Hyperbolic Horn Physics and Design

(Rev. 03/19/17 – Roy Minet)

The Problems of Reproducing Low Audio Frequencies

Physics requires a large structure of some sort in order to do a truly good job of reproducing the long wavelengths of very low audio frequencies. Anything else is a compromise (e.g., resonant "bass reflex" systems) and still will not reach frequencies as low as 15 Hz. Also, design problems multiply when a single transducer has to handle more than two or three octaves of the eight or nine octave full audio range. It is especially advisable to split off the very low end for handling by a subwoofer. Another advantage of doing so is that the very low frequencies are not directional and no stereo imaging will be lost if the low bass from the left and right channels is combined. Thus, only one physically large structure is required instead of two.

Quality bass drivers are manufactured in sizes of 12" to 18" diameters. At low frequencies, a moving piston of this diameter must move a large distance in order to transfer the required power into the low acoustic impedance of the listening room's air. Design limitations hold the cone excursions of bass drivers to less than is required, and to maintain good efficiency and low distortion, excursion distances need to be kept even smaller.[1] They are inherently high impedance devices which are better suited to driving higher acoustic impedances (than the room air) with larger pressure differentials.

Two Design Approaches

One brute force solution would be to use an array of multiple drivers mounted close together in an infinite baffle (say, a wall). As long as no driver is as much as a quarter of a wavelength (at the highest frequency to be handled) from any other, their outputs will combine in phase and they will tend to behave as would a single driver having the total area of all their cones. Excursion distance and power handling requirements are thus shared and reduced by the number of drivers employed. This actually is a viable solution. However, it is far from the best. Efficiency cannot possibly be better than 50% (as half the power is radiated and lost on the back side of the wall), and is likely down in the 25% to 40% range. Still, power transistors make lots of amplifier power available at an affordable cost and can make this design work satisfactorily.

A more elegant design would use an "impedance transforming device" which could better match the high impedance of the driver to the low impedance of the room. In the electrical realm, such devices are appropriately called "transformers." One side of a transformer

---

[1] Drivers built for long excursions have been made for "acoustic suspension" woofer designs. Their efficiencies are extremely low and linearity over such long excursions is somewhat problematic (can lead to higher distortion). Such designs do not attempt to achieve low-distortion bass which is flat all the way down to 15Hz!

matches a high impedance circuit (high voltage/low current) to a lower impedance circuit (lower voltage and higher current) on the other side (or vice versa) with very little power loss in between. The transformer's analog in the acoustic realm is called a "horn." The most familiar example probably is the simple megaphone. A properly designed horn can transform high impedance sound (large pressure differentials and a small volume of air movement) at its small end to a lower impedance (smaller pressure differentials and a larger volume of air movement) at its large end.

If we bolt a good bass driver to the small end of a properly designed horn and bolt the listening room onto the large end, we find that the driver can couple much more power into the horn load without large cone excursions. As the sound waves travel out through the horn, their area gradually increases and they are transformed to a lower impedance that better matches a room, again with almost no power loss. As an extra, added bonus, we find that much more than half of the driver's power is coupled into the horn's high impedance than is coupled into the relatively low impedance of the air behind the driver. Instead of 50% maximum theoretical efficiency, we can achieve an 80% theoretical maximum. Actual efficiencies of 55% and 60% are realistic.

"Proper" Horn Design

A good fundamental design begins by matching the area of the small end of the horn, called the "throat," to the effective area of the driver. To calculate this measure the diameter of the stiff cone and include about a third of its surrounding suspension which supports the cone, but allows it to move. (Driver selection will be discussed later.) My driver has an effective area of 135.3 square inches (a circle with a 13.125" diameter). If this initial area increases gradually and smoothly forever, the horn will present a constant impedance to the driver and there will be no resonances. Any sudden changes to the area will reflect a portion of the sound energy and introduce resonances which will ruin the otherwise flat frequency response (and phasing).

Because of the impracticality of building an infinite horn, we'll unfortunately have to chop it off somewhere. That will be the large end, called the "mouth." So, even if we make the horn very gradual and perfectly smooth, we are necessarily going to have a big discontinuity in the area at the mouth where the horn abruptly ends and dumps out into our listening room. Fortunately, the horn will still behave as though it were infinitely long if the mouth is large, at least a quarter of the lowest wavelength we intend the subwoofer to reproduce. Since the plan is to go down to 15Hz, a quarter of that wavelength is about 224 inches (almost 19 feet); ouch; impractically large!

We will have to resort to a "trick" to make the horn think its mouth is larger. If the horn mouth opens out into the room in a corner smack up against a wall, the wall will support the emanating wave-fronts and make the horn behave as though there were another identical horn right up against it; think of it as the acoustic image of the horn as reflected by the wall. Furthermore, if we move the opening up to the ceiling, we can gain the advantage of the first two horns being reflected by the ceiling. So, we have a real horn plus three virtual horns which quadruple the area and double the dimensions. An actual mouth opening of 112 inches is still quite large. However, there is one more trick that we will pull out of the bag later and that will enable us to decide on a mouth opening of about 55 inches as well as to completely avoid any resonance problems.

