5^{th} Elements guide to making useful loudspeaker measurements

I'll start off by assuming that you've already got a decent microphone, mic preamp and a sound card with decent quality line inputs. I will also be using ARTA for this guide, there are a number of free/share-wear programs out there and they all work in a similar way.

If you have read anything about making loudspeaker measurements before then the two phrases you've probably heard used more then most are far-field and near-field. These are two measurement techniques and are vital to the end goal of making accurate and meaningful measurements of loudspeakers in less then ideal conditions. Ideal conditions would be something like an anechoic chamber or a very large open space, such as an empty car park on a windless day with your loudspeaker several meters up in the air.

First of all lets discuss the far-field measurement technique. As it's name implies this is done with the microphone positioned a reasonable distance away from the loudspeaker enclosure. How far away you need the microphone depends on the size of the loudspeaker under test and ideally you'd position the microphone at the listening distance, but this isn't always suitable for a number of reasons. It is probably worth mentioning here that the only real downside to far-field measurements is their susceptibility to contamination due to sound reflections from objects within the measurement environment.

As an example here is an impulse response taken in a far-field environment. You can see the initial impulse is inverted, but importantly note that the impulse decays quickly to nothing. After a short time however we see a wiggle in the impulse response and this is a reflection from a nearby surface.

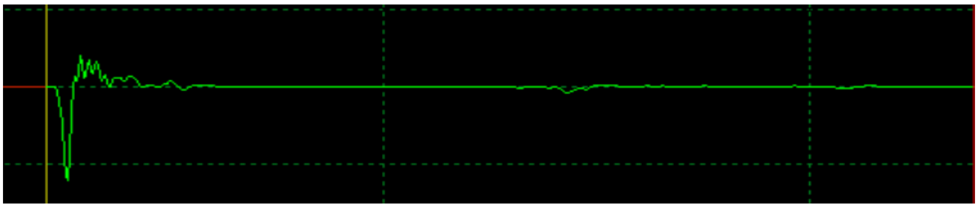

If I gate the measurement to block out the reflection then measurement looks like this.

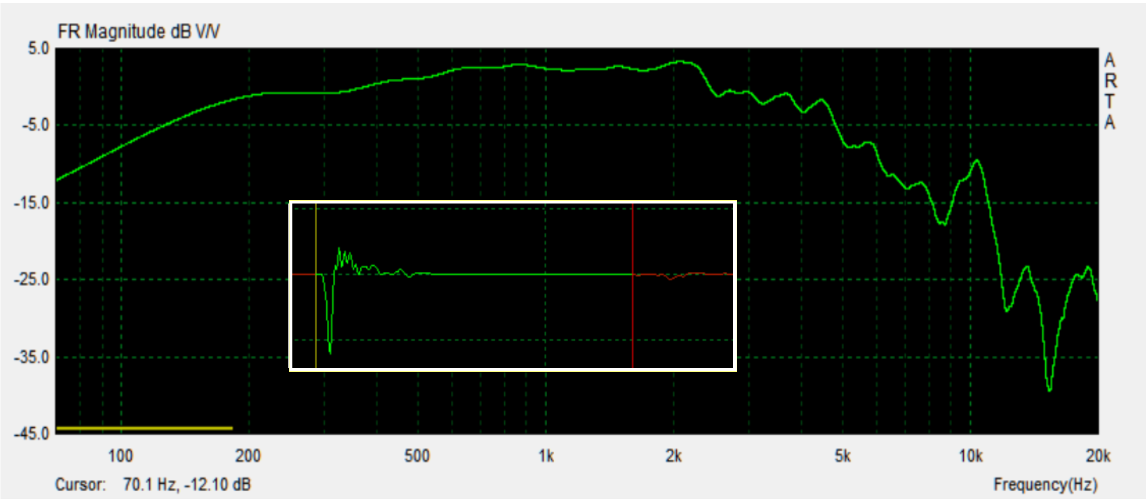

If however I now increase the length of the gate to include the reflection then the response ends up looking like this.

