# Subject: Notes for the DIYer
Posted by Wayne Parham on Tue, 01 Nov 2011 22:48:44 GMT
View Forum Message <> Reply to Message

loudspeakers. So I'm getting lots of PMs and posts asking what I'd do about this and what I think about that. Most seem to be concerned with the crossover, things like what frequencies and slopes to choose and whether or not I think passive is as good as active. I've begun copy-and-pasting the same replies in my emails, which now I'll turn into a forum post.

Crossover frequency and slope
Some people think going with sharper slopes helps reduce off-axis nulls. The idea is that a brick-wall filter has less overlap, which is true. The problem with this philosophy is that a sharp slope also has with it a sharp phase change, and that is a problem. It usually makes it harder to blend. What is nice about the slopes from about second-order to fourth-order in this frequency range is they tend to have a medium/slow phase change in the crossover band. This forms a natural delay, which of course, changes slightly with frequency. But through the fairly small overlap region, it doesn't shift much and so can be used to match the fore-aft difference in the positions of the acoustic centers. This to me is far the best approach. Whether done passively, actively, analog or digital - I don't care - but the second-order to fourth-order slope seems to work best in this range, and needs no other delay for horns of this size. It all fits together.

I've tried slopes much higher, sixth, eighth even tenth-order slopes. The measurements don't look as good, and the speaker doesn't sound as natural. On the other hand, first-order is always too shallow. I've never seen a CD speaker that worked well with first-order. It is great for on-axis, but terrible off-axis. Even second-order is usually too shallow, but I have found it good in a few speakers, because what really matters is the acoustic slope anyway. That's usually one to two orders higher than the electrical filter.

All in all, I almost always find a second-order to fourth-order filter works best, and it is almost never symmetrical.
Crossover optimization for DI-matched two-way speakers
Crossover optimization for DI-matched two-way speakers, revisited
Constant directivity, compression drivers and crossovers
Speaker motors and passive crossover filters
Crossover Electronics 101Box tuning and construction
In a box the size needed for a matched-directivity speaker, a couple things happen that are fundamentally different than a bookshelf speaker or a mini-monitor on a stand. It isn't just the directivity in the tweeter range that's different, but also the behavior at midrange frequencies, below 300Hz.

When your cabinet is bigger than about three cubic feet, internal standing waves line up that are too long for acoustic insulation lining the walls. The damping material becomes ineffective below about 300Hz. This is fine for a mini-monitor or a bookshelf speaker smaller than a couple cubic feet, because standing waves are higher in frequency, and the damping material effectively attenuates those higher-frequency modes. But in a larger speaker, it generally can't. So don't just assume that the nice smooth curve shown in the simulation software will be realized in the physical model.  Chances are, you'll measure a blip or two in the midrange.  A couple ways to mitigate this are to use a sheet of insulation spanning the cross-section of the cabinet, laying across a brace or something.  Midwoofer and port position are also important, because if either lies in a pressure node, it can make matters worse.

Also, don't assume that a trapezoidal cabinet will help.  It won't.  The strongest modes are axial, and it only takes one.  Measure the cabinet to make sure it doesn't have excessive ripple in the lower midrange, from about 100Hz to 300Hz.  If it does, move the woofer position and/or port.

Baffle-step and room modes

Another thing that comes into play in this frequency range is the transition between omnidirectional radiation and half-space radiation because of baffle dimensions.  In a small speaker, that transition can happen in the upper midrange or overtone region, several hundred Hertz.  In this case, it might make sense to provide on-axis equalization in the form of a baffle-step compensation filter.  But in a cabinet the size of a DI-matched two-way, the baffle transition usually happens lower than that, often times below 200Hz.  When this occurs, baffle-step compensation is ill-advised because it would put more power into the room modes.  What's worse, these kinds of speakers are often used on stands which make floor bounce, rear-wall bounce and vertical modes all conspire to create sharp peaks and nulls in the same region.  For this reason, I would strongly discourage the use of BSC and would encourage the use of flanking subs instead.

Baffle step is caused by directivity change, after all.  Since these speakers are designed to provide uniform directivity, it hardly makes sense to equalize the on-axis response to correct for a beamwidth shift.  That sort of thinking seems counter to the goals of uniform directivity, making the power response worse in an attempt to improve the on-axis response.  It makes more sense to solve the problem acoustically, when possible.

