

High-Fidelity Uniform-Directivity Loudspeakers

Lots of people are experimenting with waveguides and constant directivity loudspeakers these days. This has been a π Speakers flagship design for decades, so naturally we understand the attraction to this approach, and the surge in popularity amongst manufacturers and DIY hobbyists alike.

This paper explores the technologies and development history of waveguides and constant directivity loudspeakers, with emphasis on high-fidelity designs rather than sound production or public address. We'll also explore some of the different design approaches that have been used, illuminating the strengths and weaknesses of each.

While this paper primarily addresses *loudspeaker system* design rather than horn or driver design, an understanding of the acoustic properties of various horns and waveguides is needed to properly implement them. Towards that aim, we shall explore some of the kinds of horns used in high-fidelity constant-directivity loudspeakers.

After a brief history of the evolution of the state of the art, we will explore the advantages of uniform-directivity in a high-fidelity loudspeaker, and why it is a desirable trait. There are those that are unconcerned with directivity in a home high-fidelity speaker, and who would place little emphasis on it. To them, the whole concept of constant directivity is a goal only for sound reinforcement speakers. Many audiophiles hold this belief, and in fact, until somewhat recently, convincing them otherwise was an uphill battle.

Of course, once the unconvinced audiophile heard a true high-fidelity quality sound system that produced a uniform sound field, all academic argument was made moot. But to even get them to listen to a demo system sometimes proved difficult, as many of them were convinced that constant directivity speakers were only good for public address and sound reinforcement, not for high-fidelity.

An audiophile's reluctance to constant directivity speakers is usually largely due to past experiences with sound reinforcement speakers that placed a high-priority on coverage and much less on sound quality. There were many developments in the late 1970s and 1980s that favored coverage over quality. Most involved using constant-directivity horns that really didn't sound very good, but could produce a uniform sound field. They essentially fragmented the sound to put it where they wanted it. Good for coverage, but bad for quality. These products put a bad taste in the mouths of audiophiles, so horns got a bad name during those years as a result.

To understand why development took this direction, we have to step back a few decades, to the early years of sound production. In the 1920s and 1930s, sound systems in theaters required horns to provide the necessary SPL for the audience. Amplifier power was at a premium, so speakers had to be efficient. This meant horn loading was a requirement.

Increased efficiency, by way of impedance matching, was the primary goal of horn speaker designers in the early days. Quoting Harry Olsen in *Acoustical Engineering*, "*The principal virtue of a horn resides in the possibility of presenting practically any value of acoustical impedance to the sound generator.*"

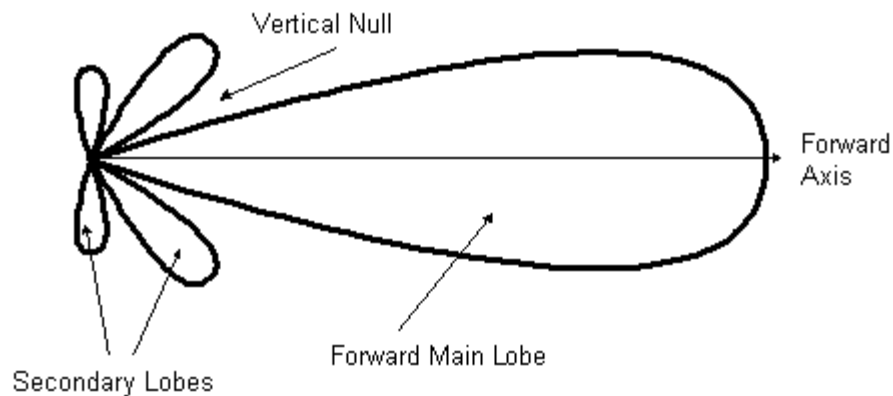
Horns offer resistance to diaphragm motion because they create pressure in the front chamber, forcing sound to exit through a relatively small orifice in the throat. This limits diaphragm excursion. The expansion from the throat to mouth allows a transformation from high-pressure/low-displacement at the throat to low-pressure/high-displacement at the mouth. This is why horns increase the efficiency of a loudspeaker, because they match the impedance of the driver to the air.

One problem of early theater horns was the path length differences between sound sources caused noticeably different delays. An example is the Western Electric "Wide Range" system, comprised of two WE 555 compression drivers on a twelve-foot-long 12A "snail" horn and four 18" woofers on an open baffle. There were also two separate Bostwick tweeters, forming a three-way system. The lengthy horn delayed the midrange response so much that when a tap dance routine was played through this loudspeaker system, each individual tap was heard twice.

The path length problem was solved with offset, or in some cases, by truncating the longest horn or omitting it entirely. This solved the biggest time delay problem. By the mid 1940s, most theater loudspeakers used horns for HF, and horns or reflex/horn hybrids for LF. Using truncated basshorns or reflex/midhorn hybrids kept path length differences and delay problems to a minimum.

Another problem was the non-uniform directivity that early horn loudspeakers generated. Listeners in the area where the speaker was aimed enjoyed relatively good quality sound, but audience at the sides would not. Directivity control was simply not the main goal of early horn designers. But this quickly became more and more important, because without adequate coverage, only a small portion of the audience received the full benefit of the sound system. The search for the solution to the directivity problem became an ongoing quest that continues to the present day.

The directivity problem is multi-faceted. One issue is the fact that there are several sound sources in 3D space instead of a single point source. This creates interference in the bands where the drivers overlap, generally in the crossover region. Where interference causes cancellation, the sound level is low and this marks the edge of a lobe or “pocket” of sound.



A second issue is the directivity characteristics of the drivers, themselves. Direct radiators and some horns tend to become more and more directional as frequency goes up. Other horn shapes remain constant in directivity. Some features like hard edges and slots cause the pattern to widen. And finally the electro-mechanical parameters of the drivers and the crossover components tend to make phase shifts that modify the way the sound sources interact, sort of bending their interference patterns, changing where the forward lobe is pointing as well as the outer nulls and secondary lobes.

The horns used almost always were exponential, meaning the cross-section area of the horn is exponentially proportional to its length. These have long, narrow throats with ever expanding mouths. The advantage of this kind of horn is it tends to have large bandwidth, providing a good acoustic load at both high and low frequency. One can expect an exponential horn to provide adequate SPL, even response and low distortion over its passband.

The problem with an exponential horn is it tends to beam at high frequencies. As a result, only listeners that are directly in front of the speaker will hear high frequencies. Ironically, a simple conical horn, one with straight sides, provides constant directivity, low distortion and it does not beam.

A megaphone is an example of a conical horn and while primitive, it can be expected to direct sound equally within the arc of its flare. Basically, if you can see down the throat of a conical horn, you can hear its output just the same as if you were directly in front of it. This is the definition of constant directivity, which is the goal. If that were the *only* goal, it would be the hands-down winner. However, the conical horn has its share of problems, namely, a lack of acoustic loading at low frequencies which affects efficiency and can also have an impact on response.

Because of the conical horn's reduced acoustic loading properties, horn designers almost always chose exponential flares in the early days. They would deal with its directivity problems in various ways.

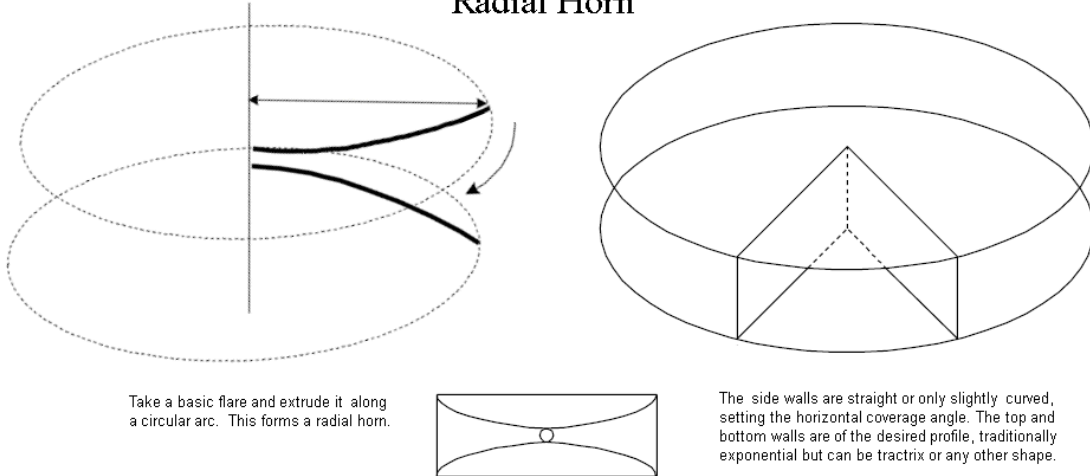
The first attempt at a solution was the multicell horn, which was essentially an array of splayed exponential horns. This worked better than a single, large exponential horn because the high frequency beaming is narrowed in scope, each segment being already somewhat small anyway. But they still tended to beam in the midrange and they were very expensive to manufacture.

Multicell Horn



The next step in horn evolution is the radial horn. It provides constant directivity in the horizontal plane, because it has straight or nearly straight side walls. Vertical directivity is determined by the horn profile, the curvature of the top and bottom walls.

