IMPROVING YOUR AIR SOUND WITH AUDIO PROCESSING

by Robert Orban
Chief Engineer
Orban Inc.

This presentation is an update of a talk I gave at NAB more than a dozen years ago. Even though we now have digital source media, a digital signal path between studio and transmitter, and all-digital audio processing, the fundamentals haven't really changed all that much. Fortunately, I don't have to talk about how to play those funny black vinyl things on the air any more, because if you still play audio that started out as vinyl it's almost certain that you transferred it to some digital medium in the production room.

To start — I've said this time and time again: garbage in, garbage out. In the digital age, you have to option to feed your processor with a signal that has been through lossy digital compression. If you use this sort of data compression, I advise using it very conservatively because the transmission audio processor can cause the assumptions under which the lossy data compression was designed to become invalid. For example, the psychoacoustic masking of quantization noise in a given frequency band may fail if the channel is passed through an element having non-flat frequency response. Multiband transmission processors are always dynamically re-equalizing the audio, so they can cause this problem. Bottom line: if you are using a system like MPEG Layer 2 or Dolby AC-2, it's a good idea to operate at a data rate of 380 kilobits per second or above. This typically leaves at least 10dB of headroom before quantization noise becomes unmasked, and gives your audio processor some leeway.

Audio processing is only going to improve your air sound if your unprocessed audio is immaculately clean to start with. Digital source media and signal paths have not eliminated the need to watch headroom everywhere in chain, starting with mike preamps, and ending with the analog-to-digital converter feeding your new digital audio processor or STL. Broadcast audio systems love to clip, and will do so given half a chance because of their complexity. Since signal-to-noise ratio in the digital audio path is usually determined by the signal-to-noise ratio of the analog-to-digital converter, and since this converter usually has dynamic range of 105 to 110 at best, it's crucial to align it so that it doesn't clip even when jocks over-drive the board. But don't go turn down the gain further than that or you'll lose dynamic range and bring up noise. Remember that compression will exaggerate any noise in the A-to-D, so that 105dB signal-to-noise ratio is going to end up being effectively much lower by the time your signal hits the air. One of the most common customer service calls we get with digital processors is that "the sound is clean when the source material is quiet, but gets real distorted when things get loud." That one's almost always caused by clipping the input A-to-D converter in the processor.

Keep as many level setting controls as possible out of the hands of DJ's and other non-technical personnel. Use your oscilloscope or peak-reading meter throughout the chain with normal program material to make sure that clipping levels are not even being approached in all analog signal paths before the A-to-D.

The other problem with audio processing that doesn't relate to the processor is the transmitter. FM transmitters and STLs are subject to problems with low-frequency bounce, as Greg Ogonowski has documented over the last several years. In most cases, relatively low-cost modifications are possible that make the transmitter essentially transparent. Composite clipping is done at your own risk because it can cause severe crosstalk between the stereo main-channel and subchannel, and can also wipe out SCAs.

Unfortunately, AM plants are another matter. Particularly when you are broadcasting highly-processed audio, absolutely everything must be done right if you are going to get the most out of the processor. There are major audible differences between AM transmitters, and it is only in the last few years that we have had transmitters that could even begin to be considered transparent in the sense that the word is applied to almost all other audio equipment. In addition to the antenna's being broadband, the transmitter must be down less than 3dB at 0.15 Hertz, have a high frequency audio rolloff that is gentle enough to avoid ringing with a band limited square wave, and a power supply that does not sag or bounce under heavy modulation. Distortion must be very low so that the transmitter doesn't add its own audible distortion to the clipping you are going to have to do in the audio processor. The state of the art in high-powered AM transmitters is about three-tenths percent (0.3%) THD, which sounds *dramatically* better than the typical 2% THD of older designs. Unless you are a transmitter genius and love to do major modifications, most of the older transmitters out there deserve an early retirement.

Without trying to sound like a salesman for the transmitter companies, let me say that you are likely to hear a clearly audible improvement even on low-line auto radios, which, thanks to the work of the Motorola AM stereo people and to IC technology in general, are finally being equipped with reasonably linear detectors.

Even the best of the new-technology transmitters will sound like junk through a narrow band antenna. I can't give you any hard-and-fast rules for when you can consider yourself in trouble, but if the station sounds brighter and/or cleaner into a dummy load than it does into your antenna system, it's time to do some work straightening up the common-point impedance.