Baffle Step

Room modes, multisubs and flanking subs

Helper Woofer LocationHorn/waveguide size and directivity

A common misconception has sprung up among many new CD / waveguide enthusiasts.  Some think it is important to use a horn that has pattern control well below the passband.  It seems to be borrowed from the idea that a driver should have good acoustic response well outside the crossover band.  These things just aren't true.

At least the idea that driver response be good outside the crossover band has some merit, in that you don't want to push a driver too far.  But even here it is a slight oversimplification that borders on being mistaken.

Every driver is a passband device, so what portion of the region isn't used simply limits the usable range.  This is particularly true of horns, which are decade devices.  That's only three octaves.  Tweeters can be pushed to four octaves sometimes, but when you get past that, you're looking at increased distortion, breakup, IMD, etc.  So it is unreasonable to limit a device to an octave on each side of its passband;  Even more than unreasonable, it is unwise.  It is not good practice, and will lead to substandard performance.  Too many crossover points, too many devices, too much wasted bandwidth.

Still, you do need some distance between the passband and the absolute maximum ratings of the device, otherwise it will be strained.  A half octave is generally plenty though, and often just 10%

to 20% or so of the range is a usable limit. So if a tweeter response starts to droop at say 1kHz, then maybe 1.2kHz should be its lower limit. You definitely don't have to wait until 2kHz for crossover, unless of course you're running first-order. A midwoofer breakup peak gets bad starting at 1.2kHz, maybe run only to 1kHz. Going all the way down to 600Hz is probably lower than needed. This depends largely on the slope - you just don't want to excite that cone near its breakup frequency. It's not hard to verify - the proof is in the pudding - if you go too far, measurements will show an anomaly. It will have excessive response ripple or distortion or both. Those kinds of metrics should drive the decision, not an arbitrary rule of thumb.

This is even more true of pattern control. Much more true, in my opinion. You absolutely do not want a horn that is larger than required. The ability to provide pattern control down low is strictly a function of size, so to maintain pattern control below the passband means the horn is larger than needed. This is not only a waste, it also ensures horrible polars because the device will be too big. Integration with other devices is adversely affected. One should use a horn just large enough to control the pattern down low, and no larger.

Also, in a related note, the metric to be concerned with here is the horizontal, not the vertical. If you size for pattern control in the vertical, the horn will be way too large. Waveguides with 90°x40° to 90°x60° aspect ratio should be sized to provide horizontal pattern control just below the passband, no further. The vertical pattern will begin to widen (get taller) above that point, but the vertical nulls will cut into it. So the vertical pattern of the waveguide in the octave above crossover is less important than the pattern higher up. What's most important in the vertical is that it limit the pattern at HF (like in the top two octaves) and that when it does start to open up, that it does so gradually and without abrupt change. Abrupt change in directivity manifests itself in bad response, both amplitude and time response. So we want smoothness, and we want reasonable size. That is best overall.

Consider what you are asking of the horn/waveguide. It should provide directivity control in the horizontal that is constant through its passband, all the way down to the crossover point. It should be matched to the directivity of the lower frequency device so there is no deviation in off-axis response at any horizontal angle. Then also, it is desirable to limit HF at large vertical angles, because floor and ceiling reflections are unwanted, especially at high frequencies. But it should not be so tall that vertical nulls cut into the forward lobe, making it too thin. This is more important than the vertical pattern control down low, since the nulls will cut into the pattern anyway.

It may make sense to choose a horn that provides horizontal pattern control slightly lower than needed, in order to gain some extra vertical pattern control. In other words, choosing a horn that is a little bit oversized to get vertical control closer to the crossover point may make sense. But don't go crazy with that. Use a good sense of balance with the respect to the competing priorities of null angles and vertical pattern control. Carefully select a horn sized appropriately for the drivers and crossover frequency.

It's a fine balance to strike, choosing between vertical pattern control and positioning of the nulls. What improves one worsens the other. One thing you definitely don't want to do is to get a horn that is large enough to hold a tight vertical beamwidth only to have the null angle so small it cuts into it. It would be better to have the horn begin to lose control, because the null angle sets the patten in the crossover region anyway.

