Wheatstone
WHITE PAPER

New Findings on FM Stereo Multipath Control

By Jeffrey A. Keith, CPBE, NCE
Senior Product Development Engineer

Wheatstone Corporation
August 2017

Abstract

The use of FM Stereo in the United States was approved in April 1961. Since that time, the technology for generating the complex signal required for FM stereo has matured. Achieving 'lab grade' performance through the entire system from transmitter to consumer receiver is no longer a challenge. Ongoing experimental research conducted by this author over a period of more than thirty years has revealed that lab-grade stereo performance of the broadcast signal chain is not explicitly required in order to deliver a high fidelity listening experience to the FM station's audience. Surprisingly, under certain conditions such high performance may even be a detriment. This paper will explore my ongoing research and will offer possible conclusions worthy of consideration by FM stereo stations concerned with achieving maximum audience reach.

Background

In 1980 while performing on-air tests of a stereo enhancement idea I was developing, an apparent correlation between enhanced FM stereo separation and increased multipath-type interference was observed. It was already assumed that stereo enhancement for FM must be carefully done in order to minimize loudness loss on mono radios. This is because in the FM stereo modulation scheme, as the level of the L-R difference signal is increased the L+R mono must decrease so as to hold total modulation constant. An automatic means for dynamically controlling the maximum amount of stereo enhancement as program conditions changed would appear some years later when commercially available stereo enhancement products became available.

Multipath

In this paper the term 'multipath' interference will be used to describe the audible effects of stereo receivers blending to mono in the absence of clean reception, not the distortion that occurs to even mono reception when multiple delayed signals arrive simultaneously at a receiver's antenna. While the only remedy for the latter is modification of the physical environment, i.e., intervening terrain between transmitter and receiver, it can be shown that the former can be managed via program controlled modification of the total amount of stereo information within stereo program sources.

FM Stereo

In 1961 the FCC selected the Zenith/GE stereo encoding method for transmitting stereo audio and did so mainly because of that system's 100% compatibility with existing FM mono receivers and its need for very little additional bandwidth. In this method the stereo baseband signal, which then later modulates the station's FM carrier, is created in a four step process:

- The left and right channels are added to create an L+R mono sum signal;
- The left and right channels are subtracted to create an L-R stereo difference signal;

- The L-R difference signal modulates a 38 kHz carrier to create a double sideband suppressed carrier, or DSBSC stereo signal;
- The signals representing the L+R, DSBSC, and 19kHz ‘pilot’ which is phase locked to the 38kHz carrier, are added together to form the stereo multiplex baseband signal. This signal is also known in the industry as ‘MPX and/or ‘composite’.

This paper concerns only the modification of the stereo information within the program audio that eventually makes up the L-R stereo difference signal.

Alternate Modulation Methods

Since the time the Double Sideband Suppressed Carrier method went into use other modulation schemes have been invented, proposed or reborn, all which promised to fundamentally improve the signal to noise ratio of the L-R subcarrier. The most recent to resurface was the reintroduction of the idea of using single sideband suppressed carrier, or SSBSC (or SSB) for the L-R subchannel as proposed by Gillman in his 1997 paper [1]. While the generation of a high quality SSBSC signal was an unruly and difficult problem in the analog domain but as pointed out by Tarsio [2] it could be tamed by very careful circuit design. Today the precision available from DSP signal processing makes textbook-perfect SSB much easier.

Three of the most well known broadcast audio processor manufacturers [3]-[4]-[5] have introduced an SSB modulation option in the stereo generators of their high end FM audio processors. Each manufacturer has encouraged end users in the broadcast industry to experiment with the technology on the air to see if it is helpful in their mitigating multipath issues. To date the idea remains promising, however field testing has been unable to prove one way or the other whether SSB does what it was hoped it would do. In my experience, when listeners aren’t aware of which modulation scheme is in use, they found it quite difficult to accurately identify which modulation method was in use.

It has been widely recognized that the SSB modulation method can cause stereo decoding issues with certain consumer receivers; some don’t decode it at all and blend to mono instead. FM stereo receiver tests that I’ve conducted over the past two years showed that approximately 26% could not properly decode SSB. However, two important things must be mentioned about this percentage: first, the 100 or so radios making up the test sample is statistically small. Second, the more radios that were tested the smaller the percentage of non-conformers became. One could take this to mean that if all possible stereo FM radios could be tested for SSB compatibility the percentage of non-conformers would be much smaller.