So, the horn will have to transition gradually and smoothly from a 13.125 inch circle to a quarter of a circle with a 55 inch radius. The area must grow from 135.3 square inches to a little over 2,400 square inches. It turns out that how you get there is very important and, in this case, much more than half the fun.

Every horn has a cutoff frequency. Below the cutoff frequency, the horn no longer "loads" the driver. That is, it no longer performs as an impedance transformer and the driver is then just pushing on the low impedance room air. The cutoff frequency is determined by how rapidly the area increases, especially at the beginning of the horn where the pressure differentials are large. The more gradually the horn flares out, the lower the cutoff frequency will be. It was quickly learned in the early days of horn design, that for a given cutoff, a much shorter horn with an exponential flare would work as well as (actually better) than a longer conical one (megaphone shaped). For an exponential horn, the area increases along the horn axis in accordance with the below expression:

$$A = A_0 e^{fz}$$

Where A is the area at any point along the axis z;

A_0 is the throat area at z = 0;

And f is the flaring constant which sets the cutoff frequency.

An exponential horn flares at a smoothly increasing rate as sound moves from the driver toward the mouth. The hyperbolic horn adds the flexibility to fine-tune the shape and achieve even better characteristics. For a hyperbolic horn, the area increases along the horn axis in accordance with the following expression:

$$A = A_0(Cosh\ \Theta + T\ Sinh\ \Theta)^2$$

Where A is the area at any point along the axis z;

A_0 is the throat area at z = 0;

Θ = z/z_0;

z is the distance along the horn axis and z_0 is a reference distance (flaring constant);

T ≥ 0 and is a constant which adjusts the horn shape;

Cosh is the hyperbolic Cosine function, Cosh Θ = ½ (e^{Θ} + $e^{-\Theta}$);

Sinh is the hyperbolic Sine function, Sinh Θ = ½ (e^{Θ} - $e^{-\Theta}$).

Yes, the hyperbolic horn is a bit more complex than the exponential horn, but it's not quite as bad as it may appear and the added complexity is worth it. (Not to worry anyway, as we certainly plan to have a computer crunch all the numbers.) The constant, T, allows us to tweak the shape of the horn to obtain better performance in the critical bottom end of the range near the cutoff frequency. First note that when T=1, the hyperbolic equation reduces to exactly the exponential one. As T is increased above one, the horn shape becomes closer and closer to conical (megaphone). It is the range 0 ≤ T < 1 that is of particular interest.

A problem with the exponential shape is that the acoustic impedance (and therefore the frequency response) begins rolling off too far above the cutoff frequency. For a 15Hz cutoff, the response drops to half between 17 and 18 Hz. Setting T to a value less than one will cause the impedance (response) to actually rise from its nominal value to a peak and then fall very quickly as the cutoff is approached. Obviously, a large increase is not desirable, but a small one helps to hold up frequency response at the very bottom. I chose T = 0.59, which results in a rise of only 5%. This gentle "peak" occurs just below 22 Hz after which the impedance begins its decline. It is still hanging in at 90% when we get down to 16.5 Hz and 75% at 15.75 Hz, but then crashes. It is down to 39% at 15.15 Hz. Thus, the shape adjustment lets us maintain and support our frequency response to within a quarter of one Hz above the cutoff frequency. The below graphic may aid understanding.

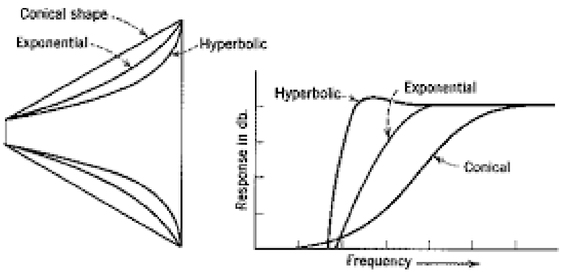

FIG. 4–14. Effect of horn shape on cutoff characteristic.

That is as far as we go with horn physics. Anyone wishing a more comprehensive and in-depth treatment of the subject is referred to an excellent two-part article by BjØrn Kolbrek titled, "Horn Theory: An Introduction," published during 2008 in audioXpress.

Solving Some Remaining Problems

Now, for the bad news: In order to have a 15Hz cutoff, we need to make the flaring constant or reference distance z_0 = 142.71 inches. This is a very gradual flare indeed and means that the horn length must be approximately 16 feet. The rear side of the wall in which the horn is to be mounted is a garage, so a 16-foot horn along an upper corner would not be completely out of the question. A more compact structure would be less obtrusive.