The larger horns look great when measured by themselves, pattern control can be maintained way down. But that matters very little since the integration with other devices changes the landscape entirely. Polars of a single CD horn are easy to make look good, they maintain constant beamwidth in the passband, then dip slightly at the bottom end and open wide up. If the horn were used alone, then by all means, go large. The bigger the better.

But again, when used in a loudspeaker system, the interaction between components sets the vertical pattern in the crossover region much more than the top and bottom wall angles of the waveguide. The wall angle sets the pattern at HF, but the driver interaction is primarily responsible for the verticals near crossover. This should not be overlooked by anyone that wants to design a good sounding speaker.
Matching directivity in the vertical and the horizontal planes
High-Fidelity Uniform-Directivity Loudspeakers

---

Subject: Re: Notes for the DIYer
Posted by skywave-rider on Wed, 02 Nov 2011 03:56:07 GMT
View Forum Message <> Reply to Message

Bookmarked.

---

Subject: Re: Notes for the DIYer
Posted by davewantsmoore on Mon, 07 Nov 2011 00:07:19 GMT
View Forum Message <> Reply to Message

Wayne. Your time and knowledge is very much appreciated.

---

Subject: Re: Notes for the DIYer
Posted by Aki on Mon, 07 Nov 2011 01:17:27 GMT
View Forum Message <> Reply to Message

+1

Not knowing much about the topic, this is quite enlightening.

---

Subject: Re: Notes for the DIYer
Posted by Wayne Parham on Tue, 29 Nov 2011 20:35:23 GMT
View Forum Message <> Reply to Message

As most of you know, I've been building CD/waveguide speakers like these for over 30 years. My first models all used a simple crossover, and we continued to build them that way through most of

the 1980s. I always have tended to embrace Occam's Razor when designing circuits, especially loudspeaker crossovers. A couple early minimalist crossovers are described in my crossover document, as are other improved versions that followed it, including the current design:
Speaker motors and passive crossover filtersThe "Crossover Electronics 101" lab handout is also probably useful, in this regard. It is a related document, one I use at trade show seminars to help illustrate the principles used when designing crossovers for loudspeakers like these:
Crossover Electronics 101As I said, my first crossovers were minimalist designs and they remained that way for years. I was satisfied with them, not noticing their deficiencies for quite some time. You can make a good crossover with just a handful of parts, but I have learned that you can't make the best crossover for a waveguide speaker with just a handful of parts. There are a few reasons for this, which I will try to describe.

While I wouldn't suggest a minimalist crossover for a waveguide, I still would suggest a simple crossover for some speakers, like cone/dome monitors where the tweeter is designed with sufficient travel to allow relatively low frequency energy without stress. They don't have issues with phase plug interference, so as long as the suspension allows and the voice coil stays in the gap, there is little problem. A minimalist crossover works great for that kind of speaker.

Compression drivers on horn/waveguides are a different matter entirely. They don't respond well to minimalist crossovers for a variety of reasons. For one thing, the low-orders allow excessive excursion, which causes a compression driver to distort. The typical 1" exit compression driver is designed to be used with a ~1.5kHz crossover using third-order or higher slope. They just can't handle much low frequency energy at all. It isn't unusual to see a diaphragm contact the phase plug when driven hard, and it's real easy to drive 'em too hard with a shallow crossover slope. And even long before you see this kind of catastrophic damage, you'll see harmonic and intermodulation distortion rapidly rise.

It isn't like a dome tweeter either, where distortion rises proportional to excursion. It's more like a shelf - above horn cutoff, excursion is low so distortion is practically non-existent but below cutoff, excursion rises rapidly and distortion shoots straight up. Horn loading limits excursion above cutoff but below, the diaphragm is free to move. Sometimes is becomes dangerously underdamped, right as the horn transitions at cutoff, especially in conical horns and other CD designs.

With low-order crossover slopes, the amount of energy down low can be significant, so there is little to prevent the driver from over-excursion. Especially since the impedance spikes right around cutoff - a quarter-wave mode - which tends to leave the driver even more vulnerable just below that spike. You really don't want much energy presented to the driver below that point, because it usually makes intermittent ugly sounds and can be damaging.

