Audio Processing From the Ground Up

by Cornelius Gould

Part 6 – FM Stereo

Cornelius has been discussing the internal workings of a typical audio processor, section by section, from a designer's point of view: the compromises necessary to make an audio processor work, and the balance between legal operation vs. the demands of the programming department. This installment looks at the final stage of processing. [Ed.]

[CLEVELAND, Ohio - June 2003] For Stereo radio stations, the end of the audio path is the Stereo Encoder. We will concentrate mostly on the FM Stereo system, since that is one of the few areas where the stereo encoder is typically "wrapped" inside an audio processor for better performance.

As a kid, I often wondered how a radio station got a "left" channel and a "right" channel into *one* radio signal; it seems this question started me on my path of learning all about the odd world of Audio Processing. I did not actively pursue the answer until years later, when I had my nifty neighborhood FM radio station. Tired of being in mono, I picked up a copy of the Code of Federal Regulations (CFR) at the library to leaf through the FCC Rules for answers. (Yes – I was a young GEEK!)

Anyway, in the CFR I learned about what you have to do to broadcast FM Stereo: There is this thing called the 19 kHz pilot tone that is sent at about 10% modulation on your program signal, as a reference for your radio. When present, it triggers an FM Stereo radio into "FM Stereo mode." I also learned that the "stereo" signal comes in the form of a 38 kHz sub-carrier. So, off to the basement I went to build a 19 kHz oscillator for my first experiment. I mixed it into the program feed to my transmitter and – viola! – the FM stereo light came on in my studio radio. Of course, the audio was still mono.

STEREO AUDIO

The CFR went on to describe the process of encoding audio for FM Stereo. First, a component called L+R feeds the main channel. This L+R signal is created by adding the left channel (L) to the right channel (R), hence L+R. The result was Mono. Easy enough, I had already broadcast that.

Then I read on, and saw talk about L-R, as in Left *minus* Right. Now that was a head-scratcher for me, but eventually I figured out how to subtract the right channel from the left channel. What you get is a signal that contains material *only* appearing on the left or right channel, and nothing from "center." This basically nulls out any mono material, while anything panned even the slightest bit left or right is very noticeable.

As an example, the L-R signal from a typical rock song will usually contain the stereo reverb effects placed on vocals and the drum kits, but no actual vocal material (or any other sound "in the middle"). In essence, this is how a lot of boom boxes perform the "karaoke" function on normal CD's.

The next challenge was to figure out how to make a sub-carrier. This is needed because the L-R audio signal has to be encoded onto a modified AM sub-carrier. (A sub-carrier is basically a radio signal created at a frequency outside our range of hearing, which is still embedded inside of the main radio signal.) The FM Stereo sub-carrier is centered at 38 kHz, compared to the "golden ear" whose hearing tops out at about 20 kHz.

Actually, the 38 kHz sub-carrier is modified in that the AM carrier is removed from the modulation, leaving only the "modulation pieces." Technically speaking, this "Modified AM signal" is called Double Sideband Suppressed Carrier (DSSC). AM signals, such as good 'ol AM broadcasts, have two sidebands whenever audio is being transmitted over the air, one above and one below your main frequency. If you were to look at an AM signal (as you can with a spectrum analyzer) you will see the main carrier, and some "stuff" moving around to either side of it, "growing" up and down with program audio.

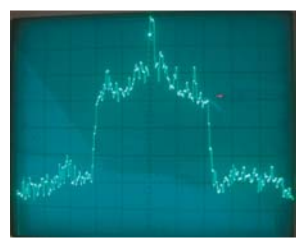

That "stuff" is the two sidebands. (You can also do AM effectively by removing one of the sidebands, but that is an entirely different subject.) When there is no modulation, the AM signal will only consist of the carrier, and *no* sidebands.

With DSSC, the main carrier part is removed, leaving only the two smaller sideband signals. This saves the FM stereo station from having to reduce modulation to fit the carrier in the signal. Another benefit is when there is no L-R audio (whenever there is pure mono audio, e.g. a DJ talking on the air without a music bed behind him), there is no signal to be found in the 38 kHz sub-carrier area. The DSSC sub-carrier method makes more efficient use of the FM signal, and eliminates the need to turn down modulation substantially for FM stereo operation, so FM Stereo stations can basically remain as loud as mono ones.

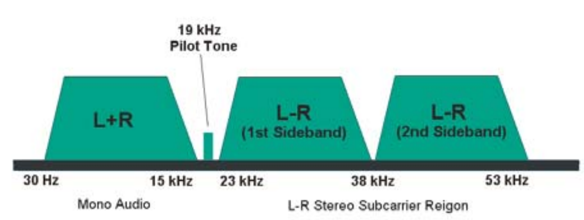

The components that make up FM Stereo.

