The J-Low

By Nelson Pass and Dana Kruse, (c) 2003 Pass Labs

Intro

The aesthetic appeal of a single wide-bandwidth loudspeaker driver is obvious in its simplicity. There are no crossover networks, little or no phase shift versus frequency, and a single acoustic source location. As a concept, what's not to like? What could be more perfect than a nice little cone that could do it all, from the deep bottom end to tinkle beyond human hearing?

Unfortunately there are good reasons why such speakers are uncommon – they are very difficult to make, relying more on hard work and good taste than your average loudspeaker part. They have to be a lot more than a stiff piece of paper, plastic or metal.

The "rigid piston" model of a mass-controlled loudspeaker diaphragm works remarkably well over a specific range of frequencies. For a given physical displacement, the acoustic output of a loudspeaker cone increases with the square of the frequency as long as the wavelength of the sound is large compared to the cone size. At approximately the frequency where the wavelength equals the circumference of the cone, this output levels off.

What does this mean? That for a given excursion distance a rigid circular type of loudspeaker will have a frequency response which rises 12 dB per octave until the wavelength is three times the diameter of the circumference. For a 12 inch speaker cone, this frequency is about 400 Hz, above which the response should be relatively flat - for a given excursion distance.

That's very nice, except that the excursion of the cone is not independent of frequency. The electromotive force provided by current traveling through a voice coil in a magnetic field might be constant, but this force creates acceleration that must be translated to velocity, which is not the same thing, and then excursion, which is twice not the same thing. At the risk on mentioning calculus, we would say that the velocity of the moving assembly of a loudspeaker results from the integration of the force applied, and further, that the excursion is the integration of the velocity.

Well, what does that mean? Assuming that your loudspeaker voice coil and cone are much heavier than air (and they are) that the excursion falls off at a rate of 12 dB per octave as the frequency increases.

I would like to have been there on that happy day (probably at Bell Labs) when somebody noticed that the acoustic output rising at 12 dB/octave nicely cancels the excursion falling off at 12 dB/octave.

After they ran out of champagne, this same scientist probably saw the flaw – this only works up to the frequency where the wavelength is larger than the circumference. Of course, perhaps back then they weren't quite as critical as we imagine ourselves today, and were pretty happy with response up to a couple thousand Hertz.

There was another fly in the ointment in that the excursion of the cone had to increase by a factor of four for every lower octave of output. At some arbitrarily low frequency this becomes a problem, somewhere less than a couple inches.

So we end up with a piston model of the loudspeaker that for a given size is bracketed on the bottom frequencies by its excursion limits and on the top by a leveling off of the acoustic resistance. As we make the cone bigger, we can play louder at low frequencies, but we limit the high frequencies by the same proportion. As we make the cone smaller to get the high frequencies, we find that the cone can't travel far enough to give us undistorted bass. Alas.

If we want to stick with something resembling the mass-controlled piston model, we have no choice – we have to divide the job into multiple drivers. A one inch diameter cone that would do a nice job with the bells on Dark Side of the Moon (hint: Pink Floyd) is not going to deliver on the bottom end.

Well, that's OK, and usually that's how it's done. But it's not the only way. If we ignore the rigid piston model, we can imagine cone materials that are flexible in such a way that the center of the cone is free to vibrate at a higher acceleration than the outer edge of the cone. Actually, I believe the Bell Labs guys thought of this the next day after the party.

Of course the Egyptians, bless them, thought of paper first, but it so happens that carefully wrought paper materials have close to the right properties of not only being very lightweight, but also having the right balance between stiffness and flexibility. Combined with a certain amount of intrinsic damping, paper cones do a remarkable job of overcoming these problems.

They aren't perfect, but the audiophile attraction is there. Of the products on the market, most make some sort of effort at decoupling the low frequencies from the high. In my opinion, the ones that do the best job are the smaller ones, which is no surprise. This means they need some sort of help on the bottom end.

Kent English at Pass Labs has a job description that includes acquisition of interesting drivers (He finds them, I sign the checks). He bought a pair of Jordan JX92S's, which is a full range cone speaker with about a 3.5 inch diameter cone and some sort of metallic coating on the cone. Well aren't they cute, sez I, and they don't cost much, so one fine Saturday we put them into boxes and started playing with them.

We were most surprised. I would call them flat to 20 KHz, and remarkably, they made it down below 50 Hz in a modest box. Figure 1 shows the wide-band response curve at 1 meter, and Figure 2 shows low frequency detail with the driver in a 3 cubic foot box.

