Subject: Uniform Directivity - How important is it? Posted by Wayne Parham on Mon, 20 May 2013 18:13:45 GMT View Forum Message <> Reply to Message

I must say that I'm in the middle on this subject of directivity and how much it matters. Back when directivity wasn't even considered by most audiophiles, I made more a point of discussing its importance. But now that constant directivity has become more popular in hifi circles, I tend to find myself reminding the new converts not to throw the baby out with the bathwater.

That's why I often refer to speakers with uniform directivity rather than constant directivity. It is more important to me that the power response be uniform than it is for directivity to be constant. They are similar, but not exactly the same.

To illustrate what I mean, consider two loudspeakers, one designed for uniform directivity and another designed solely for flat on-axis response.

The "traditional" loudspeaker designed for flat on-axis response has a woofer that is omnidirectional at LF and crossed-over where convenient. Not that it is a design criteria, but at this point, woofer beamwidth has narrowed, probably in the vicinity of 90°. The tweeter is omnidirectional at crossover and its beamwidth doesn't narrow to 90° until the top octave. Measured response is flat on-axis.

The speaker designed for uniform directivity is crossed-over where woofer and tweeter directivity match. So instead of beamwidth narrowing through the woofer band, then widening as the tweeter is crossed-in, the tweeter directivity matches the directivity of the woofer and remains relatively constant at that point. Alternatively, some designs have the tweeter beamwidth continue to narrow.

Provided there is no abrupt beamwidth jump, I would consider both approaches to be "uniform directivity" designs.

Why does this matter?

Because people listening at at off-axis angles hear smooth response from the speaker that provides uniform directivity, but they don't from the other design. Also, the reverberent field is more natural when the uniform directivity speaker is playing, and that's important because that's the sound that surrounds you. It's as simple as that.

My experience with constant directivity designs started by making Klipschorn-style cornerhorns with constant directivity horns. I learned early on that they did something special where imaging was concerned. The walls confined the midrange and midbass down to the Schroeder frequency, so the sound was uniform throughout the room. Directivity is constant through the entire audio band in a design like this.

There is no way to be outside the pattern in a design like this. It is unique in that regard, and has always been my favorite design approach. But rooms with the right layout to support constant directivity cornerhorns are rare.

Another option presented itself, which is the matched-directivity approach. Physically, it is the same thing as the large Altecs and the JBL 4430. Those were my inspiration for this second design approach. They don't provide constant directivity, because the radiation pattern is omnidirectonal at low frequencies and beamwidth narrows as frequency goes up. But the directivity is at least uniform, and that means off-axis response is still flat. It may have a downward slope, but it is relatively smooth.

Speakers with uniform directivity sound more natural because the reverberent field has spectral balance.

One problem presents itself, and that is horns that provide constant directivity sometimes don't sound all that good. The best example is early CD horns with sharp edges inside. The discontinuities created by the sharp edges cause internal reflections and response anomalies. So I switched back to radial horns early on. These provide nearly constant horizontal beamwidth and gently collapsing vertical beamwidth. That is a very good characteristic for home hifi, in my opinion.

Prosound horns have a different set of priorities where constant directivity is concerned. Prosound horns not only need high-efficiency and constant beamwidth, but they must also be arrayable. This means they have to concern themselves with things like astigmatism, pattern flip and waistbanding. As an example of why those things are important, consider an installation where horns are splayed. In this arrangement, the primary lobe of one horn gets interference from the secondary lobe of the adjacent horn. So in this case, the characteristics of the sound radiated outside the pattern is as important as the sound radiated within the pattern.

A prosound horn design will allow some response ripple and other anomalous behavior within the pattern for "good behavior" outside the pattern. This makes sense, because horns in a prosound environment will be arrayed, so it doesn't make sense to optimize the horn for single use.

But there are a different set of priorities for a horn/waveguide designed for studio monitors or for home hifi use.

For home hifi, we want the response in the pattern to be as smooth and clean as possible, and ideally we want the beamwidth to be as constant as possible too. But where a trade-off must be struck, it doesn't make sense to optimize the response at the extreme edge of the pattern or outside, like we might choose to do in a prosound horn. The home hifi horn will not be arrayed. So the best approach is to make the radiation pattern uniform and to optimize response within the pattern.

This brings me back to what I said in the first paragraph about some people throwing the baby out with the bathwater. I'm all for uniform directivity, and found myself regularly arguing its virtues over the years. But it seems like lately, I see guys posting polars and sonograms, looking for a holy grail in its virtues. Some have gone way too far with that, in my opinion, and are using prosound techniques to build home hifi speakers. Their polars look wonderful, but their response curves and distortion performance is only so-so.