Well, let's suppose that your source audio would make an audiophile drool, and your transmitter passes DC through light with 0.001% distortion. At this point the buck-passing stops and we have. to talk about the processor.

With the implementation of DSP-based transmission processing, broadcasters have gained programmability, higher stability over time and temperature, and cleaner sound. In FM, digital processors offer dramatically improved baseband spectral cleanliness, better overshoot control, and the ability to create a brighter sound for a given amount of perceived distortion. Our digital processors provide much better SCA, RDS, and stereo pilot tone protection than did their analog predecessors, and sound substantially crisper.

Before the advent of DSP-based transmission processors, the canniest engineers could take one of the analog boxes, and modify it according to the specific needs of a given format. But such modifications were never recommended by Orban or probably any other manufacturer because we have seen the unsuccessful modifications limp back into our factory service center so that the processor can be restored to stock condition. Now with digital, modifications are not possible because they would involve changing the DSP code, which is orders of magnitude harder than just changing out a few opamps, resistors, or caps. That's why all of the manufacturers of such boxes have built a great deal of flexibility into the box's setup controls. All in all, the programmability of the digital boxes and the ability to daypart processing outweighs any loss of flexibility because you can't do circuit modifications. Also, with fewer and fewer stations having full-time engineers, the digital boxes' simplified setup controls greatly increase the probability that the processor will be appropriately set up at a station with limited engineering resources.

A state-of-the-art processor run aggressively will provide anywhere from 3 to 6dB increase in RMS audio levels by comparison to a processor using the technology of the '60s. However, we have definitely reached the point of diminishing returns. You can only do so much non-linear processing to audio before it becomes seriously objectionable to the ear, and it would seem that the peak-to-average ratios that we are getting out of state-of-the-art processors are getting reasonably close to the psychoacoustic limit of what can be accepted by the ear.

What do I mean specifically? In FM, we are seeing RMS audio levels in the order of 3dB below the RMS level of a sinewave at 100% modulation. In AM, where more pre-emphasis is used, the RMS level is typically down about 4dB. Above that, in either AM or FM, the audio starts to take on very objectionable characteristics, and becomes distorted, dynamically squashed, or, usually, both.

Let's talk about state-of-the-art aggressive processing for a moment. In both AM and FM, it usually starts with a phase rotation circuit, whose primary purpose is to make voice more symmetrical. This way, it can be clipped less hard, and can stay clean while retaining an appropriate level balance against the music. In FM, some people object to the sound of phase rotation. I personally find it to be quite subtle, and certainly preferable to the distortion that happens when you remove the phase rotator and try to achieve the same RMS levels. On voice, removing the phase rotator will typically cost 3dB for the same amount of perceived distortion on some of the more difficult male voices. Only if you can apply phase rotation to all voices prior to final processing should you operate a processor with phase rotation defeated.

There has been a bunch of talk recently about "psychoacoustic processing's" being the next big thing. Fact is, *all* transmission processing is psychoacoustic, because its purpose in life is to reduce the peak-to-average ratio of the audio while fooling the ear into believing that the music is still intact. You should assess any proposed addition to the processing chain with an open mind and skeptical ear, and not allow yourself to be seduced by hype and buzzwords. The fundamentals haven't changed; if you want to keep and audience listening, you've got to process for low listening fatigue. This means that when you evaluate processing, you should listen at great length to a wide variety of program material typical of your format. Only in this way will two crucial factors: listening fatigue, and consistency — become clearly perceived. Instantaneous A/B comparisons of processors are of very limited usefulness because they tell you nothing about either fatigue or source-to-source consistency. A processor which sounds flashy on one piece of program material may well fall apart on the next.

If you are evaluating an AM processor, it's also important to evaluate it on more than one radio. Don't "tune" your radio station to the PD's car — it might to totally atypical of what is in the hands of the audience. In processing, I believe in the "greatest good to the greatest number of listeners" theory. This means that you can't make the processor sound too good on one atypical radio at the expense of sounding bad on many others. You also must be willing to write off the worst of the narrowband ceramic filter radios as hopeless — there's nothing you can do to make them sound good on music.