They sounded so good that we started cranking them up, and immediately ran into the distortion from the high excursion in the bass. Alright, so they didn't play that loud, but they were still very pleasant to listen to over the long term.

And so it remained, listening at low levels in the dead

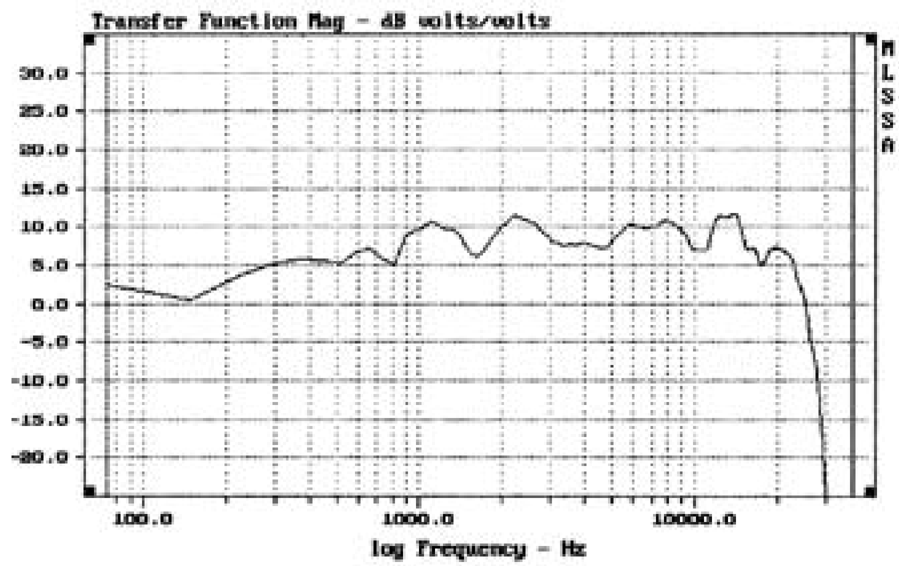

FIG 1 Jordan JX92S Full Range Response @ 1 Meter

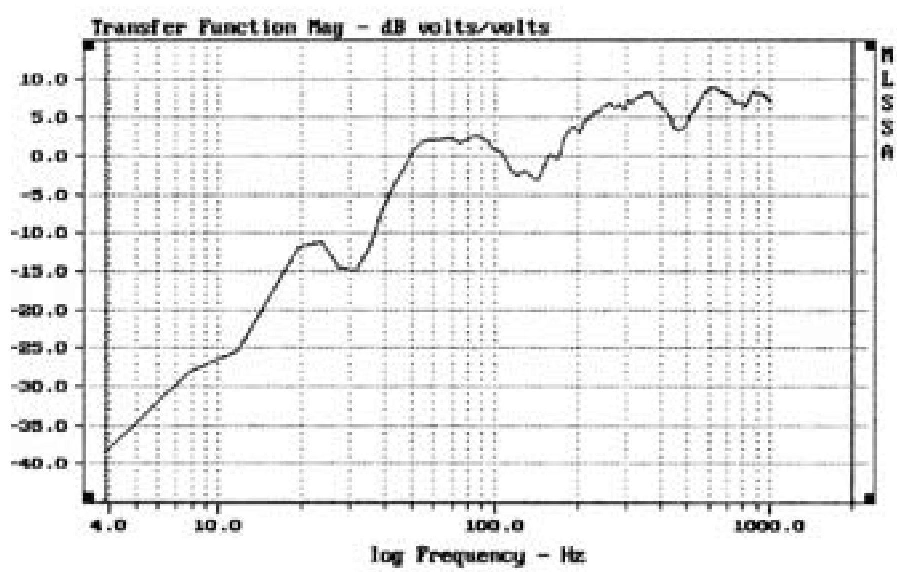

FIG 2 Jordan JX92S Near Field Bass Response

of night, until Dana Kruse showed up. Now Dana is not your average audiophile. Supposedly a successful architect in Seattle, he has spent numerous of his vacations in Foresthill sitting on the production line building amplifiers, presumably for therapy. Dana Kruse is also dangerously witty in an abstract way, so that almost none of the remarks or scatological diagram titles he has contributed will appear in this article.

It happens that he also owns a pair of Jordan JX92S's and we decided that for this vacation we would do something with them. Designing at the table saw is a particular specialty of Dana's, and so we went straight out to the woodshop.

