

VORSIS Application Note

AN2008-03A

Introduction to the Stereo Enhancement Techniques Used in Vorsis On-Air Audio Processors

February 2008 - J. Keith
Revised September 2008 – J. Keith

Wheatstone Corporation © 2008 All Rights Reserved
600 Industrial Drive New Bern, NC 28562 USA 252-638-7000

www.vorsis.com
techsupport@wheatstone.com

Background

Vorsis customers sometimes ask our sales or technical support staff “do your on-air processors have on-board stereo enhancement or should I add an enhancement unit outboard?”

The answer to the first part of this question is “yes”, and it is because of the way we designed our multiband AGC/Compressor algorithm. That makes the answer to the second part of the question “no, it’s not necessary”.

We realized when we designed our AGC/Compressor algorithms that there were some beneficial features to include and one of them was stereo enhancement. As experienced audio processor designers for both the professional sound and broadcast audio markets we intimately know the up *and* downsides of *all* of the various ways of accomplishing stereo enhancement for FM broadcast. That is why we chose to create our own.

Understanding Stereo Sound Field Enhancement

There are four basic stereo enhancement algorithms commonly found in the broadcast environment (others are simply derivatives). All implementations use a manipulation of the amplitude of the L-R, or difference signal, as a way to create a wider stereo sound field. Schemes based on this structure can be quite effective as long as the technical limitations are understood and the effect is not overdone.

In each implementation there is a means to extract the stereo difference signal from the incoming program material – this is simply the difference between the signals present in the left and right channels. Note that when a stereo signal path is carrying identical signals in the left and right channels there is no difference signal, and in fact this is exactly the condition that exists for a ‘mono’ signal.

As program conditions create a larger difference between the left and right channels, the signal amplitude in the L-R becomes higher. An “extreme” L-R example might be an early stereo Beatles recording – there was rarely anything present in one channel that was also present in the other. In this

case the sum and difference (L+R, and L-R) channels can have virtually the same signal amplitudes while carrying completely different signals.

In normal stereo programming there is a varying amount of L-R and the usual method used to increase perceived stereo separation is to alter the relationship between the L+R (sum) and L-R (difference) signals so that the average level of the L-R channel is increased.

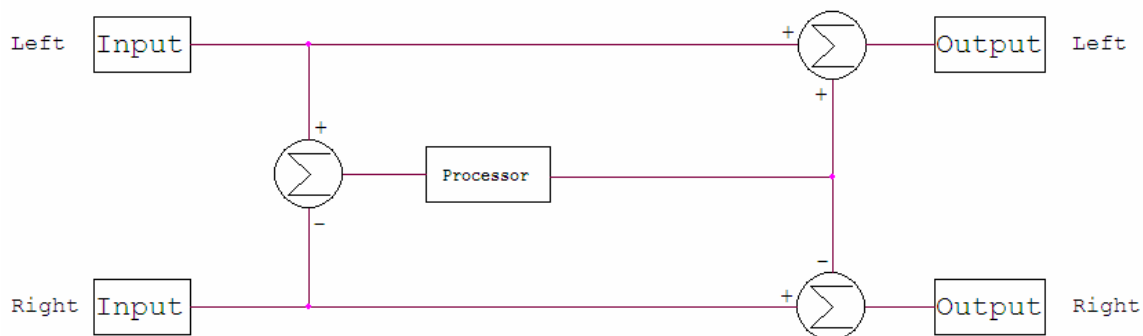
One could in fact *reduce* the level of the L+R channel instead and achieve the same thing, although stereo enhancement techniques *usually* manipulate only the gain of the L-R path so that it intensifies the directionality of the stereo sound field, making it sound wider.

Manipulating a Difference

The four general ways used to manipulate the stereo difference signal will be covered next. We will not discuss the means by which any of these schemes prevents over enhancing the L-R difference signal as they are all functionally similar and vary only in their philosophical aspects. We will also not cover schemes used to control the behavior of the algorithm in the absence of stereo programming for the same reason.

