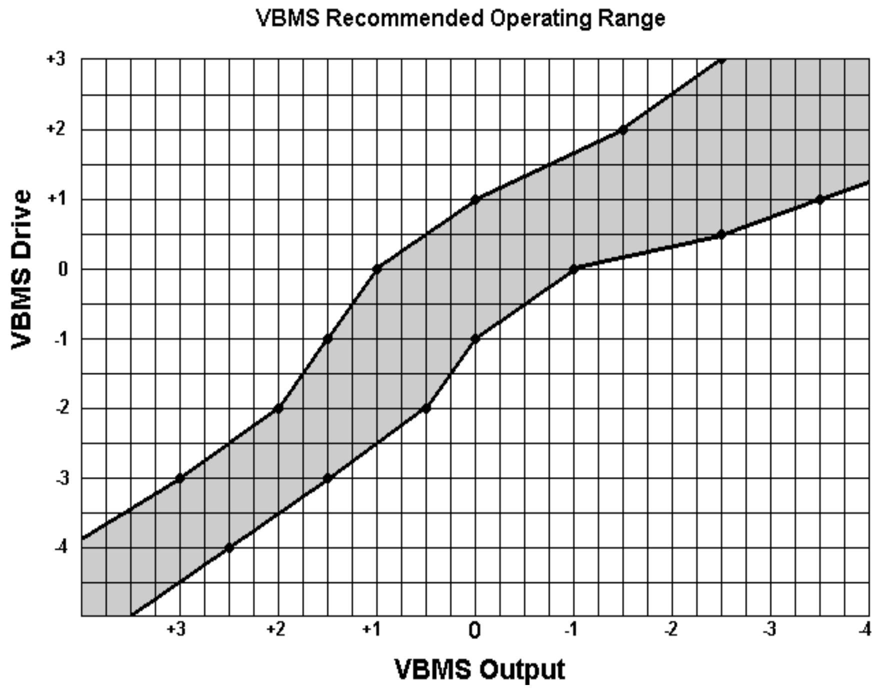

VBMS™ Output

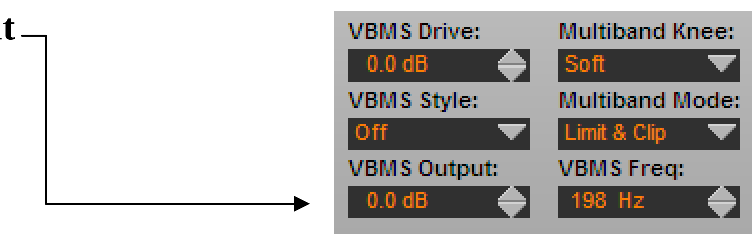

There are no particular cautions associated with adjusting the VBMS™ Output and Drive controls other than some care should be exercised when setting the Output control higher than approximately +3dB *when* the Drive is *also* set higher than approximately -2.0dB. This is because under those conditions certain program material *might* cause sporadic unmasking of spectra that the algorithm is charged with keeping hidden below the masking threshold of its adjacent critical bands.

The higher the VBMS™ Frequency *and* Drive settings the more program energy there is to be managed by the algorithm. Generally speaking there is no right or wrong amount of VBMS™ processing – whatever settings sound best for the format, market, and competitive situation are correct.

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

Multiband Mode

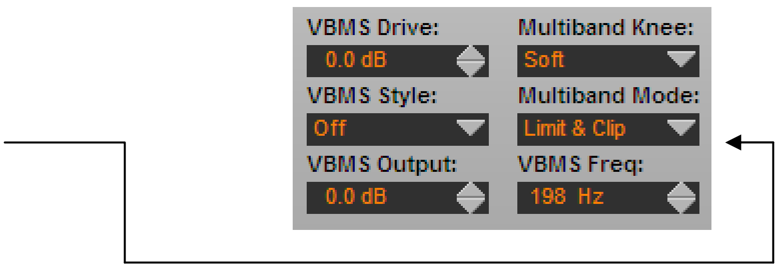

The Multiband Mode control is present ***only*** on the AP1000 product and selects one of four different operating modes for the 31 band limiter and its associated 31 band distortion masked clipper.

The FM5 product contains a ten band limiter instead of the 31 band version that exists in the AP1000 - it also does not have the multiband distortion masked clippers. Because the multiband clippers are not utilized by the VBMS™ algorithm this difference in feature set has no impact on the performance of the VBMS™ in the AM5 and FM5 products.

Please refer to the AP1000 Operating Manual for a detailed description of the Multiband Mode control.

Time Alignment

The VBMS™ design process presented some challenging time alignment requirements because certain control and audio signals within the algorithm need to be fed forward or backward to other sections of the same algorithm or conditionally routed 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 the 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.

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 VBMS™ algorithm deviates from the perfect mathematical model.

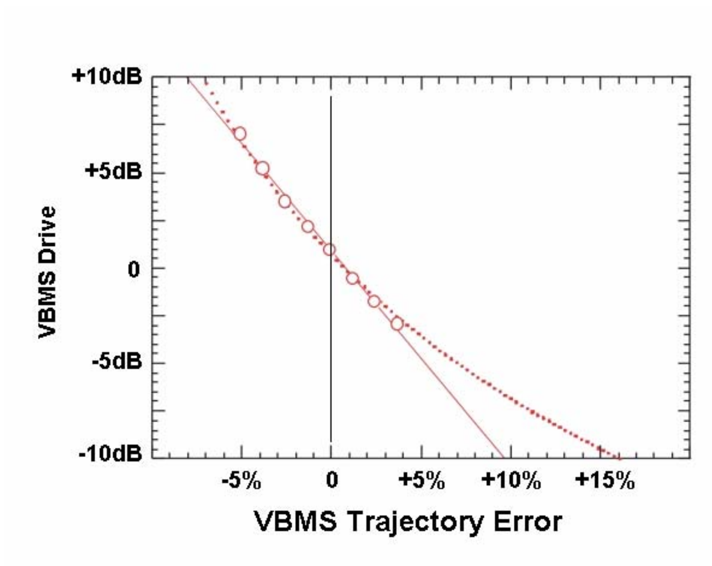

If it were possible, theoretically perfect behavior would be represented by the vertical line extending above the "0" on the horizontal axis.

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

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

Circles located along the diagonal axis depict the almost insignificant difference in error between the Hard and Soft operating modes.

Departure of the algorithm from perfect behavior causes a very slight tendency to over control bass rather than under control it. This is caused by intentional mathematical rounding in the module of the algorithm responsible for performing real-time prediction of 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 of certain dynamically derived 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 needed to operate the model "perfectly".

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