Stereo Enhancement

In the late 1980’s commercially available stereo enhancers began to enter the broadcast marketplace. The three most popular models were the Modulation Sciences StereoMaxx [6], the Orban 222A [7]-[8], and Stereo Modulation Optimizer [9]. Each of these devices modified the L+R/L-R ratio of the incoming program’s stereo content to raise the contribution of the L-R signal component. Each product was also equipped with an automatic means to manage the maximum amount of enhancement allowed. Interestingly, within each product's printed manuals was a statement that stereo enhancement '…could result in increased multipath'. This was the first time the author recognized that beyond his own field observations there may be a correlation between stereo enhancement and the perception of increased multipath.

The Role of Synchronous AM

This paper won't delve very deeply into myriad transmission side anomalies that can decrease the apparent coverage area of FM stations. However, it is worth noting that FM transmission systems which exhibit high synchronous AM will, by their very nature, have greater difficulty delivering clean reception to the station's audiences.

Because synchronous AM manifests as envelope modulation of the FM carrier, when the carrier amplitude is insufficient to cause full limiting in the FM receiver, the AM component, which is usually highly distorted, will become part of the output audio. The problem can seem even worse in stereo. Though the amplitude of the AM component is synchronous with the station's program audio, its phase may be random and therefore will add and subtract to the demodulated FM audio in destructive ways.

In a purely symmetrical but narrow bandwidth transmission system the synchronous AM component would be pure second harmonic. But purely symmetrical transmission systems are quite rare and the most common synchronous AM components are both even and odd order. It has been observed that when synchronous AM is poor, and in my experience higher than about -40dB, very distorted demodulated audio results wherever the station has either low RF signal levels or high incidences of multipath, or both.

Stereo Separation

In the early days of FM stereo it could be difficult to meet the FCC's minimum stereo separation requirement of 29.7dB at mid audio frequencies. However, FM stereo generator technology, transmitters and antenna systems have evolved to the point where today, FM stations can easily achieve stereo separation well beyond 60dB. Unfortunately achieving a high level of stereo performance on the transmit end seems almost moot when the bulk of stereo consumer receivers rarely achieve 40db to 45dB of stereo separation in the mid audio frequencies.

What isn't readily apparent to the casual observer is that near infinite stereo separation may look impressive on a stereo generator's spec sheet or on the station's audio proof of performance, but as will be shown in this paper extremely high stereo separation isn't a requirement for delivering a very convincing stereo program experience to a station's listeners. The fact that most listeners with consumer-grade receivers don't even realize that they're not hearing all the stereo separation being transmitted by their favorite FM station tends to support this hypothesis.

Stereo vs Mono

At a radio station where I was once employed it had been noted that when early ‘ping-pong’ stereo recordings were being aired the station’s competitive loudness on mono radios was compromised. During such times the station’s loudness was lower by several dB when compared to other stations in the market that were not airing such material at the time.

A series of experiments showed that mono loudness could be improved when such stereo material was being aired, and to the point where it matched the loudness of later-generation stereo programming. This was accomplished by experimentally summing to mono the early ping-pong recordings as they were being aired.

Stereo enhancement was a well understood technique by that time but what wasn’t known, at least to this author, was how much the station’s ‘electrical’ stereo separation could be reduced before it negatively impacted the acoustical listening enjoyment for those with stereo radios.

Listening Tests

In order to determine the minimum electrical stereo separation required to deliver a perceptibly normal stereo experience to the average radio station listener, a series of listening experiments were performed. The test subjects consisted of 36 males and 26 females with an average age of 42; the total sample size was 62. Admittedly a 62 person sample is small and therefore the tests might be seen in a less scientific light. Nonetheless the tests would be educational and might even lend insight into what I had suspected but had never seen documented.

Each subject was instructed to bring a few of their favorite CDs with them; that material would be used as their own familiar ‘reference’. At this point it is important to remember that the only objective of the experiment was to try to determine how much stereo separation each person required in order for them to believe they were hearing normal, natural stereo, and nothing more.