This potential distortion problem is one of two things that conspire to make minimalist crossovers generally unacceptable for horn/waveguide speakers. The second thing is the interaction between drivers in the crossover region. This is one of the biggest challenges for the CD/waveguide loudspeaker designer - to develop a clean forward lobe that is not too thin in height and to limit the amount of sound outside the pattern. With this, there is the desire to limit the anomalous sound outside the pattern and I think that's where the designer's burden really lies. This is what separates the good speakers from the really great ones. It is the difference between a decent DIY project and a world-class loudspeaker. It's all in the shape and size of the forward lobe, and in the absence of nasties outside it.

Crossovers with low-order slopes have wider overlap bands than those with higher-order slopes. Where the typical fourth-order crossover will have adjacent drivers blended for about a quarter-octave, first and second-order slopes will blend for an octave or more. The nulls only form in the blended overlap region, so narrower overlap means narrower nulls. I have found that there seems to be a "sweet spot" between (electrical) second and fourth order, because it allows enough blending to provide a smooth phase transition for good summing but also not so much that it is excessively anomalous above and below the crossover point.

I wouldn't care if the nulls were wider as a result of low-order slopes. If that's all that happened, I might prefer lower-order slopes. It could be used to provide extra "directivity control" in the vertical down where the horn was too small to set the coverage angle. That's one of my design philosophies anyway - to set the nulls just outside the vertical pattern at HF. It sort of punctuates the pattern. If the nulls could be made an octave wide, then they could serve this purpose even better.
Matching directivity in the vertical and the horizontal planesThe traditional conical horn with 1:2 aspect ratio is twice as wide as it is tall, with horizontal angle that's about twice the vertical angle. This results in a horn that cannot provide pattern control in the vertical for almost two octaves above that which it provides control in the horizontal. So you can choose a horn that's larger than needed to set the horizontal pattern in order to get the vertical control to work at lower frequency. The problem is this also increases the height, increasing the vertical distance between sound sources, and ultimately reducing the height of the forward lobe. The nulls are drawn inward, cutting into the pattern. So an oversized horn isn't particularly useful.

A compromise is usually struck, one where the horn is sized to provide control in the horizontal, but allowing the vertical to lose control in the octave or two above cutoff. The vertical nulls from the driver interaction then cut into the widening pattern, limiting it at low frequencies near cutoff to approximately the same coverage angle as the HF beamwidth. These nulls are generally a few hundred Hertz wide, so they're a quarter to half octave at the typical ~1.2kHz crossover point. However, the 1:2 aspect horn/waveguide doesn't usually begin to gain pattern control in the vertical until 3kHz. So what we have is nulls limiting the beamwidth up to maybe 1.6kHz or so, and then an octave where it is wider, settling back into the desired vertical beamwidth above that.

It seems somewhat intuitive to think maybe the excess overlap of a low-order slope might give us wider nulls, cutting the edge of the pattern through a wider range. It could be used to augment pattern control through the region where the horn beamwidth widens in the vertical. But in practice, this is not the case. What actually happens is the nulls become very choppy. Instead of two clean notches at the crossover point, one above the center of the forward axis, the other below, what you get is a notch, then a big peak, then another notch, another peak, and so on. The peaks are actually more troublesome than the notches, because they can be really loud.

When the overlap region is too wide, interaction between drivers isn't clean. This is partially because of the shifting phase through the overlap region and sometimes partially because the off-axis response of the midwoofer isn't smooth. Summing is complex, and what usually (always) happens is major nulls form near the center of the crossover overlap region, and then large peaks and dips occur above and below the major nulls. Sometimes the peaks are very large and audible, and can become the most noticeable sound at certain angles. It's pretty bad when that happens.

If not too boring in this already too-long post, here's a little bit of empirical history to illustrate the roads I traveled to learn this:

In the 1980s, I was pretty happy with most of my designs and had almost concluded that they were a fait accompli. This was especially true of the constant directivity cornerhorns, because they just sounded so good. They were all very much like the line I still have today, except all crossovers were first or second-order.