Compare the two curves on the chart below and you'll see what I mean. These are two

waveguides that are approximately the same size, designed to be used over the same passband. One is optimized to provide smooth response, the other to provide constant beamwidth. The one with smoother response has about 2dB ripple, but directivity narrows slightly below 2kHz, before pattern control is completely lost around 1kHz. It is also 3dB louder with the same input signal, so drive voltage requirements are reduced, lowering distortion. The other waveguide has about 5dB ripple, but is able to maintain beamwidth down to its lower cutoff around 1kHz.

I've measured several different horns over the years, from radials to tractrix, and a lot of them are captured in the link below. Look through the list and you can see measurements of various types of horns and waveguides:

Measured DatasetsRemember the examples above, the traditional loudspeaker with the 90°-to-omni shift compared with the speaker having uniform directivity? In this case, when we compare response at 45° off axis, we see the traditional speaker has a 6dB dip at the crossover point, because the 90° beamwidth is defined by its -6dB point. The speaker with uniform directivity has no dip, its off-axis curve is a straight diagonal line.

But what about a speaker with a smidge of waistbanding? This is something to avoid in the prosound world, because a horn creates a secondary lobe in the waistbanding region, and this secondary lobe interferes with the primary lobe of another horn in an array. Not so in a home hifi setup, because there is no other other horn to interfere with. So what other consequence do we see from wasitbanding?

Waistbanding is a "pinch" of the pattern at low frequencies. An example is a horn/waveguide that provides constant 90° beamwidth from say 2kHz up, but that narrows to 70° below that, before ultimately opening wide up approaching omnidirectional radiation at low frequencies below 1kHz where pattern control is completely lost. A sonogram will show this "pinch" in the 1kHz to 2kHz region.

What does it mean in a home hifi environment? Practically nothing, it's inaudible. What is really happening is the sound at 45° is reduced slightly between 1kHz and 2kHz, and by slightly, I mean about 2dB. It is minor, nothing like the 6dB dip off-axis of a traditional loudspeaker designed solely for on-axis response.

Does that mean waistbanding is completely insignificant? Of course not. Would we want to design to reduce it? Certainly, provided there are no other trade-offs. But there are trade-offs, there are always trade-offs. The most common mechanism to counter waistbanding is a secondary flare, and this increases horn size, which increases center-to-center distance, which in turn brings the vertical nulls closer, limiting the size of the forward lobe. Or worse, if the mouth size must remain constant, then adding a secondary flare requires shortening the main body of the horn, modifying the flare profile and possibly truncating it. This introduces extra ripple. That may be worthwhile in a prosound horn, but it probably isn't in a waveguide designed to be used in a studio monitor or home hifi speaker.

Which leads me back to the radial horns. I personally would rather have a radial horn that provided constant directivity in the horizontal, gently collapsing directivity in the vertical and

smooth response in the pattern than I would a so-called "waveguide" that had peaky response. I've seen some out there that have 5dB ripple, and that's about twice what I would be willing to live with.

A good hifi horn offers response flat within a 2dB window, and a good studio monitor waveguide is able to do this too. The waveguide usually cannot be used to as low frequency as a similarly sized (exponential) radial horn, because the waveguide's acoustic loading isn't as good at low frequencies. It has this in common with a tractrix horn, which also loads relatively poorly down low. But it need not increase response ripple like prosound horns do, otherwise it has thrown the baby out with the bathwater.

Subject: Re: Uniform Directivity - How important is it? Posted by Wayne Parham on Mon, 20 May 2013 22:08:18 GMT View Forum Message <> Reply to Message

To offer a little more information, consider the following sonograms:

First, the ideal directivity pattern would be constant through the entire audio band. A constant directivity cornerhorn comes close to this, because the walls confine the beam down low, and the midhorn and tweeter waveguides set the pattern up high. So basically, the whole room is the "sweet spot."

Of course, that's an idealized sonogram, and in the real world, room modes will break up that nice pretty picture and make pockets of hot and dead spots throughout the room below the Schroeder frequency, around 200Hz or so. That's what flanking subs and distributed multisubs seek to mitigate.

But that's a whole different subject. For now, back to tweeter waveguides.

Now the opposite end of the scale, a sonogram that is audibly deficient. This shows approximately 90° beamwidth up to 1.6kHz, then widening to over 120° to 6kHz, then narrowing to 40° above that.

The problem is the response off-axis. A listener that's sitting on-axis gets a different sonic presentation than a person sitting at 30° off-axis, and another person at 45° gets another completely different presentation. There is no spectral balance for listeners off-axis more than about 20°.