Test Fixture

In order to perform the experiment it was necessary to allow each test subject to vary the amount of stereo separation being heard while at the same time documenting their threshold of stereo perception so it could be accurately duplicated and then carefully measured later. This was best accomplished by building a special test fixture.

The test fixture was constructed to allow each test subject to modify the stereo separation between mono and full stereo using just one un-calibrated knob. The test fixture was inserted between a CD player’s left/right audio outputs and the line inputs of a consumer-grade stereo receiver. Each subject was allowed to listen on headphones or speakers, or both, at their discretion. Figure 1 shows the test setup used for the experiment.

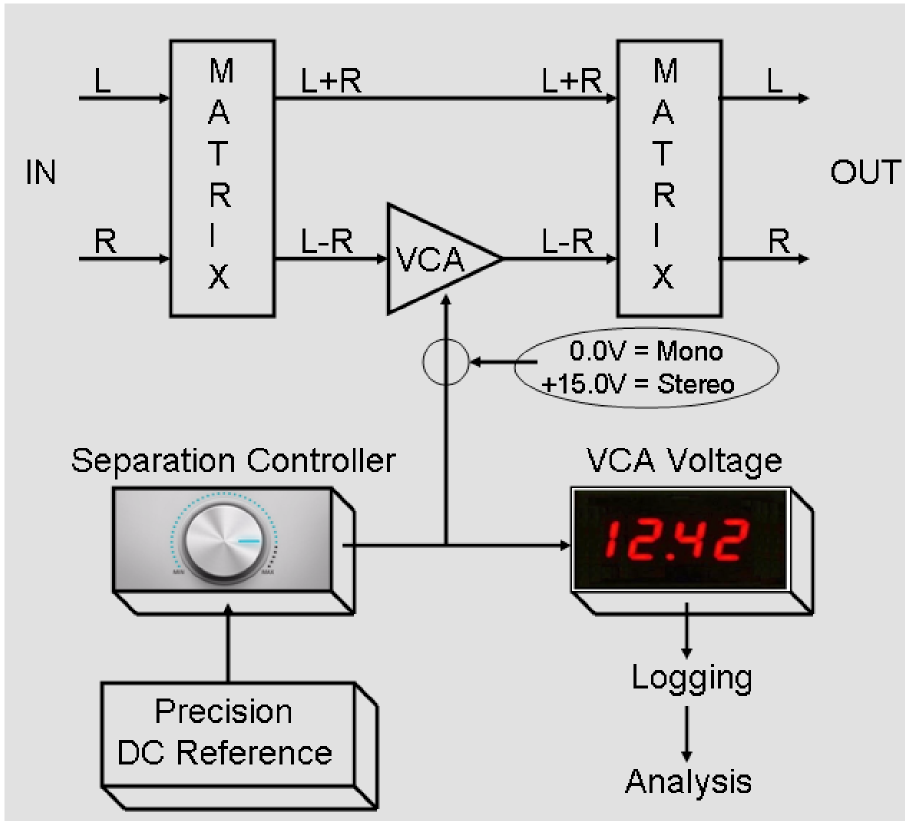

Fig 1 Stereo Separation Test Fixture

The test fixture incorporated a pair of precision stereo matrix circuits built according to a Precision Monolithics SSM Audio Products Application Note [10] to decode the incoming left/right audio into L+R (sum) and L-R (difference) signals. The difference signal path contained a low-distortion voltage controlled amplifier (VCA) whose gain could be set with a DC control voltage scaled by a linear potentiometer; this was the test subject's stereo 'Width' control. The resulting L+R/L-R signals were then reconstructed into normal stereo left and right channels.

The minimum and maximum gains of the VCA were limited by a ranging circuit so that full CCW rotation of the Width control created a pure mono signal and full CW rotation resulted in normal stereo. Each listener was instructed to start with the Width control fully CCW (mono) and stop turning it clockwise when they believed stereo separation was what they thought it should be.

When a test subject found the Width control setting that pleased them the resulting VCA control voltage was measured using a 41/2-digit digital voltmeter and then documented. Once all of the test subjects had completed their listening tests they were dismissed. Later, the results for each of the subjects were duplicated by resetting the Width control to the documented settings and measuring the resulting stereo separation at the test fixture's output.

Test Results

The data shown graphically in Figure 2 reveals that few test subjects required more than about 25dB of stereo separation in order to perceive a stereo listening experience.