We needed more output level on the bottom end for these speakers, and horns looked like the best way to get it. Not front loaded, which would obscure the high frequencies, but some nice rear loaded boxes – something which would give us intimate access to the front wave but back it up with a little authority on the bass.

Horn loading is a well-understood science. As Leo Beranek points out in his classic book on acoustics, a horn is an acoustic transformer, turning a small diaphragm into a big one without cone resonance. The most common shape for a horn is an exponential curve, where the surface area down the length of the horn is given by the equation:

$$S = T * e ^\wedge (M * D)$$

Where S is the cross-sectional area at the distance from the throat D whose cross-sectional area is T. M is the flare constant given by:

$$M = 4 * pi * F / C$$

Where F is the cutoff frequency of the horn, and C is the speed of sound.

As a practical matter, the cutoff frequency is where you have no resistive output for the horn, so a practical horn operates at - 3 dB at about 1.4 times the value of F. There is another consideration for the horn, being the requirement that the mouth of the horn should have a circumference greater than the wavelength of the lowest frequency to be amplified or have the equivalent area.

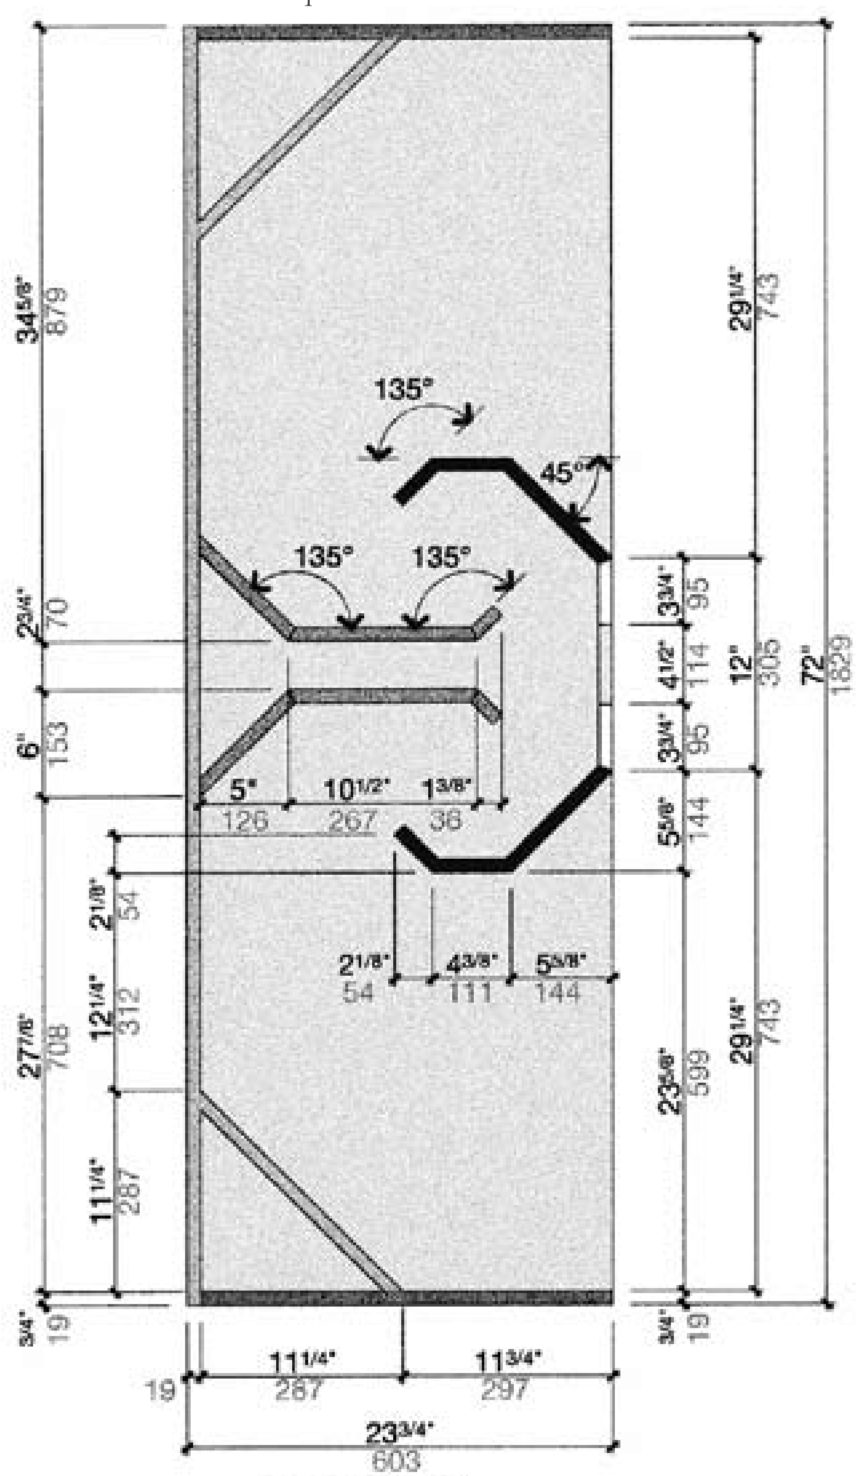

FIG 3 Side View

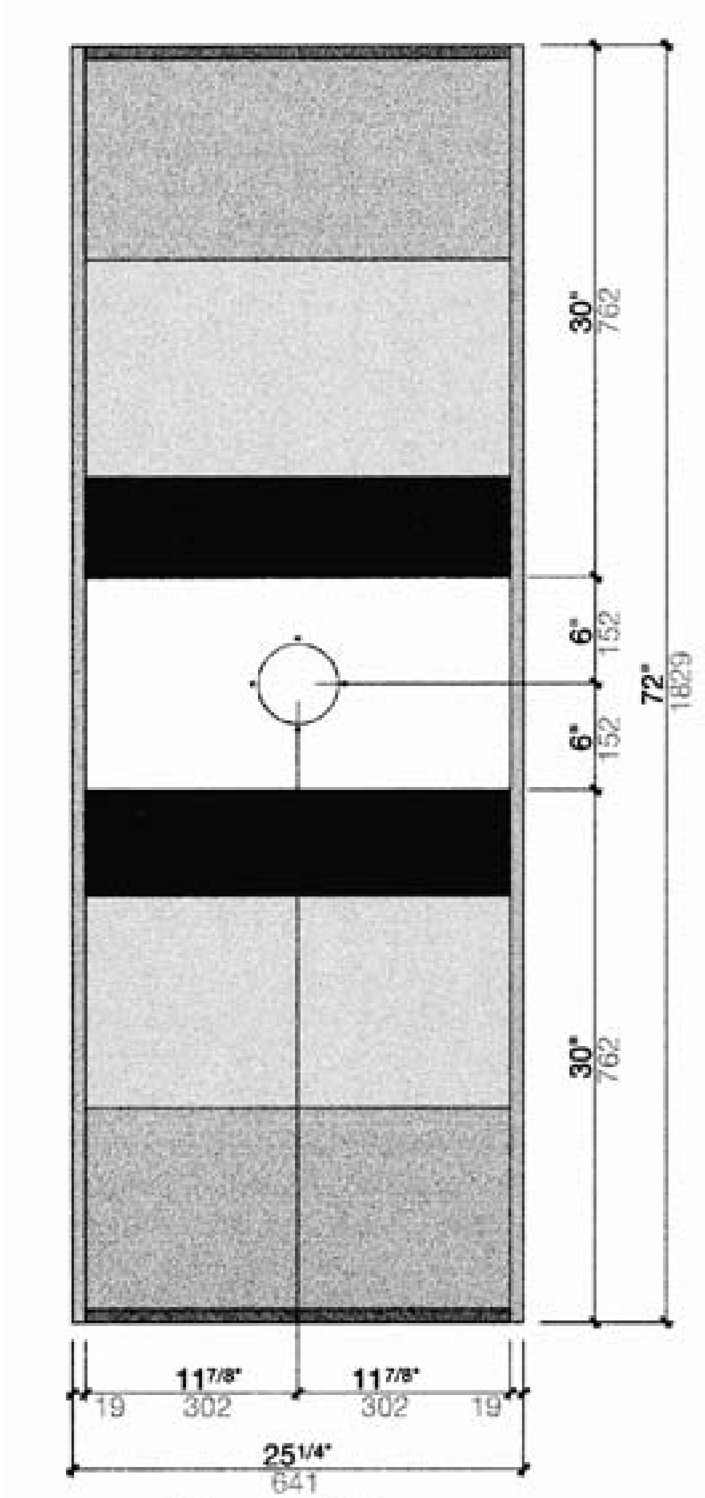

FIG 4 Front View

It's pretty straightforward - the lower the frequency you want to take your horn, the bigger it will have to be. The following table illustrates the incremental distance from the throat it takes to double the cross-sectional area for various cutoff frequencies:

60 Hz	12 inches
50 Hz	14 inches
40 Hz	18 inches
30 Hz	24 inches
20 Hz	36 inches
15 Hz	48 inches

Of course a cutoff of 15 Hz means that the horn would be usable down to about 20 Hz. The mouth area required follows a similar pattern:

60 Hz	37 sq ft
50 Hz	54 sq ft
40 Hz	84 sq ft
30 Hz	149 sq ft
20 Hz	336 sq ft

Designing at the table saw is also a well-understood science. You look around and see what kind of pieces of wood you have. Fortunately I had quite a few sheets of MDF with a nice oak veneer, and so we sketched up a nice rear loaded design, checking the expansion through trial and error with ruler and compass.

Dana comments:

After cutting the pieces on the table saw, we introduced some beer into the equation, resulting in an alternate geometry as described below:

S = T * e ^ (M* G* D) + :) + :)

With the new found efficiency this formula afforded us, we were able to take a little time away from the dust and glue to prepare an impromptu meal of Tostitos, bread, wine, bread, DiGiorno, bread, beer, bread, steak, bread and parmesan, chased with bread, champagne, wine, beer and ice cream. Reflecting upon what we had accomplished the next morning at our 8 AM Philip Kaufman screening, it seemed that the flurry of snappy banter and the complete absence of flesh-meets-whirling-steel, added up to a pretty fair day in the shop. With coffee, peanut butter, bread, honey, bread and Equal in our glue encrusted and splinter-pricked hands we toasted to a job well done and looked forward to what revelation this day would bring."

The design is pretty simple, as you can see from Figures 3 and 4. We were tempted to go the full size of the MDF sheets at eight by four by four feet, but decided to limit ourselves to something that would actually fit through the listening room door. This placed a serious limit on how low a frequency we thought we were going to get away with, and we decided on about a 35 Hz cutoff taper (which would take us down to about 50 Hz in reality) and hope that the 20 square foot mouth areas of the combined speakers would get some help from the floor and back wall in a 30 x 30 foot room.

Dana later generated a full-boat set of diagrams in glorious color, with cutting and assembly instructions which is to large to reproduce here. You can download the pdf file from www.passdiy.com

Figure 5 is a picture of Dana Kruse standing in front of the right channel loudspeaker. What appears to be a cute little on/off button at the front of the speaker box is the Jordan JX92S's. We conducted our listening tests with a 40 watt balanced version of the Zen Lite, which is the box you see with the light bulbs on top.

Going back to Figure 3, you see a chamber behind the driver that opens up into upper and lower horn throat areas. This chamber before the horn throat helps to form a low pass filter that reduces the amount of high frequencies that will pass through the horn. This is an important item, as you don't want to be listening to the rear wave above 100 Hz – it will interfere with the front wave. This acoustically capacitive chamber should also be filled with absorbent material, such as Dacron, wool or fiberglass. We chose Dacron.

You can easily tune this acoustic low pass filter by altering the volume of the chamber behind the driver, and in our case we played with the volume by altering the density of the Dacron.

The result was simple and dumb, but very effective. The JX92S's didn't go any lower in frequency than before but they picked up about 10 dB of gain centered in the 70 Hz region, and in fact became bass heavy. This was what we were looking for. Figure 6 shows the near field response of the result. Figure 7 shows the smoothed response full range as taken from the listening position. Between measurement and listening we concluded that the bottom end was about 6 dB too much at about 70 Hz.