The Original L-R Booster

First, there is the conceptually simple “L-R booster” shown schematically below.



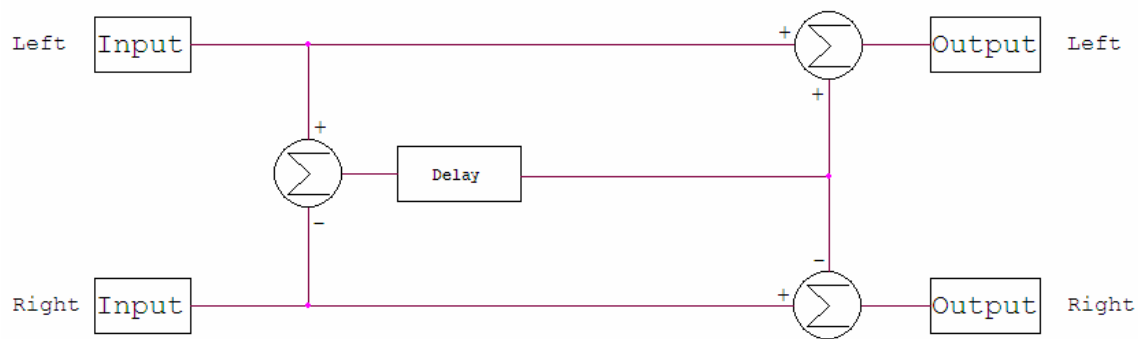
In this method a difference amplifier extracts the L-R signal and passes it on to a “processing” stage which alters the character of that signal. This

‘processing’ stage could be nothing more than a level control with some gain! The altered L-R signal is then re-injected in the proper phase to the left and right channels and a wider stereo sound stage results.

The “processor” shown in the center block can also be one that adds signal dynamics such as compression, limiting or expansion, time delay, high or low pass filtering, or combinations thereof. With the exception of the time delay case, this is probably the most “pure” of the schemes because it alters only the difference signal and therefore is 100% mono compatible.

Adding Time Delay to “Scramble” the L-R (Also referred to as “decorrelation”)

When time delay is introduced into the difference channel (as shown below by the block labeled “Delay”) an entirely different stereo effect is produced. Because of the delay (usually in the 5 to 20 millisecond range) the L-R no longer mathematically combines back with the L+R in the same way – the zero crossings of the sum and difference signals are now offset by whatever the delay time happens to be.



There is a comb filter effect produced when a delayed L-R signal is mixed back with the non-delayed L+R signal. Comb filtering, when applied in this manner to the L-R path, produces an uneven frequency response that has complementary peaks and dips in the left and right channels. That means that whenever the left channel has a boost at a particular frequency, the right channel has a corresponding dip in its response at that frequency. These peaks and dips are related to the delay time – in very simplistic terms shorter

time delays cause a higher number of peaks and dips throughout the audio frequency spectrum and longer time delays produce fewer of them.

Overall, the sonic effect of adding time delay to the L-R signal is a 'smearing' of the stereo difference signal across the left to right soundstage. Some find the very wide and spacious effect odd-sounding but still others find it quite pleasing and hence the popularity of the "Maxx" type device.

One of the downsides of the "Maxx" technique is that accurate stereo image localization is replaced with a general spreading out of the stereo sound field. The reason that stereo image localization accuracy is reduced is that the L+R and L-R signals must remain in time and phase synchronization in order to place the components in each channel back in the proper location from a hearing perspective, and the time delay used in the Maxx's L-R path significantly upsets those relationships.