Radial Horn



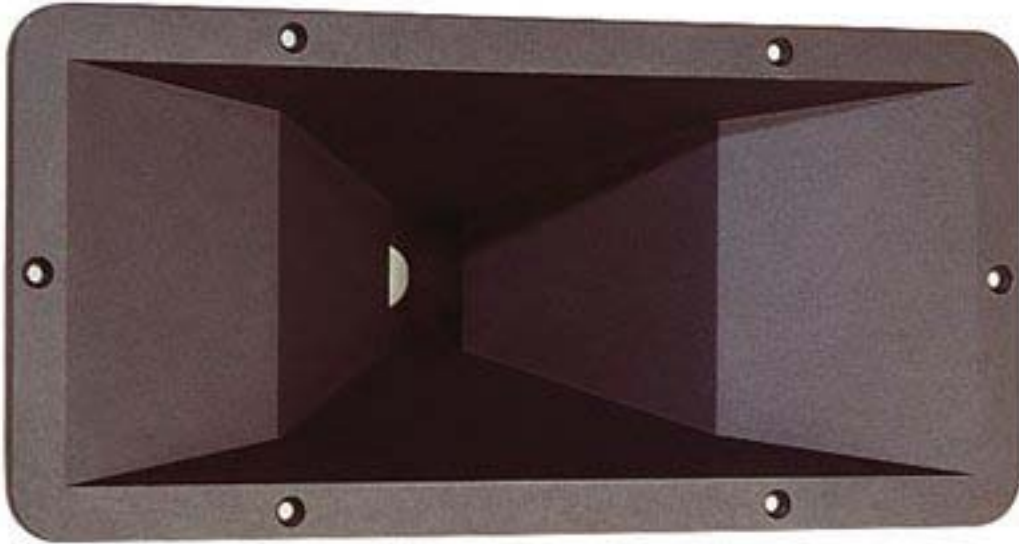
Take a basic flare and extrude it along a circular arc. This forms a radial horn.

The side walls are straight or only slightly curved, setting the horizontal coverage angle. The top and bottom walls are of the desired profile, traditionally exponential but can be tractrix or any other shape.

The radial horn was much easier to manufacture than multicell horns, and it had an advantage in that horizontal directivity was nearly constant. Since nearly all of the early radial horns had exponential horn profiles, they provided a good acoustic load. But naturally, they still exhibited collapsing directivity in the vertical plane. This drove horn designers to strive to develop constant directivity horns that were uniform in both the vertical and horizontal planes.

The first constant directivity horns appeared in the late 1970s. Several designs have been made since the introduction of the first CD horns in 1978. But all of them required a diffraction slot in the throat, after an initial exponential section that provided the majority of the acoustic load. The diffraction slot widened the pattern, allowing the final sections to set the radiating angle. The final sections were straight-sided.

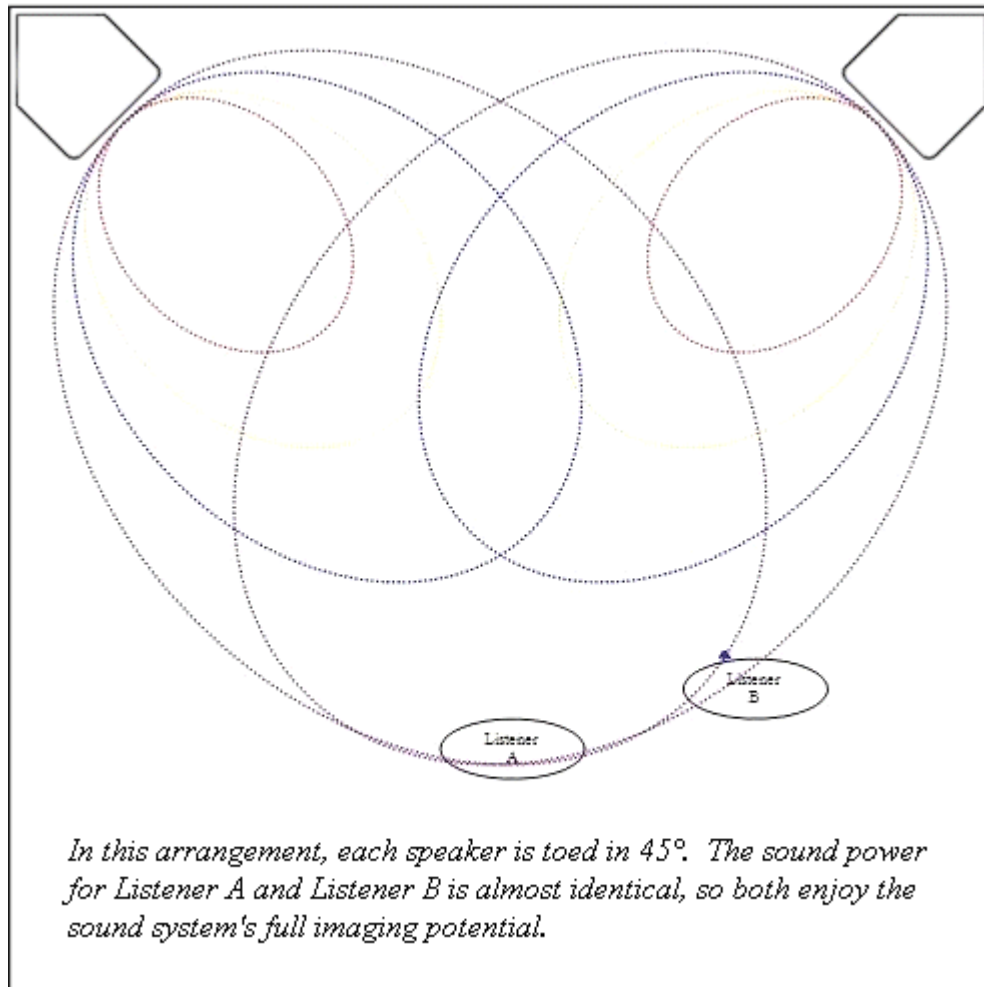
Constant Directivity Horn



This was the state of the art when π Speakers was born. Horns had gone from being primarily acoustic impedance transformers used for their efficiency to being primarily waveguides for pattern control, which are two very different things. While horns are useful for both purposes, the flare profile chosen tends to optimize one or the other, but not both.

This is what drove the first constant directivity loudspeaker developed by π Speakers, the Constant Directivity Cornerhorn. The designer, Wayne Parham, had been influenced by the Klipsch Cornerhorn, but wondered if it couldn't be improved by focusing on the directivity aspect rather than the acoustic load. He reasoned that the corner provided what amounted to a large constant directivity horn, and that the labyrinth inside might be omitted.

The first π Constant Directivity Cornerhorns were built in 1980, and were immediately successful. They proved right away that the concept was viable. Not only is the loudspeaker able to provide constant directivity, but its requisite toe-in provides a self-balancing effect for the stereo image. Lateral movement left or right of center tends to bring the listener closer to one speaker, but more aligned with the forward axis of the other. This makes the “sweet spot” quite wide where stereo imaging is good.



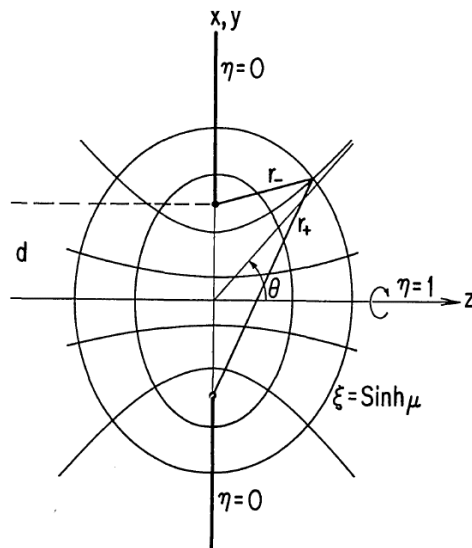
The “magic” of the π Cornerhorn was undeniable in light of its popularity. Long before the internet was around, this word-of-mouth sensation spread quickly in the Midwest. This design approach is still the best method to achieve constant directivity over a wide bandwidth. There is no other loudspeaker configuration that can match the π Cornerhorn’s ability to provide constant directivity from the Schroeder frequency upwards, basically through the entire audio band from 150Hz on up.

Not every home has adequate corners to support a π Constant Directivity Cornerhorn, so the matched-directivity two-way loudspeaker was developed shortly afterward. These use a midwoofer that covers the bass and midrange, and is crossed-over to a constant directivity tweeter at the frequency where the direct radiating midwoofer beamwidth collapses to match the 90° tweeter in the horizontal plane. The center-to-center spacing is kept close enough that the nulls are spaced outside the tweeter's pattern at HF, usually around $\pm 22.5^\circ$ above and below the center of the forward lobe. This provides a nice $90^\circ \times 45^\circ$ coverage pattern, similar to that from the π Cornerhorn.

The thing that bothered the designer was most CD horns tended to sound quite harsh. Parham assumed it was due to the sharp edges inside the CD horns, and began to use radial horns instead. They provided nearly constant directivity in the horizontal, and depending on the shape of the top and bottom flare profile, gently collapsing directivity in the vertical. This is a good solution, having good coverage and smooth sound.