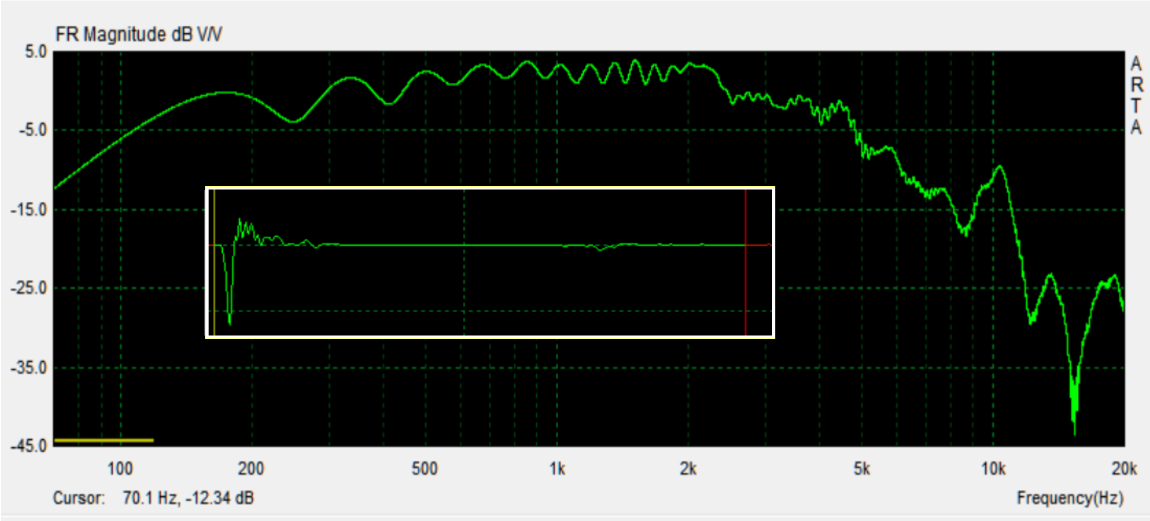

As you can see, including the reflection introduced sound that wasn't apart of the original impulse and as a result it interferes with the original. One thing to notice is the little orange line in the bottom left hand corner of the measurements. You will notice that in the last measurement that the bar reaches up to around 130Hz, whereas in the first it's a lot closer to 200Hz, this is because I changed the length of the gate. The larger the gate the better the lower frequency accuracy of a far-field measurement becomes, but as you can see, if you let in reflections all that accuracy is lost. In an ideal measuring environment (an anechoic chamber) we can increase the length of the gate significantly so as to get accurate data down to say 20Hz, but due to the real world limitations of a home measuring environment this is virtually impossible.

This however is where near-field measurements come in and as the name implies are done with the microphone placed up close, right next to the cone/dustcap of the loudspeaker. So how is this useful? When you place the microphone next to the cone of the loudspeaker the sound the microphone receives is relatively very loud. The sound still radiates out into the room the and is still reflected off of nearby surfaces and ends up back at the microphone. But in relation to the level of the direct sound, the reflections are pathetically small, this in effect renders them a non issue.

Here is a measured frequency response taken from a near-field measurement and its associated impulse response.