Expanding On an Idea

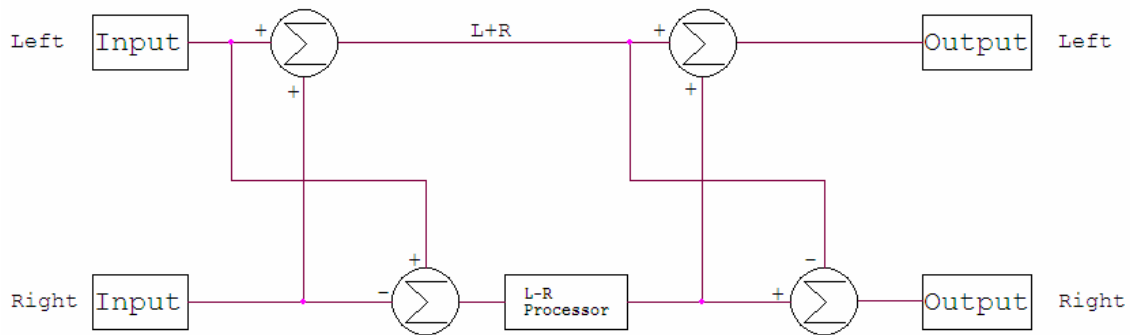
Rather than use time delay or conventional dynamic range reducing means, an expander can be used as the "processor" block instead. There are two schemes using expanders in the L-R channel:

One scheme utilizes conventional downward expansion with the expander sidechain triggered from the presence of the L-R difference signal. In this scheme when a stereo input signal is present and its L-R signal is above the expander's threshold the expander sets the L-R gain to some fixed boost or cut value in order to set a desired L+R/L-R ratio.

Another scheme uses the L+R channel (in fact, the dynamic transients occurring in that channel) to trigger an L-R expander only on the transient L+R components of the input program material. This results in a 'triggered' widening of the stereo sound stage and because it is triggering on transients only it doesn't increase the average level of the L-R. Therefore this scheme can be somewhat less prone to increasing multipath effects on an FM station that is prone to such artifacts in its coverage area.

Separate Sum and Difference Processing

Another method to manipulate the stereo sound field is to separate the left and right channels into their sum and difference signals and process each differently. The diagram below shows how the signal flow for this scheme might be arranged.



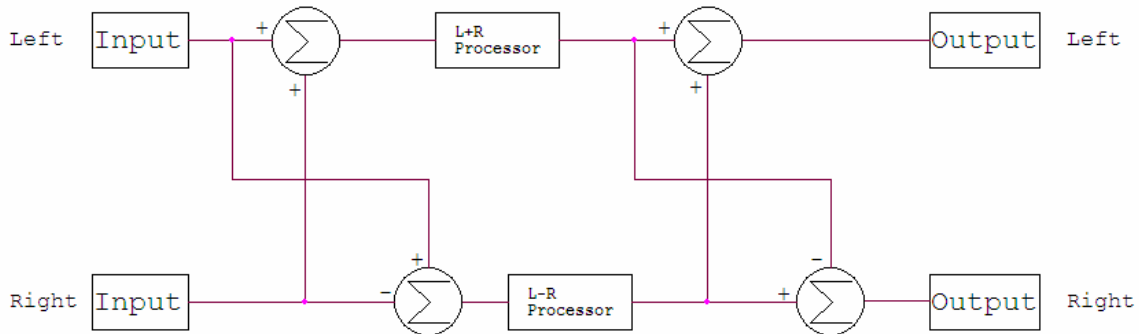
In the above the L+R and L-R signals are extracted using a stereo matrix. The L+R component is then passed without modification to the output matrix which reassembles the L+R and L-R back into the left and right channels.

However, unlike the L+R signal, the L-R difference signal undergoes processing in order to change its character. From an 'enhancement' standpoint this is functionally identical to the first scheme we discussed, however one benefit of this scheme is that the L-R signal can be fully attenuated to zero if desired – this results in identical left and right signals – a mono signal!

Coupled with program analysis intelligence this scheme can be used to produce excellent control of the L-R difference signal in real time on a program-adaptive basis to maintain an artistically stable stereo sound field as program elements change. In fact the author designed and used such a device at oldies station WMJI in Cleveland to successfully manage the wide discrepancy between the recording techniques of the various eras on the stations' oldies music format.