It perhaps goes without saying that once you are dealing with a transmission system with preemphasis — either AM or FM — the you're going to need a multiband processor. I think that it's safe to say that the single-band processor is obsolete. As Mike Dorrough so aptly put it, the single-band processor is nothing more than a voltage regulator. Its primary problem is that the ear loses sensitivity at low frequencies due to psychoacoustics and the receiver loses sensitivity at high frequencies due to de-emphasis. So if strong low- or high-frequency energy comes along, it will cause gain reduction far in excess of its audible contribution to the mix of sound energy. The ear then hears the loudness of the overall sound as unnaturally reduced, and perceives that a "hole" has been punched in the audio. Even the "smartest" multiple time-constant circuits can't overcome this basic effect caused by the ear's and the receiver's non-constant sensitivity as frequency changes.

If you are trying to do only a moderate amount of processing, then three bands are fine — bass, midrange, and high frequency. However, state-of-the-art aggressive processing requires more bands to get density without audible modulation of one part of the audio spectrum by the others. We use five bands with high-slope 18dB/octave crossovers to get a sort of "automatic equalization" effect which is one of the those surprising things that you philosophically think couldn't possibly work until you actually try it.

It is certainly important to use a slow gain-riding stage in front of an aggressive multi band processor. This seems to have become very common in state-of-the-art processors. If you don't, then operator gain riding becomes exceptionally critical, because these multi-band processors can introduce weird frequency balances if they are seriously overdriven. And if they are underdriven, they don't get loud, and the PD starts breathing down your neck!

What about equalizers in the program line before the processor? So far, the pre-emphasis we have talked about, whether FM or AM, is really to try and equalize the receiver — to flatten out its frequency response. It seems that some bass boost is necessary when you are using a multi-band processor just to keep things dynamically in balance. Orban designs these into all our processors in some way or another.

The question then becomes: "Can the air sound be improved by further equalization?" My answer is no in general, but with some reservations. First, equalizers have their place in the production studio, particularly in correcting the frequency balance of older records. It's amazing how truly lousy the monitor systems in most of the old recording studios were — the producers would literally fly by the seat of their pants, having learned through long experience what sound out of their coaxial monitor was likely to sound good at home, on the radio, or in a jukebox. The sound in the studio usually bore no relationship to any of these other sounds!

The second yes is if your PD has decided that the radio station is to be loud at all costs even if it has to sound like a telephone or honk like a 30 year-old P.A.. to do it! The cheapest, easiest way to make an FM radio station loud is simply to boost the midrange in the 2 to 4 kHz region. In AM, you boost the 1kHz region, because most radios *still* roll off between 2 and 3kHz! Unfortunately, this equalization makes the sound strident and fatiguing, and is likely to shorten average time spent listening. Try to talk the PD out of it -you can always carry the processing wars too far, particularly when everyone has good processors, but everyone keeps turning them up and up!

In audio processing the key to loudness is primarily the design of the peak limiting system. About half of our patents are in this area alone and they mostly involve moving the spectrum of the clipping distortion into the frequency ranges where it is least likely to be audible on a dynamic basis. This involves heavy-duty computer optimization techniques, and the days are long past where a wideband limiter followed by a pair of back-to-back zener diodes was all that it took! It's much easier to design a good-sounding compressor than it is to design a good-sounding peak limiter and this is one thing to a ways be aware of when choosing or designing a processing system.

So to sum up how to improve your air sound by audio processing in 1996: — First, clean it up. If you skip this step, everything else is wasted! — Next, evaluate commercially-available processors. Do it right, which means at length, with many different types of program material typical of your format, and on many different radios. Avoid depending on instantaneous A/B comparisons, because they will mask some very important long-term characteristics of a processor. — This is probably all you need to do. If you're not getting the sound you want, then consider equalizers, exciters, enhancers, etc., with the understanding that all of these gadgets will tend to compromise the basically musicality of your air sound for the sake of novelty or simply not sounding like the next guy down the dial! And *don't* add elements to the processing chain before you've made sure that the source feeding the processor is clean, that the transmitter is transparent, and that your modulation monitor is not overshooting and causing you to set your modulation too conservatively!

Well, that about wraps up the formal part of the talk. I'd be happy to take any questions at this time.