But occasionally, I would notice a screeching sound in some vocals that I thought for a long time was caused by problems with vinyl. I was always meticulous with cleaning, and kept my cartridges anti-skate and set-down weight adjusted properly, but still, I thought the rare anomaly was from the albums. It sort of sounded like what I would expect from a damaged groove, a sort of intermittent screech on some material. You could go literally months and not hear it, but some material would do it every time.

That was one of my clues - the fact that some material did it consistently. Most times it was intermittent, but sometimes it always happened at the same place. The one I remember most distinctly was Dear Diary by the Moody Blues. The Leslie-modified vocals would exacerbate the problem and you could regularly hear it even at moderate volume levels.

As I said, I thought it was a problem with the source media until one day I realized I didn't hear it

loudspeakers that had a compression driver, and played the same material through the same

Back then, I really didn't like compression drivers anyway. They didn't reach far into the top octave even with mass-rolloff compensation, and so I was in a sort of experimental phase with the whole DI-matched two-way concept. The earliest models, just prior, were all three-ways, using a cone midrange and a slot tweeter. Of course, they weren't able to provide uniform directivity, they were traditional loudspeakers. But the magic I heard from the cornerhorns compelled me to continue on this path and to make uniform directivity a design goal of all my loudspeakers. The horn mounted compression driver made this possible where nothing else would.

Most of my design tools were mathematical models. They were primitive by today's standards, barely more than formulas done by pen and paper. I wrote computer programs to help me work the mundane formulas, but they were all just providing single points of information. No graphics, no response curves. All just single-dimensional data-points. To find some kind of response (whether impedance, amplitude and/or phase), I would "run the numbers" several times, to make a sort of a "sweep".

I was also limited in measurement capability, having just an oscilloscope, signal generator and microphone. While I could make measurements of stepped sines, I really only had visibility to about a 3dB resolution and my patience level limited me to making measurements about every few hundred Hertz. Naturally, my charts were very course, and my visibility limited. I relied on my measurements mostly for finding trends.  So I had poor visibility into this intermittent screeching problem, and had to hunt for a solution somewhat blindly.

Since the loudspeaker without the compression driver didn't do this, I considered it might be just a problem with compression drivers in general.  After all, they were generally used for prosound use only.  But it seemed worthwhile to explore different crossover configurations, and I found that crossing it higher in frequency or in slope made the problem go away.  My conclusion then was that the problem was caused by excessive excursion.

Through the 1990s, I was able to make incremental improvements.  Compression drivers got better, and better computerized design tools started becoming available too. Spice was available for the PC, and with it, I could graph transfer functions.  My modeling tools gave me much better visibility.  Then by the early 2000s, I was able to incorporate the Spice models with measurements, actually making digital crossovers that were configured using Spice models.  That was huge, in that it let me make any sort of crossover I wanted, by simply describing it as a Spice model.  Then a physical model exactly matched the mathematical model.  The modeling tool became so good that I would say the model was an exact digital representation of the physical model.

By this time, I had very good visibility into what's going on.  That's the state of the art today - DIYers have a large assortment of tools to work with.  Computer modeling tools and measurement systems are inexpensive and very capable.  But you still have to spend some time and know what to look for.  There's a lot of sweat equity in a good loudspeaker design.
Crossover optimization for DI-matched two-way speakersIn hindsight, I'd have to say the audible anomalies in the crossover region caused by minimalist crossovers are as much due to the jagged vertical response off-axis above and below the crossover point as it is the over-excursion of the diaphragm.  Both these things conspire to make a minimalist crossover untenable for waveguide loudspeakers with compression drivers.  If you move the crossover point high enough that the compression driver is properly loaded, then the midwoofer is strained and directivity isn't well-matched.  If you move crossover down to a more reasonable frequency band, then you get the screech from a strained compression driver, and the jagged off-axis peaks and dips above and below the nulls.

All this to say I wouldn't waste my time with a minimalist crossover in a waveguide loudspeaker.  It isn't really an application of Occam's Razor.  When you realize the complex acoustic and electrical impedance provided by the horn/driver combination is effectively a virtual circuit of dozens of components, the simple single component crossover becomes non-sequiter.  It isn't the best solution.
High-Fidelity Uniform-Directivity Loudspeakers