By the 1990s, newer radial horn models and some of the products marketed as waveguides began to appear on the market. These offer uniform directivity without the sharp edges in earlier constant directivity horns that made them sound somewhat harsh. These smoother horns were always chosen for their superior sound quality.

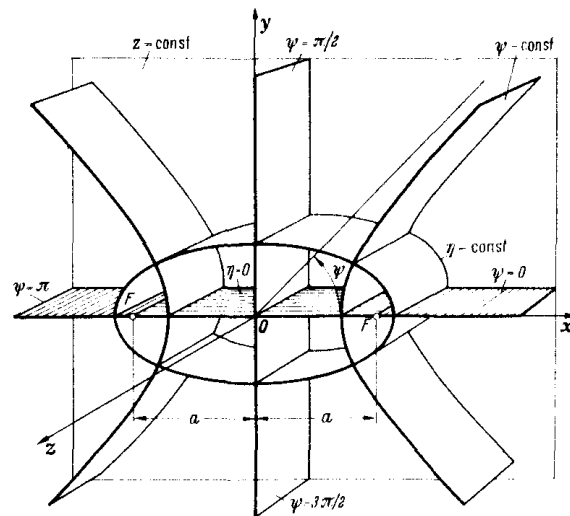
Radial horns and waveguides with smoothly radiused throat transitions are generally better sounding than horns with sharp edges, but a particular flare profile stands out as being the best. An oblate spheroid (which is the natural shape of the earth and other planetary bodies) is the basis for this unique curve.



Oblate Spheroidal Coordinate System

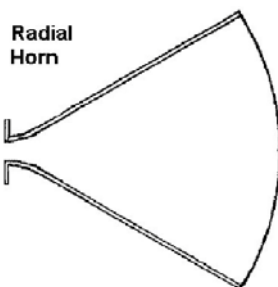
When a conoidal surface is defined using lines drawn tangent to an oblate spheroid, the resulting curve is a hyperboloid of one sheet. This forms an oblate spheroidal waveguide/horn.

The key feature of the oblate spheroidal coordinate system is the ellipse. A similar coordinate system is the elliptic cylinder, which also is based on the ellipse. Both are rounded - not spherical but elliptical – and both have an ellipse as their cross-section. The difference is the oblate spheroid is egg shaped, and the elliptic cylinder is more like a pipe.



Elliptic Cylinder Coordinate System

Devices using these elliptical coordinate systems are known to be most optimal for smooth wavefront propagation not only in acoustics, but also in the fields of optics, microwaves and ultrasonics.

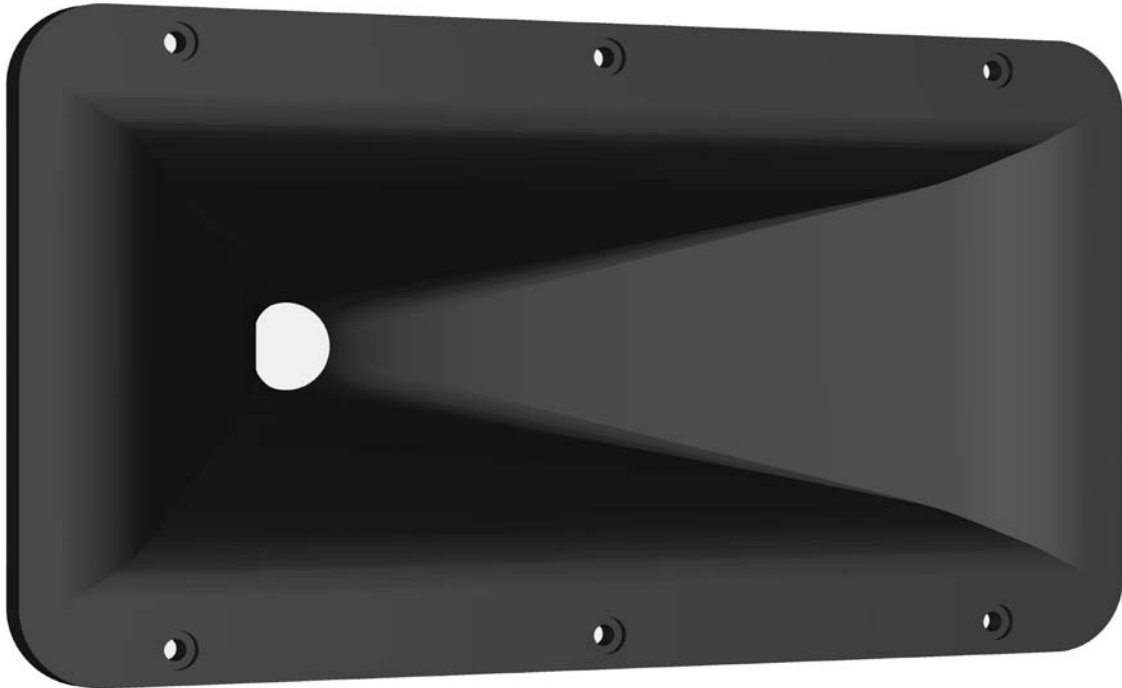


Radial horns, revisited. Notice the radial horn schematic shown to the left. It is basically straight-walled, with a radiused throat transition. This shape can be modified to use an elliptical coordinate system and gain the benefits of using this shape. An oblate spheroidal section at the throat is blended with an elliptic cylindrical section at the mouth to form a waveguide/horn having any desired coverage angle.

Where waveguide/horns are concerned, there is a continuum of useful profiles. Horns with narrower tangential angle tend to provide better acoustic loading, with correspondingly smoother response. They also limit unwanted early reflections from nearby boundaries. However, too narrow beamwidth is not particularly useful for high-fidelity or home theater loudspeaker systems.

One way to improve acoustic loading without decreasing horizontal beamwidth is to use an asymmetrical mouth having smaller vertical dimension. This allows the area expansion to be optimized for acoustic loading while simultaneously providing the desired coverage angle. It also has the added benefit of reducing vertical spacing between sound sources.

π Horn/Waveguide

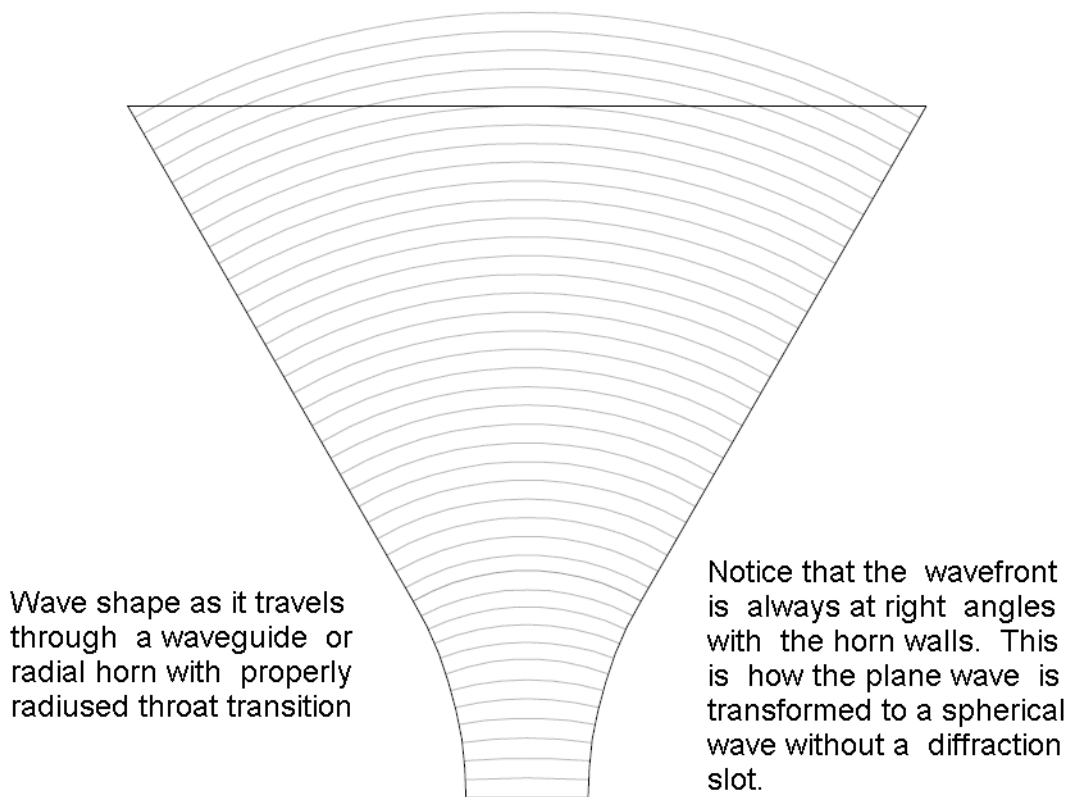


There are several trade-offs to consider when choosing a waveguide/horn. Larger horns tend to retain pattern control at lower frequencies, but they necessitate larger center-to-center distance between sound sources. Smaller waveguide/horns provide less acoustic load on the diaphragm, especially at low frequency, and sometimes suffer from resonant modes and corresponding response ripple. In either case, the flare profile and length/area ratio have a large effect on pattern control and response smoothness too. There are usually competing priorities, but using appropriate design choices, one can strike an attractive balance.