Look at the legend to the right. See the colors, and how they relate to SPL. At 45°, the sound is -6dB compared to the on-axis level below 1.6kHz. It rises nearly to the on-axis SPL, maybe -2dB in the 1.6kHz to 6kHz region, and then falls rapidly above that, being approximately -12dB at 12kHz.

Said another way, imagine the response curve for a listener sitting 40° off-axis. For that person, the response blooms about 4dB from 1.6kHz to 6kHz, and then it falls rapidly down to about -6dB by 12kHz.

This is non-uniform directivity, and it is something I would personally not want.

But now let's look at some popular horn/waveguides that provide good directivity.

This is a horn/waveguide that provide near-ideal beamwidth. As you can see, the sonogram is reasonably flat. It stays locked-on at 45° beamwidth across the band.

But what if there are design features that cause this horn to have 5dB ripple on-axis, as well as off-axis? In a way, it doesn't matter that the polars are nice and pretty, because there is a lot of ripple no matter where you sit.

Here are some sonograms of horn/waveguides that show some waistbanding, but that are still very good.

In each case, you will notice the beamwidth is 45° from midband up. But down low, there is a little bit of a pinch in the beamwidth.

Examine each chart closely, and look at the legend. Remember that the beamwidth is defined as the angle where response is -6dB down from the on-axis level. So since each of these waveguides is a 90° device, and since each shows some waistbanding down low, look and see what the SPL is in the "pinched" region. What you will notice is that instead of being -6dB, it is a little more, like -8dB from the on-axis level. What you are actually seeing in these charts are devices that have about 2dB less output at 45° off-axis in the waistbanding region.

In some cases, this may be in the crossover overlap region, where the woofer and tweeter directivities blend. If so, the waistbanding may be masked. But even if it isn't, or if it is only partially blended, we're still talking about a relatively minor anomaly. It isn't as though waistbanding is entirely trivial, but it is more important in prosound applications, where arrayability is important. For high-fidelity monitoring applications, I would prefer the waveguide have smooth response through the pattern as its primary design goal.

In the end, the matched-directivity design will look something like this, with woofer and tweeter blended:

I just can't help assuming that in your last sentence, you are tooting your own horn. Pun intended.

Subject: Re: Uniform Directivity - How important is it? Posted by Wayne Parham on Sat, 25 May 2013 15:54:41 GMT View Forum Message <> Reply to Message

Yes, of course!

The H290C is the ultimate waveguide for critical studio monitoring, home hifi or home theater applications. Its response curve is as smooth as any horn I've seen, efficiency is high, distortion is low and directivity is constant. I designed it with features that optimize wavefront propogation and smooth response, and the only thing I had to give up in the trade is a smidge of beamwidth narrowing below 2kHz.

And this 1kHz to 2kHz region is blended with the adjacent sound source anyway, the midwoofer in a matched-directivity two-way speaker or the midhorn in a constant directivity cornerhorn. That's where the directivities blend, in the horizontal they're mostly contructive in the pattern, so the sources beamwidths tend to average together, with a little bit of ripple at the edge. In the vertical, they're constructive to the edge of the forward lobe, where the vertical nulls cut in. These features modify the beamwidth of both sources in the 1kHz to 2kHz crossover overlap region.

So the H290C is the perfect set of optimizations that gives up nothing for studio monitoring and high-fidelity applications. It provides high efficiency, smooth response, low distortion and uniform directivity. I'm very proud of it.

Subject: Re: Uniform Directivity - How important is it? Posted by Wayne Parham on Wed, 12 Jun 2013 15:16:53 GMT View Forum Message <> Reply to Message

Based on my observations over the years, I have made three basic assumptions that largely drive my design choices:

My first working assumption is that speakers with constant beamwidth always sound better than speakers that don't, provided everything else is equal.

A second assumption is that speakers with non-constant beamwidth but uniform-directivity - those having just gradual change - sound better than speakers with directivity that shifts radically

somewhere in the passband. This is especially true if the directivity shift happens in the peak of the Fletcher-Munson curve where we are most sensitive.

When I say "directivity shift", I'm not talking about 20%, by the way. I'm talking about at least 50%. Beamwidth that stays constant within 20% is perfectly fine, certainly, at least for home theater or hifi. It's the shift from 90° to 180° in less than an octave that you'll hear, not the little bulge of 10° or 20° in a pattern that averages 80° or 90°. A 20% shift amounts to about 2dB at the very edge of the pattern, which is completely inaudible. But a 90° to 180° shift is 6dB at the edge of the pattern, which is most definitely audible. A shift like that screws up the spectral balance in the reverberent field.