There is another small problem with using a horn. The air itself is non-linear. This is not of consequence as long as pressure variations do not exceed a very small percentage of atmospheric pressure. But as pressure variations increase, the air's inherent non-linearity will distort audio waveforms and introduce distortion. We have just worked hard to support the driver coupling more power into the air with larger pressure differentials at the throat of a horn.

However, there is a basic change that can be made which not only preserves all the benefits advertised so far, but also makes three significant improvements. Instead of one large horn 16 feet long with a 2,400-square-inch, quarter-circle mouth, truncate the horn at the point where its mouth is 300 square inches (about 17.3 inches square) and build eight of them. Mount the eight square horn mouths at the ceiling corner in a closely-spaced array that approximates a quarter-circle. The eight-horn array, with reflections in the wall and ceiling, act like 32 paralleled horns and achieve all design goals. The length of the eight horns will now be about five feet, including the drivers mounted at the rear.

The three big improvements are: first and most obviously, a more compact structure; second, pressure differentials have been reduced by a factor of eight so distortion, even at high power, is not a concern; third and finally, possible problems with resonances have been removed. The last may not be immediately obvious. The primary and lowest resonance will be at the frequency whose wavelength is double the length of the horn. For a single large horn, that occurs at about 35 Hz, right in the middle of the range we intend for the subwoofer to handle. With the shorter horns, the resonance is pushed up to about 122 Hz. This is well above the crossover frequency as the subwoofer will only handle frequencies below 70 Hz, so to whatever extent it exists, this resonance simply will never be excited by the bass drivers.

How To Build Such A Thing

The original plan was to build a mold which smoothly transitions from the 13.125” diameter circular throat to the 17.3” square mouth, then make the horns using polyester resin. That plan was abandoned because of concerns regarding the rigidity and durability of the polyester, the difficulty of removing the horns from the mold without damaging either, and the health risks associated with working for prolonged periods with dangerous VOCs (Volatile Organic Compounds).

The final plan was to build the structures using (primarily) 3/4” MDF (medium density fiberboard). MDF is easy to work with, has no grain and can provide the necessary rigidity. The horns were designed to have two sections for two good reasons: first, to allow the massive thing to be broken down to facilitate being able to move it to a new location (which has actually been done); and second, to make the complex gradual transition from a circular to a square cross section in the rear section as easy as possible. The front section then is more simply an expanding square cross section. Two thicknesses of ¾” MDF were laminated to make 52” x 52” mounting and support panels for each end of each section. The rear of the rearmost panel has the eight 13.125” diameter circular openings that match the drivers. The front 1.5” thick panel, as well as the two 1.5” thick panels for the middle, each have eight square openings of the correct sizes. The two middle panels will be bolted together using seven half-inch bolts to hold the two horn sections together.

Rear Horn Sections

A Java program was written to very precisely calculate the cross-sectional area at quarter-inch intervals along the axis from the driver’s 135.3 square inches to the mouth’s 300 square inches 54” away from the driver end. The smallest square structure that could accommodate mounting the circular drivers has inside dimensions of 13.125” by 13.125”. Thus, the area at the front end will be a little over 172 square inches. That fixes the length of the rear horn section at a little over 27 inches and means that the rear section will be a bit more than half of the total length.

The first step is to smoothly transform the square front opening to an octagon at the rear by means of MDF inserts in each of the four corners. The second step is to turn the octagon at the rear into a 16-sided regular polygon by casting polyester “wedges” in each of the eight octagon corners. These smoothly transition the 16-sided polygon to an eight-sided structure about half way through the rear section. The construction photos will clarify this description. The minor transformation from a circle to the 16-sided polygon is accomplished through the 1.5” thickness of the rear panel. The rear section now has the correct areas at its front and rear openings, and it smoothly transitions from circular to square. However, its shape does not accurately match the required hyperbolic equation for the area in the middle. Its area is too large everywhere except at the ends.

The Java software was modified to calculate the actual area of the structure as built (again at quarter-inch intervals) and then to subtract the required hyperbolic area from it. This showed the amount by which the area needed to be reduced at each point along the axis. The software was then modified to calculate the dimensions of inserts that could be mounted along four of the sides to make the cross-sectional area match the required hyperbolic shape everywhere. A mold was then constructed to cast these complex shapes from polyester resin. Access to CNC (computer numerically controlled) equipment would have been very helpful, but lacking that, templates were printed on an inkjet printer to assist in making all dimensions and shapes accurately.

Front Horn Sections

Compared to the rear section, the front horn sections were easy. These are flaring square sections which mate up with the 172 square inches of the rear sections and have 300 square inch mouths. As before, the areas are only accurate at the ends and must be corrected in between. The required corrections are very much smaller in the front sections. They were calculated in a similar way by the software and consist of carefully shaped pieces of MDF cemented into each of the four corners of each of the eight front horn sections. These are somewhat hard to see (especially when painted black), but are visible in the photos.