π Speakers have asymmetrical waveguide/horns with constant 90° horizontal coverage and approximately 45° vertical coverage. The horns used have smooth features, so sound is not harsh like earlier CD horns. Horizontal coverage must be uniform and sound quality must be high. Vertical coverage is good floor to ceiling, but no more, since ceiling reflections are undesirable. The crossover and driver spacing put vertical nulls just outside the coverage angle, where they punctuate the edge of the pattern in the crossover region.

The primary distinction between constant directivity horns and waveguides is the mechanism for introducing and directing the wavefront. Early constant directivity designs used a diffraction slot in the throat to transform the planar wave coming from the compression driver into an expanding spherical wave. The walls of the horn then constrained the spherical wave and ultimately set the pattern.

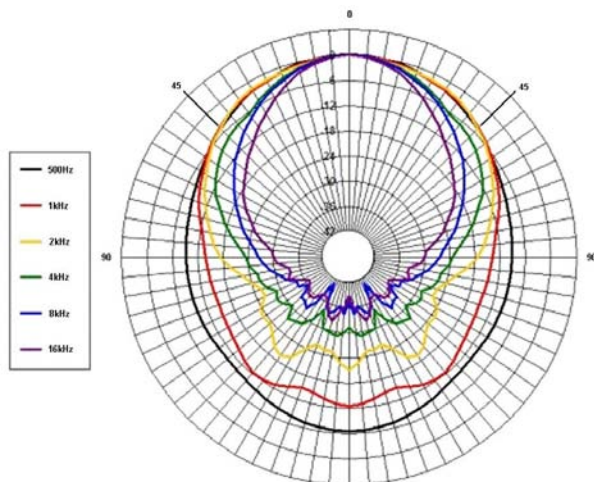
Waveguides bend the planar wave as it exits the compression driver and enters the throat. Rather than using a sharp transition that rapidly transforms the plane wave to a spherical wave section, the gently radiused throat causes the plane to bend, ever closer to a section of a spherical wave as it progresses down the throat.



π Speakers waveguide/horns provide beamwidth of approximately of $90^\circ \times 45^\circ$. The horizontal beam is constant 90° through the entire passband, and the vertical beam gradually narrows through the crossover region, reaching a constant 45° above the crossover point. This was the design goal, as it allows us to develop a crossover that provides the desired pattern and generate a nice forward lobe with widely spaced vertical nulls.

There are three important reasons why $90^\circ \times 45^\circ$ is the optimum coverage pattern for home high-fidelity:

1. Rooms are generally proportioned this way, with relatively low ceilings and 90° corners. So when a loudspeaker provides 90° horizontal beamwidth, it can be angled inward 45° which covers the whole room yet limits reflections from adjacent walls.
2. You don't need (or want) much HF output at large vertical angles because it does nothing but create unwanted floor and ceiling reflections. But the forward lobe should be tall enough to be useful, and nulls should be outside the listening area, directed out of the way.
3. The geometry of a $90^\circ \times 45^\circ$ horn tends to naturally make a shape that is wider than it is tall. This allows tight center-to-center spacing, which when accompanied with proper crossover design, puts the nulls and secondary lobes far outside, widely spaced. This gives a nice, useful coverage pattern.



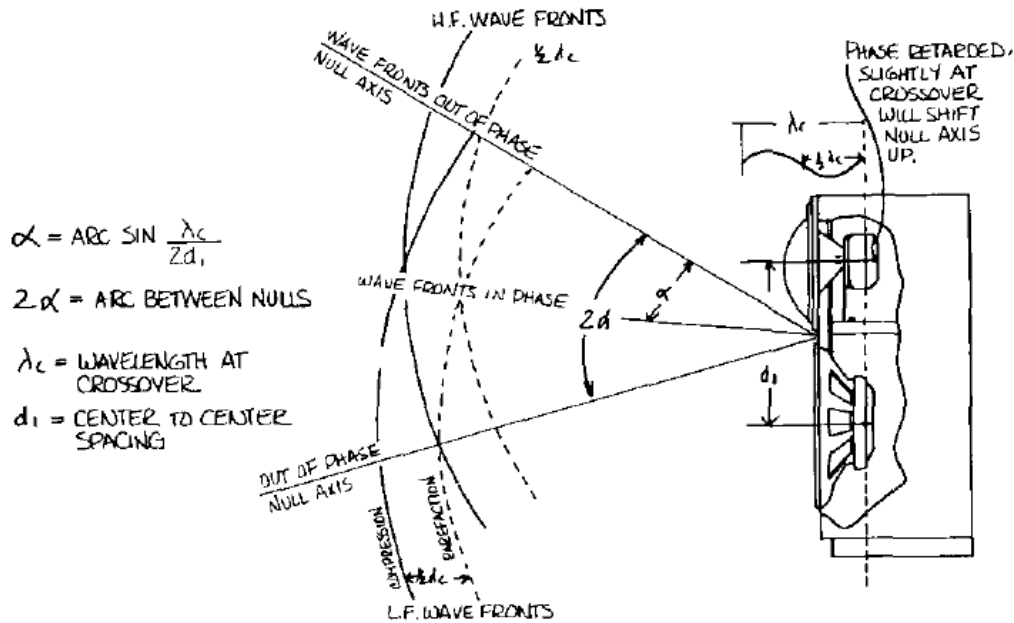
Horizontal Coverage of Matched-Directivity Two-Way Loudspeaker Crossover at 1.3kHz - Notice Directivity is Matched and Uniform

This makes a compelling argument in favor of HF horns with an asymmetrical $90^\circ \times 45^\circ$ pattern. The relatively wide horizontal angle provides a large coverage area. The 90° angle works out well for another reason, it matches the narrowing directivity of a large-format midwoofer at around 1.0kHz to 1.6kHz. And the narrow vertical angle matches the null angle from vertical spacing of the drivers.

If two sound sources are stacked vertically, then movement along the horizontal plane doesn't change the difference in distances between the listener and the sound sources. So as long as summing is in-phase on the forward axis, it will be also be in-phase off-axis in the horizontal plane up to relatively wide angles.

But movement along the vertical plane changes the distance between the sound sources and the listener, delaying the more distant driver by a small amount. At frequencies and positions where the delay represents a 180° shift, nulls appear in the polar response.

No loudspeaker system with more than one driver on a baffle can have the two sound sources in-phase at all locations and at all frequencies. Somewhere, there is going to be destructive interference. But the idea is to place the anti-phase nulls outside the coverage angle. The best loudspeakers are those where nulls are outside the wall angle of the HF horn where they do nothing more than punctuate the cutoff at the edge of the pattern.



Geometric causes of vertical off-axis response nulls.

It is important to position the drivers vertically where the off-axis nulls are just outside the angular coverage of the horn. The first thing to do is to set the centerline for the forward axis. This is done by aligning the acoustic centers on the Z-axis, front-to-back.

Alignment of acoustic centers is simple in principle, but sometimes rather complex in practice. The acoustical load reactivity changes as a function of frequency and the electrical phase of the crossover changes too, both of which shift the apparent position of the sound sources. The designer must take each of these things into consideration, acoustic reactivity, physical position and electrical phase.

Aligning the sound sources on the Z-axis sets the forward axis centerline, which biases the arc of in-phase summing in the vertical plane. The nulls that form above and below this centerline are found at angles determined by the distance between drivers and the frequencies where they are both used.

When using a horn with approximately 45° vertical coverage, it is best to set the arc between nulls around 40-60° degrees through the crossover region, which is accomplished by vertical center-to-center separation on the baffle and selection of crossover frequency. Closer spacing and/or lower crossover frequency makes the null angle wider, and further spacing and/or higher crossover makes the null angle narrower.

This relationship makes horns with taller vertical coverage unattractive. Even if it weren't for the problem of ceiling slap, vertical nulls and secondary lobes make large vertical coverage problematic for this loudspeaker configuration.

As an example, consider a 90° HF horn that's large enough to be used down to 1kHz. It should have mouth dimensions of at least a foot. If used with a direct radiating midwoofer, the woofer will need to be 12" or 15" diameter, because anything smaller has too wide a pattern at 1kHz. The collapsing directivity of the woofer should match the horizontal beamwidth of the tweeter.