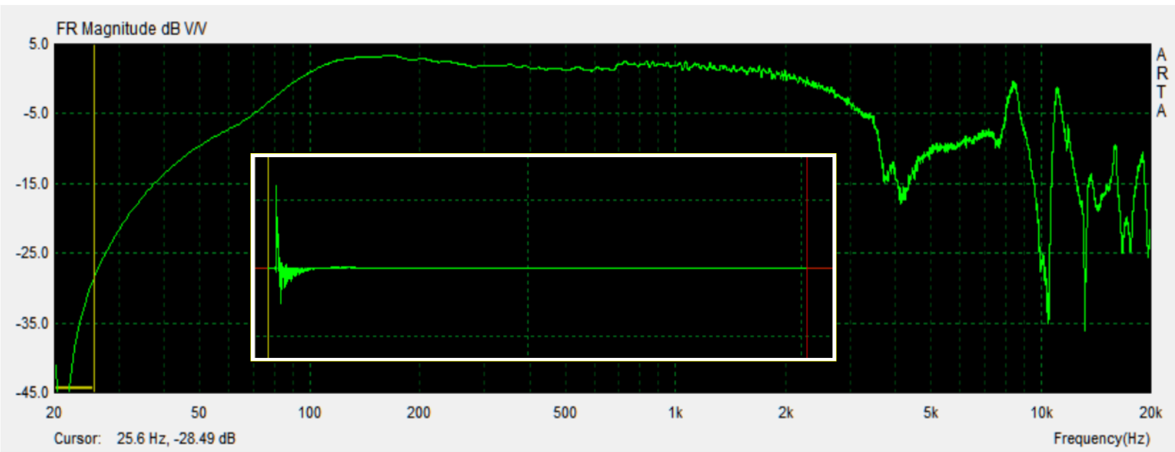

As you can see the tail end of the impulse response is clean and allows us to increase the length of the gate, in this I set it to ~40ms. This allows us to take measurements with very good accuracy down low, in this case to 25Hz, but it is useless for giving accurate data in the high frequencies. As the microphone is placed right next to the cone it will not measure any effects due to cabinet diffraction and it will not measure any effects due to baffle-step, it will also not give an accurate representation of the drivers own high frequency response. As a result we can combine the two measurement techniques to get the best of both worlds.

At first glance this seems like a relatively easy thing to do, but making measurements that are actually useful for proper crossover design is a little bit more complicated and if neglected your measurements might look nice, but will be virtually useless for accurate simulations.

The most important thing about loudspeaker measurements are the far-field responses. In a way you can look at it from one simple point of view – if we all had anechoic chambers all we'd ever need would be far-field measurements. As we don't have anechoic chambers we end up having to augment the far-field with a near-field splice at a convenient frequency. Note that this is not the same thing as augmenting a near-field with a far-field splice, this is going about it the wrong way round.

Why is this the wrong way round? You might ask and it's a good point, but here's why. When performing loudspeaker measurements we are actually after two things. The first is the frequency response data and the second, but of no less importance is the phase response. We have all seen what phase responses look like and if I take the first measurement and add in the acoustic phase it ends up looking something like this.

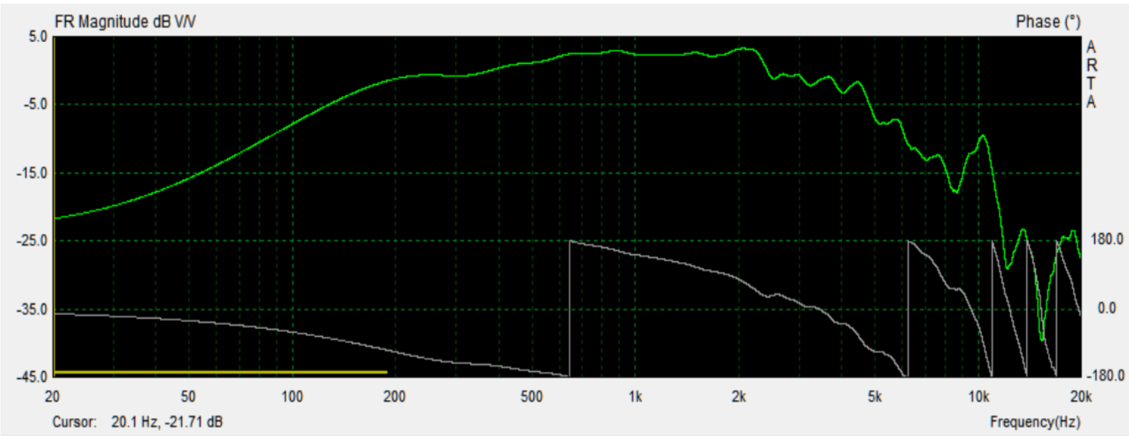

Here we can see that the grey trace has been added and you may ask, well what use is that? The simple answer is this, on its own it's relatively meaningless. One could move the microphone elsewhere, perhaps off axis or alter the distance to the microphone and we'd end up with a completely different phase response. What's important about the phase response is that it allows us to see how things compare or relate to one another. In other words the phase allows us to see what relationship there is between the tweeter and the mid/bass of our system so that we can properly integrate them when designing the crossover. If the measurement of the acoustic phase is done incorrectly then this relationship is lost and we're stuck. So how do we do it correctly?