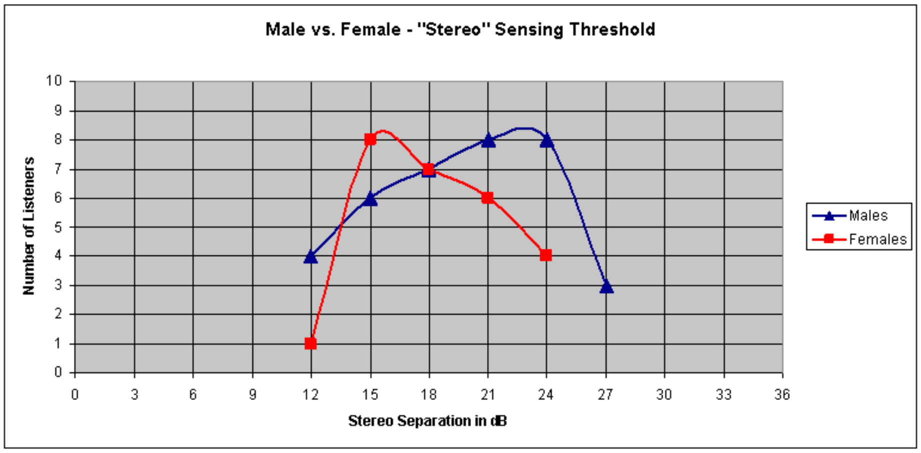

Fig 2 Male/Female Stereo Sensing Thresholds

The curves also show that males tended to prefer slightly more stereo separation than females, and by about 8-10dB. Of note was that the shapes of the male and female preference curves are almost complimentary though what it might mean isn't known.

It is important to note that it is not the curve height that is important in these graphs (because the sample size is both small and not equal in males and females), but where the maximum stereo separation and greatest number of test subjects are clustered.

The preference curve for the 26 females can be interpreted like this; there was 1 subject that felt that around 12dB of stereo separation was normal, 8 subjects needed around 15dB, 7 needed around 18dB, 6 needed about 21dB, and 4 needed around 24dB.

The curve for males is interpreted the same way; 4 male subjects found 12dB of separation to sound 'normal', 6 subjects needed about 15dB, 7 subjects needed 18dB, 8 needed 21dB, 8 needed 24dB and 3 subjects needed 27dB.

The data could be summarized by stating that the male subjects required more stereo separation than females with a difference of about 12dB and that all of them required far less than the >90dB separation available from the test fixture and audio source. A larger number of test subjects may have provided more data and more insight, but at the time it was not possible to arrange for it.

It may have significance that none of the test subjects considered themselves to be an 'audiophile' and instead preferred to be called 'average' radio listeners. Someone with critical listening skills – that is, a true audiophile – would probably have skewed the test results farther towards higher stereo separation, especially if they intuitively understood that full CW rotation of the 'Width' control would result in a full stereo image. Though I work with audio for a living it is difficult for me to tell the difference between 30dB and 90dB of stereo separation with most program material using stereo speakers arranged as they would be in a typical living room environment. In an automobile listening environment where a lot of radio gets listened to, it may be even more suspect.

Supporting Evidence

While searching for more data to either prove or disprove the discovery made about the perception of stereo separation by typical FM radio listeners two papers [12]-[12] were discovered that address the perception of stereo separation in a multichannel sound environment. Figure 3 is a summary of the data from [11] showing how little stereo separation is required for a given amount of perceived stereo image width.

The perception of a stereo listening experience as the amount of inter-channel difference is increased is addressed in both papers and the data appears to correspond well with what was measured during the experiment related to this project. It is apparent from the data that that our perception of the degree of stereo separation and the actual electrical stereo separation are only indirectly related. In fact all the evidence suggested that the electrical stereo separation of typical program content may be reduced to almost any sensible level in order to solve the mono loudness problem and most likely without compromising the stereo listening experience for typical radio station listeners who are not audiophiles.

Stereo Separation in dB	Apparent Stereo Image Width in %
15.2	74
12.8	67
10.9	60
9.4	53
8.2	48
7.1	42
6.2	37
5.4	33
4.7	29
4.1	26

Fig 3 Stereo Separation vs Perceived Image Width

Stereo Width Management

No product existed that could perform the function of 'Stereo Width Management' so the development of a new type of audio processor was undertaken. The goal was simple: manage the overall stereo width of program material in real time while it was being aired in order to even out the station's loudness on mono radios. This was its sole original purpose.

