It's Just a Phase

QUIBBLES AND BITS

By Richard Murison

I want to devote a column to Phase Response, and this is actually a good time to take that particular detour. Phase is in essence a time delay. If a system has linear Phase Response, then all frequencies that go into it come out of it at the exact same time. Sound propagation through air displays linear phase response. If I play a musical instrument at one end of a room, when the music travels across to the other end, all of the music's different frequency components arrive in synch with each other. If sound propagation had a non-linear phase response, then the different frequency components of the music would all arrive at different times. [Interestingly, by way of contrast, light traveling through glass has a non-linear phase response, and the different wavelengths of light do in fact emerge from the glass at different times. This is how a prism splits white light into a fan of rainbow colors.]

Phase distortion – applying a non-linear phase change – has some interesting effects. It can change the shape of the waveform without changing its frequency content. We're used to seeing charts showing the frequency response of some audio component or other. We look for flat frequency response and consider it a desirable characteristic to achieve. This is because research has shown that the human ear is sensitive to deviations from a flat frequency response. We perceive such deviations to colour the sound, usually in an unacceptable manner. But if a system has a flat frequency response in combination with a non-linear phase response there is a lot of uncertainty as to whether – and if so, in what way – the resultant phase distortion is audible, even though the result can manifest itself in dramatically visible distortions in the actual waveform itself.

Simplistic experiments, using synthesized waveforms, pretty much always show the same result – that even the most gross phase distortions, accompanied by the most gross distortions of the waveform itself, are apparently inaudible. Phase distortion is therefore given relatively little attention. A body of opinion is emerging, though, which holds that phase distortions make their presence felt audibly not through tonal colorations, but rather through a loss of precision in the ability of a stereo system to create a stable 3-dimensional image. At this point, though, it is fair to say that such notions remain speculative.

However, the notion of time alignment in loudspeaker design is not new, and has been around for decades. If you adjust the precise fore-aft position of a loudspeaker drive unit in its enclosure, you are adjusting the relative times taken for the sounds emitted by each drive unit to reach your ears. This adjustment process is known as time-alignment. What it is actually doing is making a coarse adjustment to the loudspeaker's overall phase response. Normally this is done at the design stage, and by the time the product is manufactured the preferred alignments will have been baked into the design. However, Wilson Audio – some of whose iconic loudspeakers are priced like automobiles – has introduced (on their Mercedes- and Lamborghini- priced models at least) a facility to fine-adjust the time alignment of the midrange and high frequency drivers.

I have twice heard these über-Wilsons being set up, and when the alignment is in the Goldilocks Zone there is absolutely no doubt – the stereo image just 'snaps' magically into place. What is interesting is that if (and this is a big if) the effect we are hearing is solely due to getting the time alignment right, then it enables us to put some specific numbers on the amount of phase distortion that is potentially audible. And those numbers are surprisingly small.

So where does phase distortion come from? Mostly, in electronic equipment, it comes from components or circuits that are frequency-sensitive. Suppose you make an amplifier that contains a low-pass filter. This filter can be there for any number of reasons from stability enhancement, to RF rejection, to noise reduction, to specific sonic tailoring. Generally there will be region where the frequency response is pretty flat (usually the audio band), and a region where the filtering kicks in and the frequency response takes on a desired characteristic. What you tend to find is that the phase response is pretty linear over the flat region, but goes wild wherever the frequency response fluctuates. This kind of behavior is pretty much unavoidable, and is largely explained by various theories with names like Kramers-Kronig, Sohkotsky-Plemelj, and Hilbert Transform. Check 'em out on Wikipedia....

The important takeaway is that if a filter exhibits a dramatic perturbation of its frequency response, then you can count on it introducing a concomitant dramatic perturbation of its phase response. The one causes the other. Now, if the filter's corner frequency – where the frequency response starts to do things – is a long way from the audio band, then there is a good chance that the phase distortions themselves will also be confined to frequencies a long way from the audio band, and the phase response within the audio band may remain close to linear.

But sometimes that's not possible. A good example is the brick-wall anti-aliasing filter that an analog signal is required to pass through before it is digitally sampled. For standard 16/44.1 "Red Book" audio, the audio band finishes at 20kHz, but all frequencies above 22.05kHz need to be scrupulously filtered out. This requires a low-pass anti-aliasing filter that will exhibit serious perturbations in its frequency response [0-100dB in 2.05kHz has a certain kinship with 0-100mph in 2.05 seconds.], which in turn will cause serious perturbations in its phase response. But it will be making those perturbations at frequencies immediately adjacent to the audio band, which will cause the resultant phase response non-linearities to extend down into the audio band itself.

Even if you can persuade yourself that the frequency response of the anti-aliasing filter is sufficiently flat throughout the audio band, you really need to take proper account of any related phase distortion before you can definitively claim that any such a filter is inaudible. This is an area in which I suspect great progress will be made in the coming decade.

Richard Murison enjoyed a long career working with lasers, as a researcher, engineer, and then as an entrepreneur. This enabled him to feed his life-long audiophile habit. Recently, though, he started an audiophile software company, BitPerfect, and consequently he can no longer afford it. Even stranger, therefore, that he has agreed to serve in an unpaid role as a columnist, which he writes from Montreal, Canada.