

White Paper

How to Achieve Accurate In-Room Quasi-Anechoic Free-Field Frequency Response Measurements Down to 10 Hz

**By
Jeff Bagby
1/3/2014**

This paper is based on a presentation given at the “Loudspeaker Design Workshop” in Kokomo, IN, on December 14, 2013 by Jeff Bagby. It presents a new methodology for merging near-field and gated-impulse far-field (quasi-anechoic) measurements that result in a much more accurate final summed frequency response than the traditional hard-splice method produces.

Preface –

In taking loudspeaker measurements very few builders have access to an anechoic chamber. Some builders have resorted to elaborate set-ups outdoors, or have tried to take complicated ground plane measurements on their driveway. However, these methods have their shortcomings and issues too. In the end, it is often felt that the most useful and accurate measurement process involves taking a gated (quasi-anechoic) far-field frequency response and then splicing it with the near-field response of the driver to obtain a summed response that is an accurate combination of the two. However, what I have found is that the traditional method of splicing this near-field and far-field data is not, in fact, accurate, and tends to exaggerate the response amplitude in the low-end. This inaccuracy can lead to errors in the design of crossover networks if the data around 200 Hz is not accurately represented.

In this paper I will present a new method of merging near-field and far-field data that results in an accurate frequency response to as low as 10 Hz. This is accomplished by using specific methods to capture near-field and far-field data and then using a new Excel Spreadsheet loudspeaker design tool called the “FRD Response Blender” created by Charlie Laub and Jeff Bagby to merge the measurement file data and extract the minimum phase response. When following the method that I will describe in this paper the new combined frequency response will result in more reliable crossover modeling and a more accurate frequency response presentation.

Before detailing this new method for merging near-field and far-field data it is important to review the basics of loudspeaker measurements and the correct methods for performing and understanding this data.

The Impulse Response –

The most effective measurement method is to use an Impulse Response, and because of this, nearly every measurement system will use it (or a very short Sin sweep similar to it). The Impulse Response is a loudspeaker’s response to a sharp, narrow, short pulse that contains a uniform distribution of all frequencies in the audio band. Due to its length of time, which is only a few milliseconds in duration, it is less likely to excite room modes and reflections than using longer term signals like Pink Noise or Warble tones.

The Impulse Response is a time domain response plotted along a time scale. This allows you to see the arrival of later reflections being picked-up by the microphone as ripples in the Impulse Response plot. By examining the impulse response you can “window” the data to remove these reflections. Once the software has captured the speaker’s Impulse Response the frequency response, phase response, cumulative spectral decay, and energy storage data are calculated from it using a Fast Fourier Transform. All of this information is contained in the single impulse response. The following discussion of frequency response measurements will be based on using the Impulse Response method.

Ground Plane Measurements –

Ground plane measurements can be effective, but do create some concerns. To perform ground plane measurements we place the mic and the speaker on a flat, smooth, highly reflective surface, with the microphone lying on the ground. The speaker will sit on the ground at least the minimum far-field distance from the mic (to be defined below). You may need to tilt the speaker downward toward the mic in order to maintain the proper axis during the measurement. The reflection from the ground creates a virtual mirror image of the speaker, so ground plane measurements are always 6 dB higher than half-space measurements due to this reflection.

One of the problems with ground plane measurements though are due to the fact that the virtual baffle is now twice as large as it really is and this will change the low frequency response some. Another problem is if the surface is not perfectly flat or very highly reflective, these will alter the response too. The advantage to the ground plane response is that if performed outdoors can provide response that is accurate to 20Hz or below. There is also the problem of ambient noise when measuring outdoors. In my neighborhood it just does not seem to lend itself to taking good measurements, and no one wants to take measurements in bad weather. There are techniques that can be used to get around some of the issues with ground plan measurements but they require some very precise distances and placement of the microphone and speaker to be accurate that other methods don't require. For this reason I tend to stick with near-field and far-field measurements indoors and merge the data.

Far-Field Measurements –

To begin with we need to understand what a far-field response really is. To be in the far-field is not as far as you may think. **You are effectively in the far-field when your microphone is at a distance that is 3-5 times the radiating diameter of the driver.** So, for a 6.5" woofer that has a radiating diameter of approximately 5" you will be in the far-field once your mic is placed 18" or so out from the driver. In order **to make sure that you are picking up the baffle step correctly you need to make sure your mic is at least twice the width of the baffle away from the speaker as well.** I bring this up because usually smaller drivers are on smaller baffles, so if we meet the criteria based on driver piston diameter we are probably fine, but if the baffle is wide you need to take this into consideration. Also, large speakers with multiple drive units that are spaced quite a bit apart need some distance to integrate. So, bigger speakers can become more problematic when measured in rooms and may require a greater measuring distance to be effective.