So speakers with collapsing directivity (like DI-matched two-ways) can sound very nice provided the directivity change is smooth and gradual. The DI-matched two-way approach is a worthwhile compromise where constant directivity is impossible or impractical.

My third assumption has to do with the "provided everything else is equal" part. This assumption is where sound radiators are concerned, those with truer (flatter) amplitude response sound better than those with peaks and dips. This is true not only of direct radiators but also of horns.

Waveguides offer the promise of smoother response than constant directivity horns, but at the expense of slightly less pattern control. For example, without a diffraction slot in the throat, they cannot maintain beamwidth in the top-octave, narrowing instead to the compression driver exit angle. They sometimes waistband a little at the bottom end of their range too, depending on the shape of the flare nearest the mouth. But in general, a waveguide provides constant directivity and also provides much smoother response than a constant directivity horn.

Waveguides are designed to provide smooth wavefront propogation. The wave, where it contacts the waveguide, is always perpendicular to the surface of the flare. This provides a nice, clean spherical section as the wavefront exits the mouth. It makes them act something like tractrix or LeCleach flares, but with nearly constant directivity. But different shapes and flare profiles offer different optimizations, and correspondingly different performance metrics. Some geometries provide smoother response than others.

An argument can be made that as long as response ripple is constant across all axes, then it can be equalized flat. The idea is that if directivity is constant, then the power response is the same shape as the on-axis response, so equalization in one plane is appropriate to all planes. I think there is merit in that argument, but I do not agree that just because a horn is equalized flat, it will sound as good.

There is a big difference between equalizing for mass-rolloff and using a series of tank circuits to tame response ripple. The conjugate filter for mass-rolloff is a simple single-pole high-pass, and is not a resonant condition. That is quite different than the conditions that cause ripple, and I have not found any cases where the underlying mechanisms that create this ripple come without additional penalty. Sound quality suffers.

You can always take a constant directivity horn and EQ out the ripple. Take the JBL 2370 and 2380 horns, for example, which exhibit 5dB peaks in the passband. They can be equalized flat, but even so, those kinds of horns still sound harsh. And we have seen that some waveguides

also generate a peaky response chart, such as the SEOS family of devices we compared with earlier. The SEOS device produces some ripple, a result of its geometry. Here's some discussion about it, where Bill Waslo claims it to be audible, but suggests a way to correct for it in the crossover, using multiple notch filters:

Bill Waslo wrote on Sat, 3 March 2012 21:12lt's those Inductor-Capacitor-Resistor (LCR) strings that go across the CD driver. Can't just take them out, other stuff would need to be adjusted to compensate or it would sound awful. I did run some designs without those LCR, but really think they should be left in. One of the bumps they deal with is at 2kHz, which is a terrible frequency to have a bump at (near where ears are most sensitive). Why go to all the trouble of waveguides and CDs and then cheap out on a few components?Bill Waslo comments on the audibility of the ripple inherent in the SEOS deviceConsider that 5dB represents a 3x increase in power. Equalization requires a significant power shift - To remove a 5dB peak means the power is cut 3x at the peak, which also means that it must be raised in comparison by 3x everywhere else. This also means excursion is increased and everything else that goes with it. That is not the only issue, in fact, it may not even be the most significant issue. But whatever it is, there can be little doubt that a constant directivity horn is nowhere near as smooth sounding as a properly designed waveguide.

I have said many times before, I even prefer a good radial horn to many constant directivity horns, purely because of their sound quality. I can remember so many discussions over the years with tractrix horn guys, many that use a simple first-order capacitor and nothing else. They trade everything to get smooth response - out goes directivity, power response, excursion at the low end, etc. And when I say "out goes directivity" I don't just mean the horizontals, but even more so the verticals, because with a single cap, the forward lobe becomes a paper-thin strata. But still, they love the pure sound they get in that one pinpoint spot.

What I like about a good waveguide is we can achieve this kind of smoothness, and still provide nearly constant directivity. It really is a design approach that has one foot in the constant directivity world and the other in the audiophile response purity world. Of course, there is a continuum of optimizations one can choose, spanning between those two worlds. The waveguide can be more constant directivity or more smooth, or somewhere halfway in between.

Which brings me back to the first observation/assumption, that speakers with constant beamwidth always sound better than speakers that don't, provided everything else is equal. Or more precisely to the qualification, the "provided everything else is equal" part. It's why I said earlier that a good waveguide should not "throw the baby out with the bathwater." We do not want to use a waveguide that is excessively peaky in a high-fidelity loudspeaker. I'd say a worthwhile criteria is no more than 3dB variance in an octave, i.e. +/-1.5dB. Above that, and the response ripple becomes audible. Why settle for audible ripple when you don't have to?