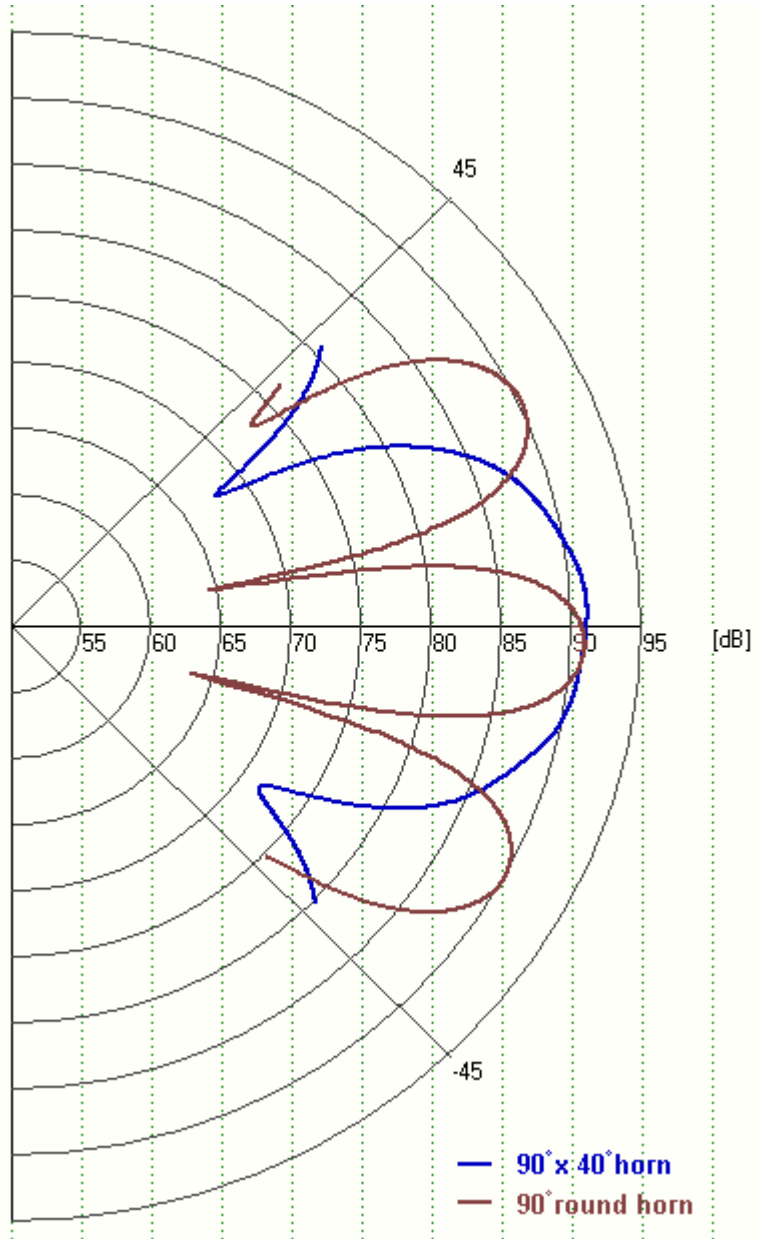
If considering an axi-symmetrical horn, the problem is the size. Both woofer and tweeter horn have to be 12" to 15" diameter to match directivity at the 1kHz crossover point. That makes the absolute minimum center-to-center dimension be at least 12" and more likely 15" to 18", which puts the null angle right in the middle of the pattern. There is no constant directivity possible outside the main forward lobe. So it is not reasonable to attempt a constant directivity design of this configuration having a vertical coverage angle this tall. It is much better to limit the vertical pattern to a beamwidth smaller than the null angle.

This is the fine art and science of directivity matching. Horizontal directivity should be matched at the crossover point and as long as the sound sources are stacked vertically on the baffle, this is largely determined by frequency. Phase interaction between sound sources doesn't change with movement along the horizontal plane, so horizontal pattern matching is fairly straightforward. But vertical directivity is influenced by phase interaction, so a little bit more design effort is required to get the vertical pattern right.

A real-world example is shown to the right. The blue line shows the vertical polar response of a two-way matched-directivity speaker with 15" midwoofer and 90°x45° radial horn, crossed over at approximately 1kHz. Notice the nulls are around +/-30°, making a useful forward lobe.

The brown line shows the vertical polar response of the same 15" midwoofer and compression driver, using the same crossover frequency. The only difference is it uses a 90° round horn, which increases vertical spacing.

Both speakers have good horizontal directivity, but it is clear that the speaker with the round horn has a very thin forward lobe, making it undesirable. The nulls fall dangerously close to the forward axis and listeners may easily find themselves positioned in a null or secondary lobe, where sound quality is poor.



To many people, this is unexpected. Most people see the round horn and expect response to be uniform in both the vertical and horizontal planes. They are usually surprised to find that the round horn actually creates a much thinner forward lobe than the shorter rectangular or oval horn does. This is primarily because of the vertical distance between sound sources.

All drivers have some things in common, as do all constant directivity horns. There are differences too, and those must be addressed. But the core features are common with all of them. One is the 6dB/octave low-pass filter function that diaphragm mass creates. Another is an additional 6dB/octave from inductance. The first pole is mass, and this starts around 4kHz in most 1" exit compression drivers. The second pole, from inductance varies from driver to driver, but it usually starts to become significant between 10kHz and 16kHz. There is a third pole too, which is the compliance from the front chamber, usually forming a knee around 16kHz to 18kHz.

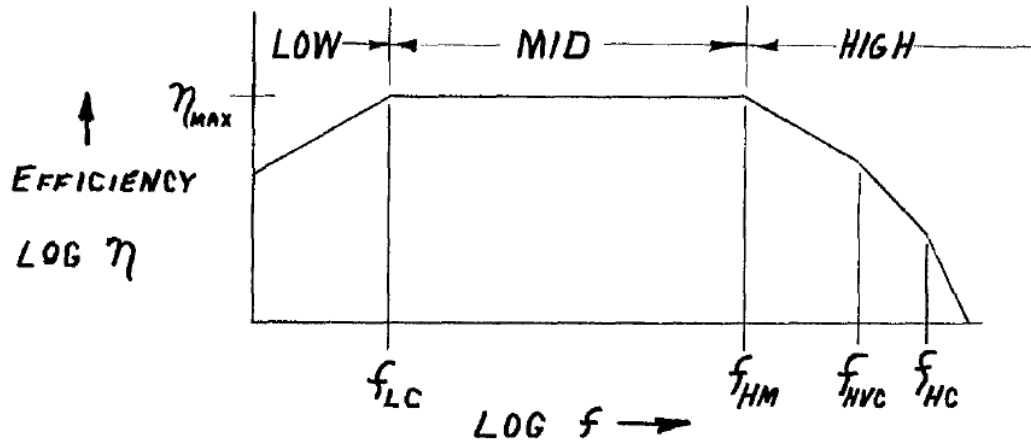


Fig. 3. Idealized frequency response of horn-driver system. Horn cutoff frequency f_c is assumed to be very much lower than f_{LC} (C_{MET} very large). The midrange band is defined primarily by driver and back cavity compliance rolloff on the low end ($f < f_{LC}$) and driver effective moving mass rolloff on the high end ($f > f_{HM}$). Secondary high-end rolloffs due to driver voice-coil inductance ($f > f_{HVC}$) and front cavity compliance ($f > f_{HC}$) are exhibited.

Horns that have collapsing directivity become more directional as frequency rises. These are said to beam, and they provide acoustic equalization, boosting high-frequencies on-axis because the energy is concentrated more as frequency goes up. This tends to provide a natural conjugate for mass-rolloff, meaning the on-axis sound level remains flat even though power response begins to droop.

Direct radiators do this too, because they are omnidirectional at low frequencies where wavelengths are large compared to radiating diameter. The sound source is acoustically small at LF, much like a point source. But when frequency rises, as the radiator becomes acoustically large, it becomes more and more directional. This is because sound is emitted from all points on the surface of the diaphragm, and path length issues begin to come into play.

Horns that develop a constant directivity pattern deliver the compression driver output without any additional filtering. If a horn generates a pattern with perfectly constant directivity, then the power response is exactly the same as the on-axis response. Power response is total output in all directions. The constant directivity source sounds the same in all directions within its coverage angle.

Of course, real-world horns deviate from this ideal. Sharp-edged constant directivity horns tend to have response spikes where diffraction causes ripple. But they usually tend to generate a pretty constant radiation pattern, especially at higher frequencies so the on-axis response tracks the off-axis response. Radial horns and waveguides, on the other hand, do not use diffraction to widen the pattern. So above 10kHz - where the throat of a 1" exit compression driver becomes acoustically large - beamwidth will begin to narrow.

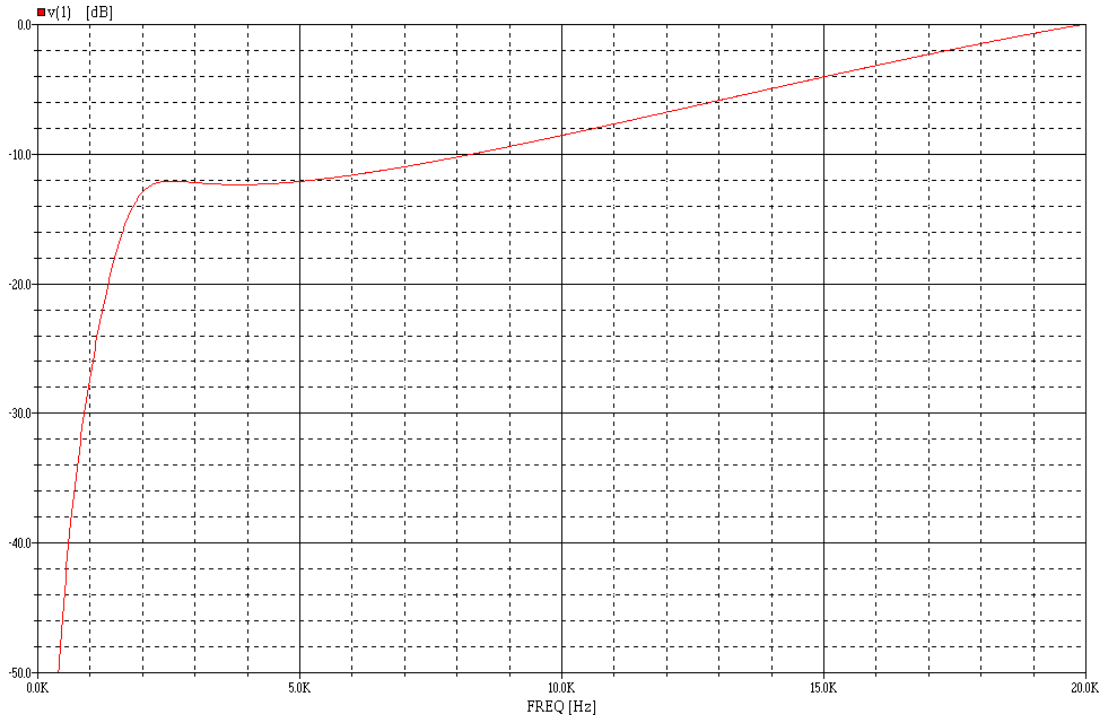
At the low end of the frequency scale, horn resonances may cause response ripple. This is because the acoustic load becomes increasingly reactive. At higher frequencies, the horn can be depended on to be primarily resistive but down low, it's not. Also, the horn/waveguide will lose its ability to provide pattern control at some point. If it is asymmetrical, it will likely lose control in one plane before it does in the other, usually maintaining pattern control to a lower frequency in the horizontal than the vertical.