Since the device would be managing the stereo width there was no reason it couldn't be designed to also provide stereo enhancement whenever stereo-shy program material needed it. Because the 'blend' and 'enhance' functions were mathematically related the processor would be able to reduce or increase stereo separation on the fly without listeners being the wiser. Obviously the processor would need to do absolutely nothing to, or in the presence of, mono program content.

In order to accomplish these tasks as invisibly as possible as far as the station's listeners were concerned there was a requirement to reliably detect mono material regardless of the absolute audio levels in the incoming programming. An active tracking stereo detector was designed which could reliably detect the difference between mono and stereo audio over a 60dB range. A simplified block diagram of the detector is shown in Figure 4.

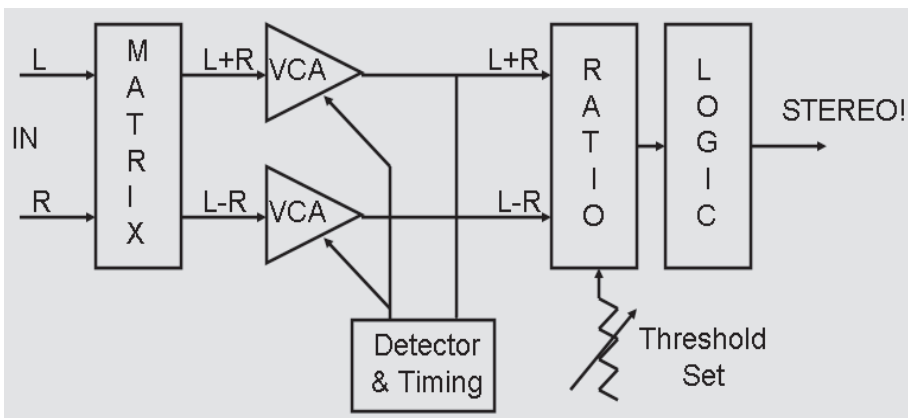

Fig 4 Autoranging Stereo Detector

One feature that makes the stereo detector unique is the fact that the L+R and L-R signals fed into the ratio measuring circuit are simultaneously controlled by the absolute peak level of L+R signal at the output of the stereo matrix. In the tracking arrangement the levels of L+R and L-R, or rather their ratios, are 'servoed' together by the L+R program level. Since in effect this is an AGC, program gating was also provided to ensure that the servo gain would hold at the last current value if the input audio level fell to a certain threshold or disappeared altogether; the detector would not 'hunt' in the absence of program audio.

A similar technique was used to stabilize the L+R and L-R levels into the width measuring circuit and this allows the processor to perfectly measure and then manage the L+R/L-R ratios over a very wide dynamic range of incoming program audio. It should not be necessary to provide an AGC function ahead of the Width processor.

The modification of stereo separation, or L+R/L-R ratio, is not done in the usual manner which is with an L+R/L-R matrix; in this processor such matrices exist only in the stereo detector and width measuring circuits. Instead, a pair of VCA's in a special cross-coupled circuit act in concert to adjust the stereo separation up or down as required. The circuit could be termed a 'virtual' matrix because of how it works, but the audio signal is never placed into in the L+R/L-R domain. While this topology does increase the propagation delay through the unit by a few hundred nanoseconds and increase the noise floor by almost 2dB, it has the benefit of being able to precisely match the phase behavior between the virtual signal paths which, coincidentally, maximizes the 'normal' stereo separation of the processor to essentially that of no processor at all.

It was discovered during on-air testing of the completed prototype that the transition from stereo blend to stereo enhance could sometimes be noticeable due to the time constants required to mask the operation of the processor during its blend function. To remedy this, an asymmetrical time constant slew limiter was designed and incorporated into the VCA control circuit. The slew limiter modifies the transition between blend and enhance in a highly program dependent way, making such transitions virtually inaudible. The timings and slopes of the slew limiter had to be empirically tuned for the most invisible operation because the calculated values were close, but could be improved.