Two (and admittedly quite unexpected) benefits of the design once placed in the air chain at WMJI were a dramatic reduction in the audible effects of multipath-type disturbances in the FM signal and an overall increase in loudness and 'smoothness' of the station's on air sound.

The fourth scheme shown below separates the left and right channels into their L+R and L-R components and then processes each separately.



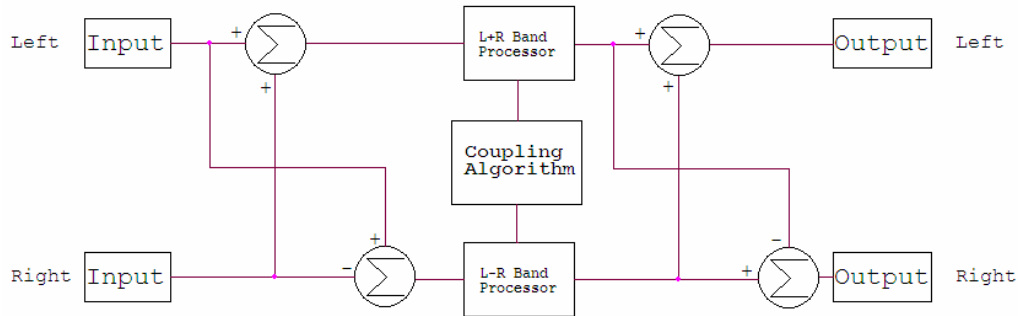
This scheme is different from the preceding one in that the L+R signal is *also* being manipulated. Although this scheme can be quite a bit more complex it has the potential to achieve several desirable goals at the same time:

1. It can control the width of the stereo image as desired;
2. It can control the stereo image width on a frequency-conscious basis;
3. It can use L+R to L-R signal ganging to prevent the stereo image from wandering uncontrolled.

With major modifications this is the scheme utilized in all Voris on-air processors and in our special implementation, the sum and difference channels are processed separately while in the AGC/Compressor domain. This means that depending on the multiband operating mode these signals are being processed by a three or five band processor. Normally this would create some very strange stereo sound field artifacts, but in our implementation it doesn't. Here's why:

For a moment consider that the previous block diagram applies to just one of the five AGC/Compressor bands. Then consider what would happen if the AGC gain control signals for the L+R sum channel were logically ganged to the AGC for the L-R difference channel so that their gains could not wander with respect to each another. And the end result? Not a thing....

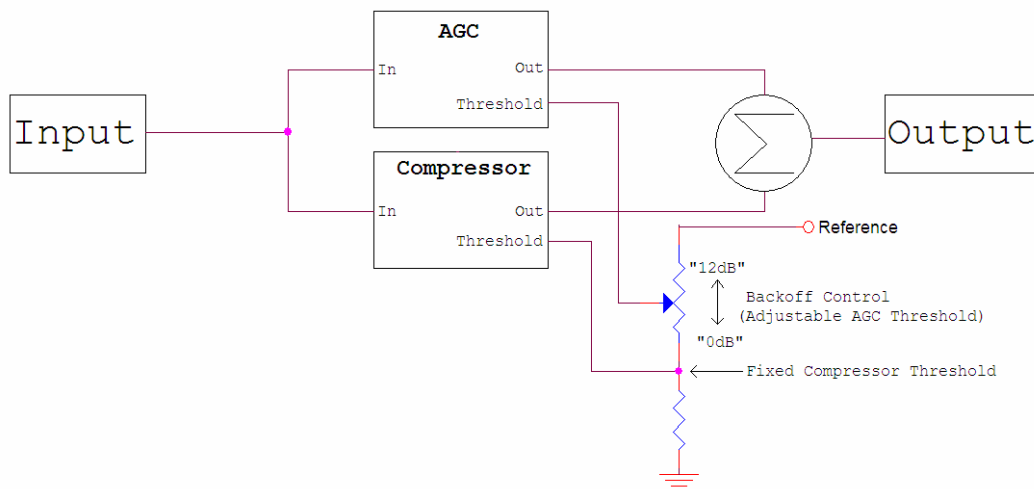
If the control signals for the faster-operating Compressor were not ganged between L+R and L-R it would allow the compressor's faster time constants to independently work on the L+R and L-R while the slower AGC time constants being ganged holds everything reasonable. This will widen the stereo sound field in a very controlled and program dependent way. Schematically it looks like this:



With proper intelligence in the L+R/L-R Coupling Algorithm that ties the processing sections together very careful and artistic management of the stereo image can be achieved without *any* of the artifacts that plague earlier and more simplistic approaches.

Back Off and Get More?

There is a special control in Vorsis on-air processors affecting L+R/L-R relationships called "AGC Backoff". Found on the AGC/Compressor screen it changes the operating threshold of the AGC compared to that of the Compressor. When the Backoff control is set to "0", the AGC and Compressors have the same detector threshold. See the simplified schematic below:



When the Backoff control is set below “0”, the AGC threshold is *raised* by that amount over that of the Compressor. A higher AGC threshold causes it to do less work and therefore the faster Compressor must do more. While this increases the average level of the L-R it can never exceed the L+R because the Compressor can never have more ‘gain’ than the AGC. This prevents the hole-in-the-middle effect commonly associated with blind increases in L-R energy.

When Compress Really Means Expand...

Another aspect of our stereo enhancement technique makes it vaguely similar in some respects to the ‘expander’ scheme. If you consider for a moment the characteristics of the L+R and L-R signals you realize that most of the transient information in program material is in the L+R channel.

What this means is that when the faster time constants of the compressor work to quickly control L+R signal excursions above its threshold, because the L+R and L-R are *uncoupled* at the compression stage, the L+R ‘ducks’ its gain more and more quickly than the L-R does, providing a momentary boost in the L-R level. This behavior occurs on a band by band basis and results in a noticeable opening up of the stereo sound field on virtually all stereo program material.

Tweaking Hint:

According to the above explanation , this also means that as the AGC attack time is increased, especially above approximately 150 milliseconds, the Compressors will do more short term work than the AGC’s will and therefore the apparentness of the L+R ‘ducking’ effect will increase. The ganging of the L+R and L-R control signals at the AGC time constant allows this effect to be quite dramatic if desired and without increasing the L-R to the point where it becomes unnatural or prone to inducing multipath. This is not to say that the effect cannot be overdone – any effect can – just that it may not create the nasty artifacts that other methods of stereo image management will create at the same level of enhancement.

Because this technique does not increase the average level of the L-R above that of the L+R there is less chance for generating or exacerbating multipath-type effects in the station’s coverage area.

Mixing Things Up

Vorsis on-air processors also provide another tool to manipulate how the stereo separation sounds on your radio station. Those are the Sum and Diff mixer controls in the AGC/Comp menu of the GUI which allow the multiband outputs from the sum and difference channels to be individually set. By manipulating these mixer controls the L+R and L-R can have completely different overall frequency responses, allowing the stereo sound field spectral texture to be set as desired.

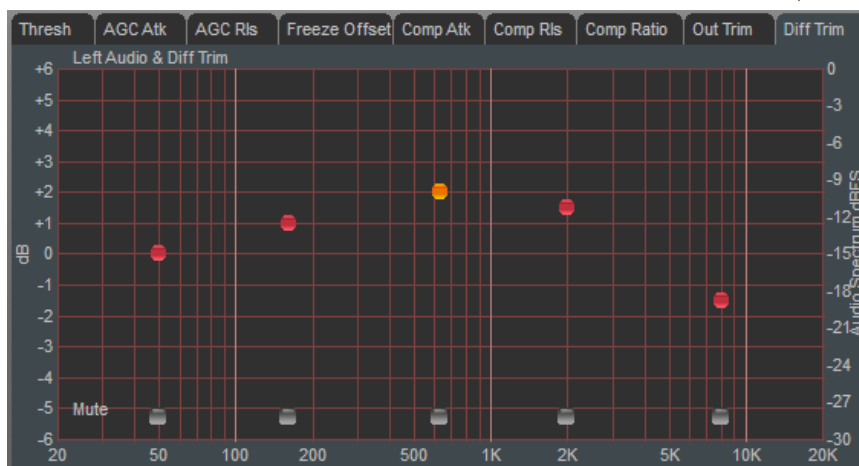
Below is a screen shot of the Sum mixer where the output levels of the five bands (the dots representing each band) have been set flat (“0dB”) except for Band 1 which has a +1.5dB boost.