These are all deviations from the constant directivity ideal that modify the on-axis SPL response by way of directivity change. Remember that if power response is even and directivity changes, then on-axis response will also change. These kinds of things must be considered in the crossover, by avoiding anomalous regions where possible, or by compensating for them.

The goal is to provide uniform directivity. The radiation pattern should be constant where possible, or at least gradually and uniformly collapsing where constant directivity is impossible. We have found empirically that low frequency radiation can collapse gently but high frequencies are best radiated over a constant area. All other things being equal, the lower frequency that directivity can be maintained constant, the better.

The crossover is critical in a loudspeaker designed for constant directivity. Not only must it be designed for flat response on-axis, but it is also partially responsible for setting the position of the forward lobe, nulls and secondary lobes. Since the crossover introduces a certain amount of delay, the beam can be "steered" with crossover phase. This makes it even more important that the designer be diligent when configuring and optimizing a constant directivity loudspeaker.

An important requirement for the constant directivity crossover is top-octave compensation. As seen earlier, the power response of a compression driver falls off at 6dB/octave above mass-rolloff. In a typical 1" exit compression driver, this starts around 4kHz, so top-octave compensation generally adds about 12dB boost total.



Top-octave compensation equalization can be provided actively or passively. In the case of the popular two-way loudspeakers, with a direct radiating midwoofer mated to a constant directivity horn, the difference in sensitivities is usually enough to make required attenuation of approximately this amount. Compression drivers on horns/waveguides are very efficient, and even the most efficient direct radiators are at least 10dB lower at the same power level. That gives the designer an opportunity to build passive equalization into the tweeter attenuation circuit.

During our development and evolution of the first constant directivity cornerhorns and of the first matched-directivity two-way speakers, the designer of π Speakers, Wayne Parham, studied several crossover topologies and circuit types. These are shown in another publication available from π Speakers, *“Speaker motors and passive crossover filters”*, often referred to simply as the *“Speaker Crossover”* document.

There are many ways to accomplish the requisite frequency splitting and equalization necessary. Some circuits are simple, others more complex. Each has benefits and trade-offs, and is useful in various specific applications. All are shown in the *“Speaker Crossover”* document.

After extensive evaluation, one basic topology stood out over all the others. The low-pass and high-pass slopes are selected according to their phase, specifically, so that acoustic and electronic phase combine to set the forward lobe in the proper position. The top-octave compensation circuit is set using a specialized π pad, having two resistors forming the input shunt and series attenuator followed by the complex impedance of the compression driver as the output shunt/load. Optionally, a capacitor may be placed across the series resistor and/or a resistor placed across the driver for additional response tailoring.

In the π crossover, the first resistor (R1) is a series resistor that sets the amount of attenuation to the tweeter, basically setting the average sound pressure level. The second resistor (R2) provides the majority of the load on the core high-pass filter which sets damping for the circuit. The right amount of damping can be used to provide slight peaking at crossover, if required, to create a shelf of flat response up to the point where mass-rolloff augmentation begins. Resistors (R1) and (R2) combine with the driver to form a π pad, where voltage division of the driver voice coil inductance and the series resistor create a 6dB/octave slope for mass-rolloff compensation.

Optional component (Rs) is a shunt resistor placed across the driver to swamp impedance peaks, where required. Capacitor (C1) can be used, placed across the series resistor (R1) to provide additional top-octave augmentation. This is often required if voice coil inductance doesn't provide enough augmentation or where (Rs) is used, swamping it. These two locations are provided for additional flexibility manipulating the tweeter circuit transfer function.

All horns and waveguides have an impedance peak near their cutoff point, and some have more than one. This impedance fluctuation interacts with crossover components and can manifest itself in the response curve. In addition, changing directivity down low modifies response both on-axis and off-axis. These kinds of things affect how much damping is required, configured using the (R2) input shunt and (Rs) output shunt resistors.

The reason this particular filter type was chosen above all others was its inherent stability and configurability. Over the years, after several loudspeaker designs were implemented using a variety of different drivers and components, Parham noticed that each tended to require a different solution unless it was done with this π pad configuration. The π crossover works for *all* constant directivity designs. It has proven to be the most robust, and least vulnerable to driver shifts, whether from thermal or environmental conditions, or even from unit-to-unit variations.

This is why the π crossover network works so well. It allows all of the “critical points” to be addressed, the lower shelf below mass-rolloff, the amount of compensation boost for the top-octave and the overall average SPL, all can be set with the proper choice of values in the crossover. It is a reliable and configurable circuit that behaves predictably with a wide variety of compression drivers and waveguide/horns.

Earlier in this document, we examined diffraction from certain kinds of horns but this is not the only cause. It can also be created by cabinet corners and even outside the loudspeaker in the environment. Any sharp edge will create diffraction. It causes a wavefront discontinuity, ultimately creating a phantom source, with the secondary wavefront travelling in a new direction. Some of the energy travels in the original direction, but some is vectored off, "curving around" the diffraction edge. Just like reflections can cause self-interference that generates time/frequency anomalies, so can diffraction. The two wavefronts may interact in the environment and develop complex modal behavior.

An aperture is said to be acoustically small when it is small compared to the wavelength of sound passing through it. Similarly, an edge is acoustically small when the radius of the bend is small compared to wavelength. The physical size of a feature compared to wavelength is very important to its acoustic properties. For example, when an aperture is acoustically large, it acts like open space to a wavefront and causes no diffraction. If a radiused corner is large, meaning the bend is slow and gradual compared to wavelength, the wavefront will travel along it without discontinuity. But as frequency drops, wavelengths get larger so physical features become increasingly small from an acoustic perspective. Once the scale changes, the acoustic behavior is different.

For a real world example, consider a loudspeaker on a baffle. At low frequencies, where the baffle is much smaller than wavelength, sound radiates omnidirectionally. The baffle is acoustically small and so the loudspeaker acts like a point source. The baffle dimensions do nothing to affect directivity. At high frequencies, the baffle is acoustically large, so it constrains the radiating angle and forces the sound to travel forward. The pattern is no longer omnidirectional as it was at very low frequencies, but rather half-space. As an aside, as frequency rises even further, where the cone of the speaker becomes acoustically large, the pattern will narrow even further because of path length differences across the radiator, itself. At that point, the baffle is no longer setting the pattern.

Higher frequencies will radiate into half-space because of the baffle, lower frequencies will radiate omnidirectionally because the baffle is acoustically small. If the speaker is in open free-space, it will generate greater SPL on-axis at higher frequencies due to the change in directivity.

There is a transition region, where baffle dimensions are approximately equal to wavelength, and in this region the behavior is somewhat complex. There is self-interference between the vectored "phantom source" wavefront and the original. The sound actually narrows just a little bit more than the (half-space) baffle angle in this region. Through this region, directivity vacillates between slightly narrower and slightly wider than the (half-space) beamwidth set by the baffle. Since directivity shifts while power response is constant, there is some ripple in the on-axis response through the transition region.

Now take the baffle and fold it inward to form a V-Shape, effectively making a waveguide or horn. All the same things happen, but the constraint angle creates a beamwidth that is narrower than a baffle. The radiation pattern is set by the walls of the horn at higher frequencies, but as frequency drops, the horn loses pattern control and the beam widens. In the transition region, it narrows briefly before it widens, and in this region, there can be some directivity and response ripple.

During our first few years, all our loudspeakers had rounded corners to reduce baffle edge diffraction. Since we noticed the horns with sharp edges sounded harsh, we reasoned that all sharp edges were probably bad, including those at the edge of the baffle. However, we relaxed this requirement because we have learned empirically that cabinet edge diffraction from large 90° constant directivity loudspeakers is inaudible.