Here's a tip, if you are measuring in your room then closer is usually better, as long as you meet the requirements for being in the far-field. There is no rule that says you need to measure at one meter to have good measurements, and usually the opposite turns out to be true – measuring at one meter may instead allow for much more room interaction in your measurements. **For small two-way monitor-sized speakers I usually measure at a distance of around 20".** One meter measurements will work fine in an anechoic chamber. It is also used as a standard for establishing loudspeaker sensitivity. But for capturing measurements that can be used for accurate crossover and loudspeaker design, this distance is not necessary.

Near-Field Measurements –

Near-field measurements are usually done to overcome the effects of standing waves and reflections in the room. This is accomplished by placing the microphone very close to the cone so that room and baffle effects are very low in relative amplitude compared to the impulse response, or they do not have adequate time to develop.

For a near-field measurement to work correctly certain rules are usually followed: The mic needs to be placed as close to the center of the driver as possible. The mic then needs to be spaced approximately 0.10 times the effective radiating radius of the driver from the cone. So, for our 6.5" woofer, the 5" radiating diameter has a radius to 2.5", so the mic should be positioned 0.25" (1/4") off the dust cap. The reason we aim for the center of the cone is because at higher frequencies sound will arrive from different areas of the cone out of phase and create cancellation nulls and peaks that you won't see in the far-field. This is caused by the relatively large distance that the edge of the cone is from the microphone compared to the distance the center of the cone is from the microphone. By placing the mic in the center we force these cancellations to occur as specific frequencies, since the edge is the same distance away in all directions, and we extend our usable bandwidth.

Because of this cancellation effect the usable upper limit of the near field response is defined as $F_{max} = 4311 / \text{radiating diameter (in inches)}$. In our example this would be $4311 / 5$ or 862Hz. Below this frequency the near-field response will be accurate and free of room and cabinet diffraction effects, above this frequency the near-field data is less reliable and less accurate.

Near Field Cone and Port Summation –

In a vented speaker we may want to sum the output of the port and the cone into one single response. My Passive Crossover Designer (PCD) spreadsheet can do this very easily if you do a little work in advance. It works like this: Measure the port and cone responses using our near-field technique described above and save these responses as FRD files. Now, once you have taken your two responses unless the port and cone have the same area (which is unlikely, unless the port is a passive radiator) we will have to adjust their relative SPL's. The smaller of the two, usually the port, will need to be lowered in SPL (dB) because its smaller area will apply more pressure on the microphone making it appear louder. Its output needs to be lowered by the ratio of their areas using this formula: Lower the port by $20 \log (\text{Port Diameter} / \text{Cone Diameter})$. So, for our 6.5" woofer, with its 5" cone area, working with a 2" port, the port output will need to be lowered by $20 \log (2 / 5) = -7.96$ dB, or roughly 8 decibels.

Next take both near-field measurements and import them into PCD with the cone response as the woofer and the port response as the tweeter in a two-way configuration. Turn on the active section for the tweeter and adjust the port's level down by -8 dB. Now, if you look at the summed response (black line) in PCD it will be correct for the summed port and cone response for both amplitude and phase response. You might note that the Impulse Response on the port measurement is negative, so you don't have to change the phase in PCD it will already be correct. Finally, save the summed response as an FRD file and you can now use this new summed port and cone response to merge with the far-field response later if you choose.

Note: for dual ports, find the total port area of both ports and determine the effective diameter for the combined area. Now, take the measurement at only one port, find the ratio of the total port diameter to the cone and make the calculation and then add 6dB. So, for our example, if there had been 2 – 2” ports then the port output would be lowered by $-7.96 + 6 \text{ dB} = -1.96 \text{ dB}$. For dual woofers and a single port we will need to do it the other way around and subtract another 6dB from the port output.

Here’s a tip – what is the best indicator of the tuning frequency (F_b) from this data? The peak in the port’s output usually appears a little higher in frequency than the actual tuning frequency, so it is not the best indicator. However, the deep notch in the woofer’s response will be right on the actual tuning frequency of the enclosure and port. Do a near-field measurement at your cone and you will find the tuning frequency very quickly.