Sum Output Trim Tab



Compare this to the Diff (difference) Mixer graphic below:

Diff Output Trim Tab



Notice that the Difference channel mixer's mid-band frequency response has *not* been set as in the Sum screen. Instead, we've adjusted the frequency response of the L-R channel for a slight midrange boost accompanied by a minus 1.5dB roll off for Band #5. We created this special curve to:


1. Slightly exaggerate the mid-frequency stereo separation;
2. Provide a slight roll off the high frequency extremes to prevent unnatural noise buildup or harsh exaggeration of coding artifacts if the program content happens to have been compressed by a perceptual codec.

If desired we can also adjust the overall drive and output levels of the L-R processing chain to accomplish an overall "Stereo Gloss" to the sound stage. This is done using the two "Diff" controls on the AGC screen.

Because the Compressors are not ganged at the L+R/L-R control signal, driving the input of the Diff chain a little harder by setting the control above "0dB" will result in a higher average input level for the stereo difference signal. This will increase 'processing' on the L-R which will widen the stereo image.

The "Diff" Drive control is located under the main Drive control and it adjusts how hard the AGC/Compressor is being driven by the L-R difference signal relative to that of the L+R Sum.

Note: when the "Diff" control is set at zero the drive settings for the L+R and L-R processing have equal gains.



There is also a "Diff" Output control which is located just below the one for Makeup Gain. This control sets how much the output of the L-R processing chain is allowed to be different from that of the L+R.

Note: As in the case of the input Diff control when it is set to "0" the Diff output gain is equal to that of the L+R path.

To review, in Vorsis on-air processors you can increase the apparent stereo separation in an intelligent and program dependent way by:

1. Adjusting the “AGC Backoff” control to settings more negative than “0 dB” which will increase the short term density of both the sum (L+R) and difference (L-R) channels. The average energy of the L-R will be increased even if the Sum and Diff mix levels are set ‘flat’ and the Drive and output Diff controls are set to “0”. The more negative the “AGC Backoff” control is set, the wider the apparent stereo separation will become.
2. Increasing the attack time of the AGC’s without changing the attack times of the compressors. This will create higher ‘instantaneous enhancement’ of the L-R without dramatically increasing its average level.
3. Using the Sum and Diff Mixer Output trims to accomplish a special spectral equalization. Be careful though when setting the high frequency boost in the Diff channel as too much can sound strident on some program material.

Have you ever heard a stereo image that sounded big and warm but without the high end splatter or shimmer that accompanies some methods of stereo enhancement? You can easily achieve this in Vorsis on-air processors by carefully setting some mild boost in the first four bands of the Diff channel!

4. Using the “Diff” controls in the AGC screen to affect either a higher Drive level or higher Mixer output level from the Difference chain, with settings up to +3.0dB being “useful” (settings above +3.0dB may sound ‘unnatural’).
5. Or... a combination of the above.

Generally the best advice for tuning stereo enhancement is to proceed slowly and carefully, making small adjustments at a time until you arrive at your target sound. It’s also advisable to use the Preset Save feature so you can return to a “before” if the last change you made is too much and you’re not sure what you tweaked last!

Wheatstone Corporation © 2008 All Rights Reserved
600 Industrial Drive New Bern, NC 28562 USA 252-638-7000
www.vorsis.com
techsupport@wheatstone.com