For large cabinets, the baffle-step directivity shift from omnidirectional to half-space happens at a relatively low frequency. Indoors, the walls become reflectors which introduce self-interference modes and these have much more impact on in-room energy distribution than baffle step. So large cabinets used indoors can be thought of as operating solely in half-space or even smaller radiating angle, depending on the radiating characteristics of the drivers, themselves. If placed in corners, the maximum radiating angle is even smaller, quarter-space or even eighth-space, depending on placement.

As for the phantom source from diffraction, and the ripples in response, we find that this is certainly audible when introduced very near the sound source. Horns with sharp edges inside are easily identified by their sonic signature, both by measurement and by ear. Most people subjectively describe them as sounding harsh, when compared with similar horns having no sharp edges. However, the further away an edge is, the less audible impact it has. Particularly when a directional horn is placed on a baffle, thereby radiating a beamwidth that is already narrower than the baffle, the cabinet edges simply have little effect on sound quality.

π Speakers are all about balance – balanced power response, balanced dynamic range, balanced directivity, balanced features of all kinds. It isn't enough for us to match the horizontal directivity of the tweeter and the woofer and claim constant directivity. We match the vertical directivity with the null angle too.

In systems with direct radiating midwoofers, we choose high-quality drivers that can easily handle 10x the power of the tweeters, because this matches the dynamic range. That's important because it doesn't make much sense to use a woofer that enters power compression long before the tweeter does. This isn't just an academic argument for "good practices", it actually has an audible impact because it sounds unnatural to have one driver straining while the other is just loafing along. A well designed loudspeaker has balanced subsystems, none that are optimized at the sacrifice of others.

Our matched-directivity two-way loudspeakers use a large high-efficiency, low-distortion direct radiating midwoofer and a tweeter horn that provides constant directivity with a $90^\circ \times 45^\circ$ pattern. The frequency split is done at the frequency where the midwoofer pattern has narrowed to 90° , matching the tweeter's horizontal pattern. This frequency, and the crossover phase and baffle spacing, are all simultaneously chosen to set the vertical nulls outside the 45° degree coverage angle.



three π loudspeaker

The midwoofer chosen for a matched-directivity two-way loudspeaker should be of particularly good quality, because it is being used to a relatively high frequency. The cone shape and material must be appropriate for use up to over 1kHz without breakup. This usually means the cone is curvilinear rather than straight-walled, and the material is a pulp that provides a good degree of internal damping. It is also desirable that the motor structure have Faraday rings for flux modulation control. These reduce magnetic non-linearities, which are the biggest source of distortion from a direct-radiating midwoofer.

The radiation pattern from a matched-directivity two-way loudspeaker is omnidirectional at low frequencies, i.e. set by room boundaries. The baffle is large enough that it begins to limit the radiation to half-space around the Schroeder frequency. Below that point, room modes set the pattern, preventing coverage from being truly omnidirectional. Since directivity in the modal region is ambiguous, there is no transition from free-space to half-space, i.e., no well-defined baffle step.

As frequency rises through the midrange, the beamwidth begins to narrow because of the physical dimensions of the midwoofer. At the crossover point, the horizontal beamwidth matches that of the tweeter, and above that point, directivity becomes constant, set by the horn. Loudspeaker placement is non-critical, however, it is usually advantageous to place left and right speakers on each side of the listeners with their forward axes crossed in front of the listeners, as shown earlier in this document.

Our constant directivity cornerhorn uses a high-efficiency, low-distortion woofer, physically set snug in the apex of the room's corner. A midrange horn and tweeter sit on top, each having 90° x 45° coverage. The frequency split between tweeter and midrange is not especially critical, since both have 90° horizontal coverage patterns. However, to ensure the nulls stay outside the vertical pattern, crossover and driver spacing are generally about the same as the matched-directivity counterparts. The woofer and midrange are blended at the low end, smoothing what would otherwise become a floor bounce notch. Multiple distributed sound sources are particularly useful for modal smoothing as frequency drops below the midrange band.

A stereo pair of constant-directivity cornerhorns sit in adjacent corners, ideally situated where the forward axes cross in front of the listeners, as shown earlier in this document. They have the unique distinction of generating constant directivity through the whole audio band, all the way down through the midrange and into the bass. There is no other configuration that can do this.

The secret to their success is the fact that they use the walls as a large waveguide. The thing that most sound systems have to work around becomes the strength of the constant directivity cornerhorn. The midhorn is able to be snuggled back into the corner where the walls act as extensions to its flare. And the woofer uses the walls as its directional device in toto.

This use of the corner as a large waveguide is excellent. Of course, not every room is suitable for this configuration but where possible, it is of great potential benefit. Instead of fighting to reduce lateral reflections from adjacent walls, constant directivity cornerhorns make them become an integral part of the waveguide, one that is large enough to be useful down to low frequencies. It really is a natural solution, one that has no equal where room coverage, imaging and overall sound quality are concerned.



**seven π loudspeaker
constant-directivity cornerhorn**

Waveguides and constant directivity loudspeakers do a great job of creating a uniform reverberant field. They are able to provide spectral balance over a wide coverage area, reducing the seat-to-seat variance in sound quality. The "sweet spot" becomes large when using high-fidelity constant-directivity loudspeakers.

There is one problem that remains though. The reverberant field in a typical home listening environment extends from about 200Hz up. Below 200Hz, room modes take over and disturb whatever pattern the speaker would generate in an open space. In a sense, the room modes set the pattern at LF, and it is not uniform at all.

Room modes are caused by self-interference from walls and other large boundary reflections. The direct sound radiated from a woofer interacts with wall reflections and where the different path lengths create an odd multiple of a 180° phase shift, sound is cancelled at that point. Move a few feet and the listener will likely be out of the null, and into a position where bass is strong. These pockets of sound line up in different places depending on frequency. Room mode positions are completely dependent on room dimensions and woofer position.

There are actually three general regions where sound acts differently indoors. The lowest frequencies are in a pressure region, where wavelengths are so long they do not form standing waves. The modal region is just above that, where standing waves form nodes that are spaced far enough apart in both space and frequency to be distinguished from one another. At higher frequencies, the modes become more closely spaced, and eventually become indistinguishable from one another. This is called the reverberant region. The approximate frequency where the room transitions between modal and reverberant behavior is sometimes referred to as the Schroeder frequency.

In the lowest frequency region, a closed room is actually pressurized by the woofer. When the cone moves out, pressure slightly increases and when the cone moves back in, pressure slightly decreases. This causes the so-called room gain, 12dB/octave boost in the lowest frequencies. Of course, the seal isn't perfect and rooms are usually somewhat "lossy" so a full 12dB/octave is rarely realized.

In the modal region, boundary reflections interact with direct sound to form a very complex pattern of lobes and nulls. Room modes make "pockets" of hot and dead zones that can be visualized something like a checkerboard. Asymmetries in placement and room features make the pockets less ordered than a checkerboard grid, but the point is that lobes are scattered throughout the room, separated by nulls. Being on-axis in the modal region has no significance, because there is no "forward lobe".

As frequency rises, the physical separation between lobes and nulls becomes smaller and so does the distance between modes in terms of frequency. When the separation starts to become "fuzzy" and there are no longer clearly defined individual modes, the interactions start to become more an averaged reverberant field.

The difference may seem trivial at first, but it is actually very important. At low frequencies, reflections sum with direct sound coherently, forming very distinct patterns. Virtual sources created by wall reflections act similarly to an array of loudspeakers, creating directivity via phase interaction. The room modes formed are much the same as the vertical lobes and nulls created from interaction of vertically stacked sound sources. The only difference is that the effect isn't limited to the vertical plane.

At high frequencies, the potential phase differences between reflections and source become much greater, so the reflected sound field becomes an averaged field. It is essentially very dense interference that causes this. There is a way to make dense interference at LF too, making it act more like higher frequencies, smoothing the modes by averaging. The way to do this is to add point sources, which smoothes the sound field at LF much as it is at HF. This is the thrust of the multisub configuration.

Conventional wisdom suggests tight grouping of sound sources is best, essentially making a point source, at least as much as space will allow. This works outdoors because there are no reflections. Outdoors, it is best to put subs close together, and to have them near the mains where possible. It is generally preferable to stack subs on the ground rather than flying them, both because it provides boundary reinforcement and because it prevents a ground reflection.

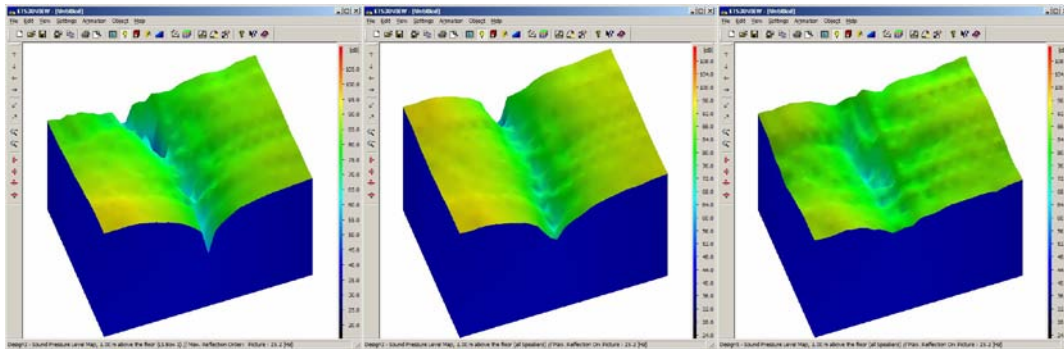
Indoors, the rules change because of room modes. A point source has little meaning at LF because of the large number and inconvenient locations of virtual sources from boundary reflections. The conventional wisdom that might suggest grouping subs together is exactly the wrong thing to do. Of course, in home hifi and home theater, many times groups of subs aren't considered but rather a single sub is usually used. This has the same problem, that room modes are worst when a single LF sound source is used.