Gating the Impulse Response –

This sounds technical, but it’s not that complex. We have already seen that the Impulse Response is a time domain response, so “gating” or “windowing” is a time window. It is a measurement of how long of duration in terms of milliseconds that we continue to capture the impulse response and allow it to develop.

Long or short time windows, what are the pros and cons? When you view the impulse response the high frequencies are at the beginning of the impulse due to their short wavelengths, whereas low frequencies make up the tail of the impulse because they take longer to fully develop. Therefore, the longer your time window is open, the lower in frequency you will be able to measure. In fact, the lowest useful frequency is the inverse of the time in seconds, for example, using a gate of 5 mSec = $1 / 0.005 \text{ seconds} = 200 \text{ Hz}$. So, a 5 mSec Gate time is technically good down to 200Hz, if you do not have reflection occurring within that time window.

The “pro” here is that the longer your gate time, the lower in frequency your data covers. The “con” is the flip side of this coin. The longer your gate time, the lower in frequency you can go, but you also allow more time for reflections of the impulse to bounce off objects in the room and reflect back to the mic. We will see these as late ripples in an otherwise smooth impulse response. We can then set a limit on the Impulse Response window; this is called a “Gate”. A gated response is also called a “quasi-anechoic response” because we are attempting to capture only the first arrival sound pressure without any echoes (reflections) included in the data, hence the term “anechoic” which means “without echoes”. Once a gate is set, the Impulse Response is truncated at this length of time, and everything beyond this point in time is rejected. The FFT is only performed on the data inside the window.

One thing that should be noted is a compromise solution: take the impulse measurements closer to the speaker, like we are doing (not at distances of a Meter or more), and if needed raise the test volume (SPL) a little – I like it around 85db – and this helps to make the first arrival from the loudspeaker more dominant over the room modes to a lower frequency. The more dominant the first arrival sound is, the cleaner and lower you can measure, hence the reason the near-field measurement is effective at low frequencies.

One interesting feature in the Omnimic software (available through Parts Express) is the “blended” response. It uses the gated response down to the frequency corresponding to the gate time and below that “blends” the room curve with it. This can be quite useful, but it will still show a lot of response variations due to room modes and reflections making a mess of things.

Response Smoothing –

It should be noted that Gating the time window of the Impulse Response results in a form of response smoothing. You will easily see this in the frequency response, even at higher frequencies, if you shorten the gate marker to shorter and shorter time intervals. In addition to this there is octave to octave smoothing that can be applied to the response as well. This is a mathematical process where the frequency response is averaged over a broader and broader band of frequencies. It is usually expressed in terms like 1/3rd Octave smoothing or 1/24th Octave smoothing, etc. Our brains tend to perceive the tonal balance of a speaker based on 1/6th Octave smoothing. In other words, a response plot shown with 1/6th Octave smoothing will correlate most closely with our perceived tonal balance of the speaker. But, this is too smooth for good crossover design work because it can allow narrow band peaks and dips to be hidden from our view. I typically use 1/48th Octave smoothing when I take measurements to design a crossover because I want to see everything that is going on and make sure nothing is sneaking past me. I would recommend 1/24th Octave as a bare minimum when designing a crossover. Crossover programs like my Passive Crossover Designer (PCD) will usually interpolate the input data to some an array anyway. In the case of PCD the data is interpolated to approximately 1/48th Octave for display and calculation, so if less smoothing than this is applied it will not be seen in PCD.

Merging Near-Field and Far-Field Responses –

Above we discussed the proper methods for collecting the near-field frequency Response and the far-field frequency response. We have also discussed Gating the Impulse Response window. From this we have seen that the near-field response is limited to a maximum usable frequency for accuracy and the gated far-field response is limited in the minimum frequency for accuracy. Wouldn't it be great if we could get the best of each together? We can!

Using our 6.5” woofer example we saw that the near-field response was limited to an upper limit of 862 Hz. And using our gated impulse response with a 5 mSec gate time our far-field response was limited to a low frequency limit of about 200 Hz. This means that these two response curves should both be good in the two octave range of 200 – 800 Hz, and this gives us the range in which we can look to merge this data into one “spliced” or “blended” frequency response. You will typically find that the near-field data is much higher in SPL than the far-field, and it is usually best to lower its level to match the far-field data.