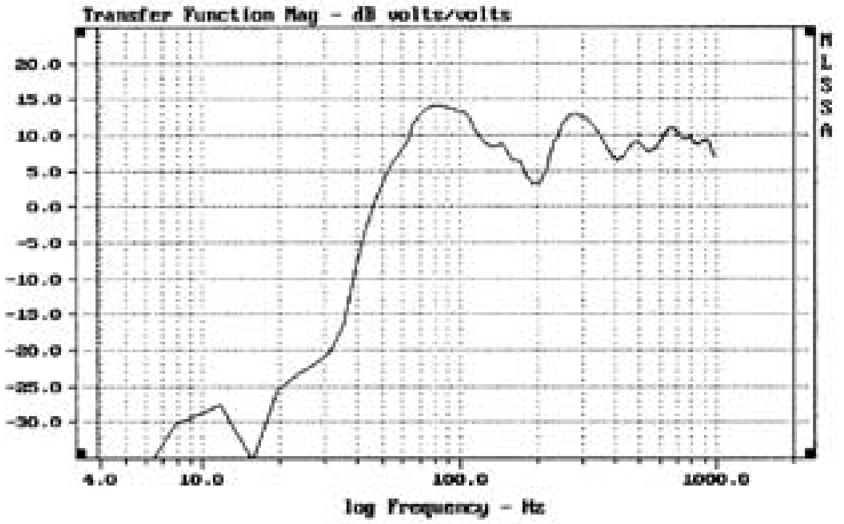

FIG 6 Near Field Bass Response in Horn

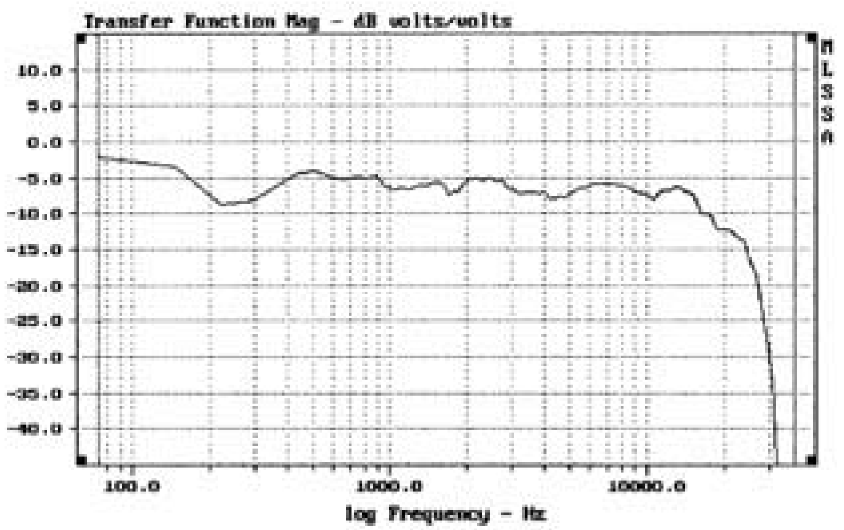

FIG 7 Listening Position Full Range Response

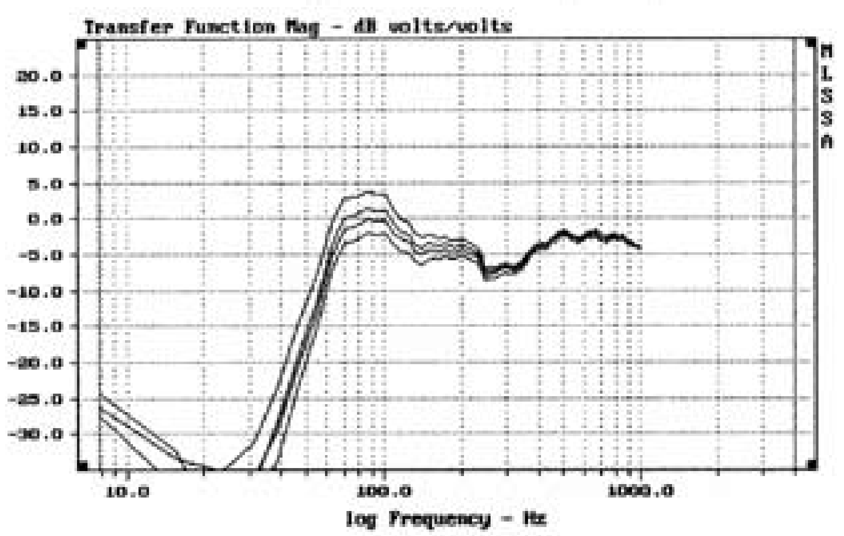

FIG 8 Listening Position Bass With High Pass Filters

This put us in a fine position to apply a high-pass filter to the system. We played with single-pole high pass filters at 50, 100, 130 and 180 Hz, and settled on 100 Hz as the most satisfying. Figure 8 shows a more detailed response of the system as seen from the listening position with the four different filters applied.

Previously the speakers delivered decent bottom end, but they couldn't play loud. Now with rear loaded horns and a high pass filter we get a similar response curve, but can play about 10 dB louder. In addition, the bass response also picked up a qualitative improvement in it's dynamic quality. I have remained happy listening to them for months now.

Dana had to go back home to design skyscrapers for Microsoft, but Kent is continuing to acquire full range drivers and we have started construction on a larger set of horns based loosely on this design, but which must be assembled on site, as they will not fit through any ordinary doors. We will be testing these drivers shortly, starting with the Mangers, and a follow up article will appear.