FOLLOW THE SIGNAL

The 38 kHz DSSC AM sub-carrier is mixed on top of the L+R (mono) signal along with the 19 kHz pilot, and it is all fed as one signal into an FM transmitter, which goes over the air to FM stereo radios all over the city. Now, how does the radio get stereo from this?

In a classic FM Stereo radio, the 19 kHz tone is not only used to turn on the FM stereo indicator, but is also used to synchronize the decoder for the sub-carrier, keeping it in "sync" with your audio processor's stereo encoder. This 19 kHz tone is doubled to create a 38 kHz signal, and the FM stereo decoder uses it to rebuild the main carrier part of the DSSC signal, creating a standard AM signal to properly to recover the L-R audio.

The recovered L-R audio signal is then used to recreate stereo by taking the L-R and adding it to the L+R signal. Remember from the old algebra classes that "L-R" actually means "+L *and* -R", and "L+R" actually means "+L *and* +R."

Thus, when you add "L-R" to "L+R" what happens is the "R's" cancel each other out, and you get "2L," which in English means "all that remains is audio from the left channel," and the sound is sent to your left speaker.

The decoder then flips the "L-R" signal so that it is now "R-L" (+R and -L), and adds this to the "L+R" signal. When this happens, the Left channel signal cancels, and leaves you "2R", or in English "all that remains is audio from the right channel." The right channel audio is then passed on to your right speaker.

In a simplified nutshell, that is how FM stereo works. (It will be obvious shortly why we had to explain all this in such detail!)

The FM Audio processor, as (explained in the previous article) has to apply filtering to the processed program audio to make sure it doesn't exceed much beyond 15 kHz to prevent interference to the 19 kHz pilot tone. If that were to happen, the radio could temporarily "lock onto" the audio material, and momentarily decode "stereo noise" as a result. Now, remember that the 38 kHz sideband energy is also controlled by the 15 kHz filtering, so if the filters are ineffective, or not there, the interference could also come down from the 38 kHz L-R sideband energy too, in effect a "double whammy."

MORE CONTROL

A very popular (and sometimes controversial) processing tool used on the encoded stereo signal is the composite clipper. I alluded to it in an earlier article, and the time has come to talk about it!

What the composite clipper does is to take the L+R signal and its superimposed 38 kHz L-R sub-carrier and feed it into a clipper stage (typically the 19 kHz pilot tone is *not* sent through the clipper – it is either extracted, and added back in after clipping, or added later).

This extra clipping stage allows several more dB's of clipping (read – loudness) than L/R clipping, without the nasty distortion. Well, at least not "normal nasty distortion."

The kind of distortion you get is not very audible at first, as it is distributed about in a not immediately detectable way. To hear this distortion, all you have to do is to remove program audio to, say, the right channel of the processor. If you listen to the (silent) right channel, you will hear "crashing" noises whenever there is high frequency energy in the left channel. As the clipping (loudness) is increased, the crashing gets worse. Plugging the right channel back in covers up the crashing noises with normal right channel audio.

When composite clipping is at its most extreme, you also run the risk of clipping harmonics falling into the 19 kHz pilot region, as well as above the 53 kHz limit of FM stereo. When the pilot area becomes contaminated, the same problem you would get from not using 15 kHz filters happens; with harmonics above 53 kHz, any sub-carriers you may be renting to clients (67 kHz, for example) will be contaminated by this "crashing" noise, too.

This effect constitutes a good part of the controversy over using composite clipping for loudness. Supporters of composite clipping will say much of the audible side effects are negated by the fact that 90% of your listeners are never in a position to have reception good enough to make a difference, but the loudness gained comes through regardless of location. Which stance is the correct one? It depends on what is important to your radio station.

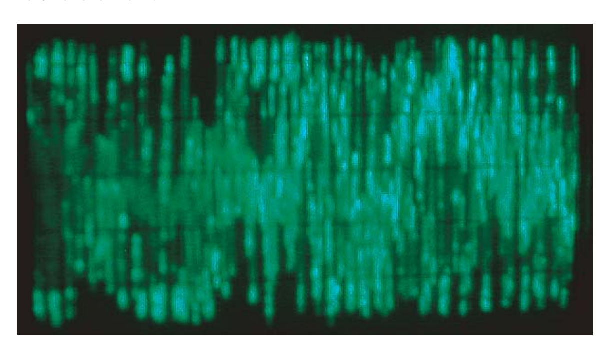

What composite clipping looks like.

As far as the negative effects of composite clipping go, the latest generation of DSP processors all address some of these issues in clever ways, including special filtering to remove most of the spectral "garbage" created by composite clipping. Despite the controversy, we will probably never see composite clipping go away. That is, unless the day comes where radio programmers stop caring about being loud.

Yeah … right!

Cornelius Gould has played with a lot of audio processors, and built some, too! He is the Senior Staff Engineer for Infinity Broadcasting in Cleveland, Ohio as well as Chief Engineer for WJCU 88.7 FM in Cleveland. You can reach him at: cg@radiocleveland.com