Hard Splicing versus “Blending” –

Traditionally, near-field and far-field data are merged with a hard splice at a specific frequency. For example, say 300 Hz is chosen; the two response lines will be adjusted to meet at 300 Hz and then spliced with a hard cut at this frequency. This can sometimes be difficult to get a good match that does leave some degree of discontinuity in the response plot. Most software works this way, including my

Response Modeler spreadsheet. However, recently Charlie Laub and I have worked together to create a spreadsheet. One of its features is that rather than performing a hard splice it can provide a softer splice and “blends” the response over a user-defined frequency band. This eliminates the risk of a discontinuity and provides a much smoother transition from far-field to near field response data. This program also allows for a modeled bass response to be blended if desired, and sometimes this is more than adequate and very similar to merging the near-field measurement data of sealed systems and maybe a little better (certainly easier) than working with the complexities of vented box data where the port and cone have to be summed as well.


Working With an Example –

Let’s work with an actual example using what we have covered already, and using the traditional method of splicing the near-field and far-field data together. Our example uses a 6.5” midwoofer mounted on a baffle that is 8.5” wide and 15” tall. I took a far-field measurement at 21” on the woofer’s axis. The SPL shown in the graph is actual and not normalized to any particular input power or distance. Here is the far-field response:





In this graph the gated impulse response is good down to 300 Hz and the response below 300Hz is blended with the room response data. The Omnimic was used to measure this response.

Next, I moved the microphone to the center of the dustcap and captured the near-field response according to the method described above. Here is the near-field response:



Based on what we learned above the data in this graph is only useable up to 862 Hz. On the next page I will show these two graphs again and highlight only the useful and accurate regions in white.




The merging or slicing of the near-field and far-field responses need to take place in the yellow frequency band that is shared by both sets of data, at a point where we can line up both sets of data together.


The Problem with the Traditional Splicing Method –

Using the “FRD Response Blender” spreadsheet by Charlie Laub and myself I imported the far-field and near-field data into the program (please see the User Guide for more information on using the FRD Response Blender. It can be downloaded here:


http://audio.claub.net/software/FRD_Blender/FRD_Response_Blender_v2.0_Tutorial.pdf . I found that lowering the near-field data by -14.7dB resulted in excellent alignment of these two data-sets in our highlighted region, as shown in the graph below:



The red line is the near-field measurement and the thin blue line is the far-field measurement. Zooming-in we see that we can line-up both responses in the region between 300 – 450 Hz. This is very typical, and the normal range you will find for splicing data.



Using 400 Hz as our splicing point, which is the traditional method, yields the following merged response plot (I have not extracted minimum phase yet, so the phase data does not show):



This graph is representative of most spliced graphs and matches what we will see in most tutorials and manuals on splicing near-field and far-field data. However, there is a problem with this data and it is right in the region that is most critical to us – in the 150-200 Hz region.

This is considered our reference level for crossover design. For almost any speaker the region between 150 and 200 Hz represents the bottom of the baffle step region and the area where the speaker is now radiating into full space. Above this frequency the speaker is transitioning to a half-space radiation due to the baffle, so the response is rising, and below 150Hz the room has taken control of the frequency response (below the room's Schroeder Frequency).


Once you have taken your far-field and near field frequency response and accurately merged them, then we need to make note of the level between 150 – 200 Hz. In your crossover design work you should attempt to bring everything above these frequencies down to this level to achieve flat response. This will result in the most accurate tonal balance in your room. BUT, if this region is incorrect in our spliced data we will not accurately set the level for the rest of the crossover and we may not arrive at the tonal balance we are targeting.

So what is causing this problem? It goes back to our nearfield measurement. As we noted in the section above on near-field measurements, not only does this method eliminate room reflections it also

eliminates baffle diffraction, so the baffle step at lower frequencies is not shown in our data. Our spliced data, therefore, does not accurately show the correct free-field baffle step response.


A New Method That Solves This Problem –

Built into the FRD Response Blender spreadsheet is a baffle diffraction modeler. This allows you to enter the baffle dimensions, the placement of the driver on the baffle, and adjust other parameters like, piston directivity, edge radius, and off axis angle. The modeler will then produce the calculated baffle diffraction response and add this to the imported near-field response. Here is a screen-shot of the modeler:



When this curve is added to the near-field response it can be overlaid with the far-field data more effectively because now both near-field and far-field response data include the baffle step information in the critical range where they overlap.