It became popular in recent years to employ small speakers for the main channels and high-pass them so they do not have large excursion requirements. Bass is summed from all channels and sent to a single subwoofer. One potential benefit is reduced IMD because of relaxed bandwidth requirements on the mains. But there is a trade-off, which is that fewer LF sound sources make room mode problems worse. The solution is simple: Use more subwoofers and distribute them around the room.

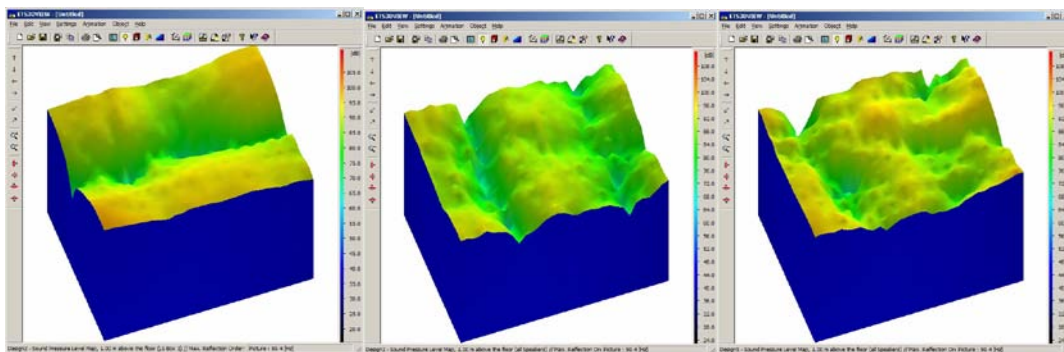
Optimal positions and numbers of subs are debated in some circles, but almost all agree that the more subs, the better. There is similar agreement that four subs provide enough smoothing that additional subs give nearly inaudible diminishing returns. The conclusion one should draw is that two subs are much better than one sub, and three or four subs are better still. But beyond that, not much additional benefit is gained.

Another observation one can make is the fewer the subs, the more important their placement becomes. If you have just one or two, their individual locations have a lot to do with the smoothness of the LF sound field. But if you have four or more, it almost doesn't matter where you put them as long as they aren't grouped together in the same spot. Spread them around.

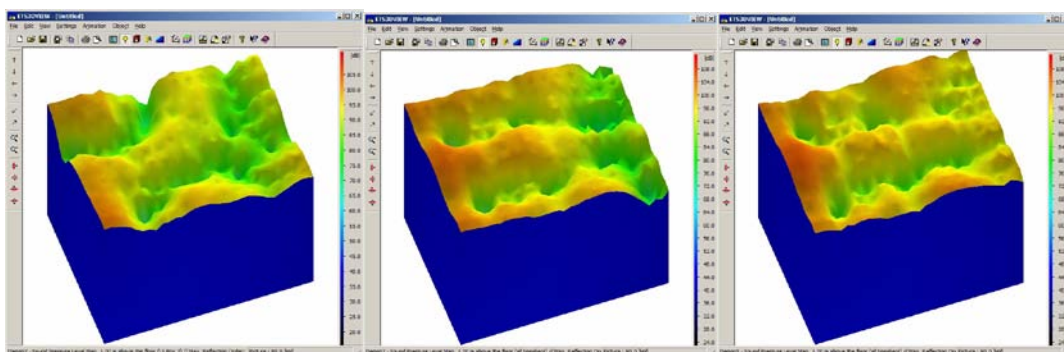
These charts show response throughout a room at various frequencies. Speakers are placed on the floor and measurements were taken at several points, all at listening height, about three feet up. The left chart shows a single loudspeaker placed near the left-front corner of the room, the center chart shows two speakers, one in each of the left two adjacent corners and the right chart shows four speakers, one in each of the four corners.



25Hz



50Hz



80Hz

The measurements of single and stereo pair are of full-range loudspeakers and the measurement of four speakers have mains in the front corners (shown left) and subs in the rear corners (shown right). The smoothing gained as LF sound sources are added is very clear.

As you can see from the charts, room modes are primarily axial, meaning they are formed by standing waves that line up between opposing walls, from side to side or front to back. This is pretty evident from the energy distribution in the room. There are other modes possible too, tangential (diagonal) and oblique (upper corner to lower corner). They have much less impact, and can almost be disregarded.

One thing that is not illustrated is the effects of axial modes between ceiling and floor. The charts shown on the previous page are all of listeners at one meter height and sound sources on the ground. However, in many cases, the main speakers are on stands, and this creates a different modal alignment in the vertical axis.

The frequencies of the vertical modes are usually towards the higher end of the modal range, often above 100Hz. This is largely because the ceiling height is usually the smallest dimension of the room.

There is also self-interference from the path length differences between the direct sound and the reflections off the floor and the wall behind the speakers. While technically different than room modes, the notches formed from path length deltas can be mitigated in much the same way as room modes, with multiple sound sources. Like the vertical modes, these are usually at the higher end of the modal region, above 100Hz.

One method that is very effective at smoothing the higher frequency modes, as is employed in π Cornerhorns, is to overlap a vertically stacked midrange and woofer in this region. Of course, this requires a relatively low woofer-to midrange crossover point, with woofer and midrange blended together, sharing the band up to 300Hz or so. Above that point, it is desirable for the loudspeaker to act as a point source.

Another method, useful when the main speakers are matched-directivity two-ways on stands, is to employ "flanking subs". Like more distant subs used in a multisub configuration, these aren't used to provide LF extension as much as they are there to provide modal smoothing. They perform both functions, of course, extension and smoothing.

Mains are run full-range, without high-pass, blended with the flanking subs and sharing the modal range with them. The distinguishing feature of flanking subs is their proximity. They are placed below mains, often behind and slightly to the side and are usually just a few feet from the mains, low-passed at 80Hz to 120Hz or so. Flanking subs run this high could be localized if placed further away but their near proximity to the mains makes them blend seamlessly.

References

- Wenthe and Thuras, "Auditory Perspective - Loudspeakers and Microphones", Electrical Engineering, Vol. 53 (1934 Jan) also, J. BST, Vol. XIII, No. 2 (1934 Apr) and J. AES, Vol. 26, No. 3 (1978 Mar)
- Beranek, "Loudspeakers and Microphones", J. ASA, Vol. 26, No. 5 (1954)
- Schroeder, "Frequency-Correlation Functions of Frequency Responses in Rooms", J. ASA, Vol. 34, No. 12 (1962 Dec)
- Damaske, "Subjective Investigation of Sound Fields", Acustica, Vol. 19 (1967)
- Keele, "What's So Sacred About Exponential Horns?" AES, Pre-Print No. 1038 (1975 May)
- Linkwitz, "Active Crossover Networks for Noncoincident Drivers", J. AES, Vol. 24 (1976 Jan)
- Hilliard, "Historical Review of Horns Used for Audience-Type Sound Reproduction", J. ASA, Vol. 59, No. 1 (1976 Jan)
- Keele, "Low-Frequency Horn Design using Thiele/Small Driver Parameters", AES, Pre-Print No. 1250 (1977 May)
- Smith, Keele and Eargle, "Improvements in Monitor Loudspeaker Design", J. AES, Vol. 31, No. 6 (1983 Jun)
- Schroeder, "Progress in Architectural Acoustics and Artificial Reverberation", J. AES, Vol. 32, No. 4 (1984)
- Benade, "Equivalent circuits for conical waveguides", J. ASA, Vol. 83, No. 5 (1988 May)
- Olson, "Acoustical Engineering", Professional Audio Journals, Inc., Philadelphia, PA (1991)
- Geddes, "Acoustic Waveguides - In Practice", J. AES, Vol. 41, No. 6 (1993)
- Schroeder, "The Schroeder Frequency Revisited", J. ASA, Vol. 99, No. 5 (1996 May)
- Beranek, "Concert and Opera Halls, How They Sound", ASA, New York (1996)
- Hughes, "A Generalized Horn Design to Optimize Directivity Control & Wavefront Curvature", AES, Pre-Print No. 5016 (1999 Sep)
- Parham, "Speaker motors and passive crossover filters", π Speakers, Web (2001)
- Welti, "How Many Subwoofers Are Enough", J. AES, Vol. 50, convention paper 5602 (2002 Jun)
- Welti and Devantier, "In-Room Low Frequency Optimization", J. AES, Vol. 51, convention paper 5942 (2003 Dec)
- Welti, "Subjective Comparison of Single Channel versus Two Channel Subwoofer Reproduction", AES, preprint 6322 (2004 Oct)
- Welti and Devantier, "Low-Frequency Optimization Using Multiple Subwoofers", J. AES, Vol. 54 (2006 May)