Unexpected Benefits

When the Width Management processor was placed in the air chain of the station and then empirically adjusted for best operation the loudness of both stereo and mono program material on mono radios was now very well matched. The ability of the processor to also enhance otherwise weak stereo separation in a program dependent way was a nice side benefit. Stereo programming still sounded completely normal, even when we noted that the processor's metering showed that it was performing a fairly aggressive reduction in electrical stereo separation.

While the processor cured the mono loudness issues that had plagued the station due to its decades-spanning oldies format, it had a second benefit, one that was completely unexpected; in areas where the station's stereo signal had previously been unpleasant to listen to because of multipath-induced receiver blending, the signal was noticeably cleaner and whenever blending in the receiver did occur it was almost unnoticeable. This was not at all expected.

By listening to the station in different geographical areas and switching the processor between Operate and Bypass we noted that when it was in Bypass the noisy receiver blending returned. In Operate blending was much less noticeable, confirming that the reduction in apparent multipath being observed was due to the processor's management of the amount of stereo being transmitted. Because it worked so well at this task the processor was later renamed from its original 'Mono Compatibility Controller' to 'Automatic Multipath Limiter'.

Conclusions

FM stereo separation does not need to be extremely high or even of laboratory-grade in order for an FM station's listeners to perceive what they believe is a completely normal stereo sound field, in fact quite the contrary. The finding that very low levels of electrical stereo separation can still result in the perception of a normal stereo listening experience was a surprise.

The data gathered using the Stereo Separation Test Fixture was initially brought into question when the task of making the actual separation measurements began and the numbers were surprisingly low. Unfortunately, the original test subjects were not available so I decided to run the tests on myself, eyes closed, to see if the data would be similar. It was.

On-air testing of the prototype of the Width Management processor took place over a period of 10 months and reaffirmed what was discovered in the informal listening tests; it is apparently perfectly acceptable to transmit stereo programming with less than lab-grade stereo separation on FM. In fact, in some scenarios the act of doing so may actually improve the station's subjective coverage.

Additional Research

Additional research and experiments seem appropriate and will therefore be ongoing. A determination will try to be made whether there are other characteristics of the stereo broadcast signal that can be controlled which would help to further reduce the perception of blending on stereo receivers. Simultaneously with my own research, our company and other processor manufacturers will continue exploring the feasibility of using Single Sideband (SSB) instead of Double Sideband (DSB) for the L-R subcarrier. Though as mentioned earlier there is an issue with non-compatibility with a certain percentage of consumer receivers, nonetheless the technology must be explored to see where it leads.

A highly refined DSP version of the original analog 'Mono Compatibility Controller' prototype is available in all but the lowest cost Wheatstone FM audio processors as the 'Automatic Multipath Limiter'. Customers using the feature have reported reductions in multipath induced receiver disturbances as well as extended stereo coverage area. As it is with all technology however, your mileage may vary.

References

[1] Gillman, W., H., “A New Method of Generating FM and Television Stereo Composite Baseband Yields Improved Performance, 1997 NAB Engineering Handbook, pp 325-330.

[2] Tarsio, R., “Response to a Method to Improve Conventional FM Stereo Performance”, January 4, 2011.

[3]Wheatstone Corporation, AirAura FM+HD Digital Spectral Processor.

[4] CRL/Orban, Optimod 8600 FM/8600 FM+HD audio processor.

[5] Omniaaudio, Omnia.11 FM/Omnia.11 FM+HD audio processor.

[6] Modulation Sciences ‘StereoMaxx’.

[7] CRL/Orban, Orban 222A Stereo Spatial Enhancer.

[8] Orban, R., A., US Patent # 4,837,824.

[9] Keith, J., A., Stereo Modulation Optimizer, SMO-900, marketed by Hit Design Ocala, Florida, 1987-1993.

[10] Precision Monolithics Inc., Application Note AN-129, “A Precision Sum and Difference Audio Matrix Circuit”,

SSM Audio Products Handbook, Volume 1, 1990, pp. 196-197.

[11] Stuart, J., R., “The Psychoacoustics of Multichannel Audio”, AES Conference: UK 11th Conference: Audio for New Media (ANM), March 1996, pp 7, Figure 16.

[12] Schroeder, M., R., “Listening with Two Ears”, Music Perception, Vol. 10, No. 3, 1993.