The next graph shows the overlay of the near-field + Diffraction with the far-field data. You can see now how well you can align them over the range of 300 – 862 Hz where the two set of data should both be valid.




The near-field plus diffraction even aligns well with the blended far-field in the data below 30Hz as well.

Now, in order to facilitate a smoother transition, rather than performing a hard splice at 400 Hz as I did before I will instead allow the FRD Response Blender to blend the two sets of response data together between 200 and 500 Hz.

Here is a graph that zooms-in on this region and shows the blending process:



The final blended response now looks like this:



I contend that this final response is an accurate free-field quasi-anechoic response down to 10 Hz. It is obtained by taking proper far-field and near-field response measurements and merging them together through a method that significantly reduces error. This new curve will give you a much better method of establishing the level in the critical 200 Hz range necessary for crossover design than the traditional method provides. There is now only one thing missing from our data - the phase response.

Phase Response –


All frequencies are cyclical; they involve an outward (compression) and inward (rarefaction) movement of the cone at X number of times a second. These show up in the impulse in the time domain. If the cone is moving forward the spike will be in the upward or positive direction and when it is oscillating backward the impulse will go downward in the negative direction. When the FFT convolves the impulse response it will not only show the frequency response but also the phase response, since this is captured in the time domain data. Loudspeaker drivers are what are called “minimum phase devices”. This means that the phase response can be directly extracted from the frequency response because there is direct causality. The mathematical means to perform this extraction is called a Hilbert-Bode Transform. Minimum phase means there is no excess delay affecting the phase, only what is attributable to the frequency response, since it is the phase at the cone itself.

Since drivers and crossover components are minimum phase, how they sum is predictable and clearly mathematically defined. If this was not the case then crossover programs would never be able to predict

the final summed response. Measurement systems can yield different versions of phase data depending on how it is captured, whether there is a distinct time marker used to begin the impulse, and whether there is time of flight included in the phase data. All of these factor in to determine whether it is best to use measured phase to extracted phase in crossover development. Programs like Passive Crossover Designer can work with either, but if extracted minimum phase data is used with accurate offsets to define the relative locations of the drivers, then the program can calculate the frequency response on various axes and show their effect on the crossover summation. This can add a powerful dimension to the crossover design process.

Since we have manipulated the measured response of our sample midwoofer by blending the near-field and far-field data together we have lost the original phase data that was measured with the response. In order to restore that phase we now need to extract it from the final frequency response. Before doing this there is something we need to do first, and that involves adding “tails” to the frequency response. These tails extend the frequency response in smooth slopes from near 0 Hz to beyond 100 kHz. Adding a tail to the low-end is rather straight-forward. Most of the time the low-end roll-off is either 12 dB/Octave or, in the case of a vented system, 24 dB/Octave. It is more complicated at the upper extreme and requires some approximations, but typically a slope of 12 – 24 dB/Octave on the upper end will work well too.

Charlie has created a very nice Hilbert-Bode transform built into the FRD Response Blender that can extract the minimum phase without requiring any macros or VBA programming. This way this feature is now accessible to those using Excel in environments like Open Office or Office for Mac. Below is the extracted minimum phase response and frequency response for our test midwoofer:



Conclusion –

The purpose of this paper was to describe a new method of arriving at an accurate quasi-anechoic free-field frequency response down to 10 Hz that anyone with a simple measurement system could make in their own living room.

This method involved following the standard procedure for taking far-field and near-field axial frequency response measurements, importing them into a new spreadsheet design tool (the FRD Response Blender), simulating the baffle, adding that response to the near-field response, adjusting levels to align, and blending the responses together over a user defined range. Finally, once this is completed all that is needed is to extend response tails, extract the minimum phase, and export the final response for use in a crossover design program.

It has been my experience that this method works very well, as was demonstrated at the Loudspeaker Design Workshop in Kokomo, Indiana, on December 14th, 2013.

Following this process should give every builder the opportunity to work with accurate and qualified data in their DIY loudspeaker project.

The FRD Response Blender Excel Spreadsheet can be downloaded here:

http://audio.claub.net/software/FRD_Blender/FRD_Response_Blender_2.0.xls

The Passive Crossover Designer Spreadsheet can be downloaded here:

<http://audio.claub.net/software/jbagby.html>

Jeff Bagby
1/3/2014