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