The first thing to note is that the measurements must be done from one fixed position. That is we measure both the tweeter and the mid/bass with the microphone and the loudspeaker in exactly the same place. If one moves either of the two, between performing the individual measurements of the different drivers, then the exact relationship between the two is lost.

As a way of illustrating this, we all know that when a tweeter and mid/bass are mounted in a cabinet, that the tweeter is generally mounted forwards of the mid/bass. This tends to move the acoustic centre of the tweeter forwards and under the correct measurement conditions would cause the impulse response to arrive faster for the tweeter then it would for the woofer. This is something we are very interested in and is captured when we perform far-field measurements in the right way. If however, I moved the microphone between the measurements of the tweeter and the woofer, this change in position would also alter the impulse response so that it would arrive at a different time and we'd be now lost as we don't know how much of the change in the time of arrival is caused by the microphone moving position, or simply the offset between the tweeter and the mid/bass. So for this reason we leave the microphone in one place.

Okay so we've got that far, we now know that we are going to do a far-field measurement at a distance of 1 meter without moving the position of the microphone or the loudspeaker, but now what?

First of all place the microphone on the listening axis, this isn't critical but as you need to position the microphone somewhere in-line with the front of the cabinet, it may as well be on the listening axis. For the sake of this example I have chosen to put the mic at the mid/bass level. It is also important to make sure that the microphone and loudspeaker are positioned clear of the floor. If you can get the mid/bass up to 1 meter away from the floor it should be acceptable.

The next step is to measure the first driver, in this case I will measure the tweeter and we end up with an impulse response that looks something like this.

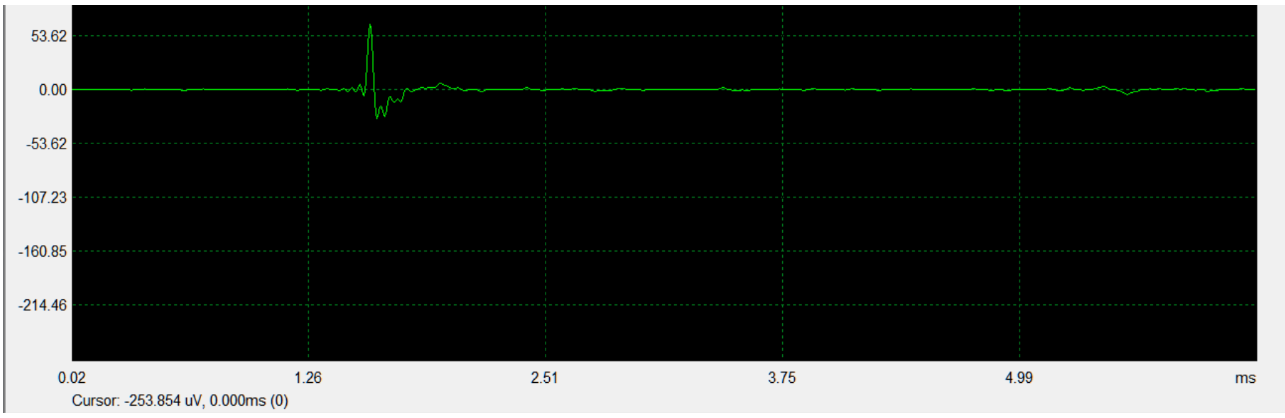

This is nice and useful and also reasonably clean. Now we've got to set the gate so that we can get some useful data out of it. One thing that is critical here is to remember where you set the start point of the gate.

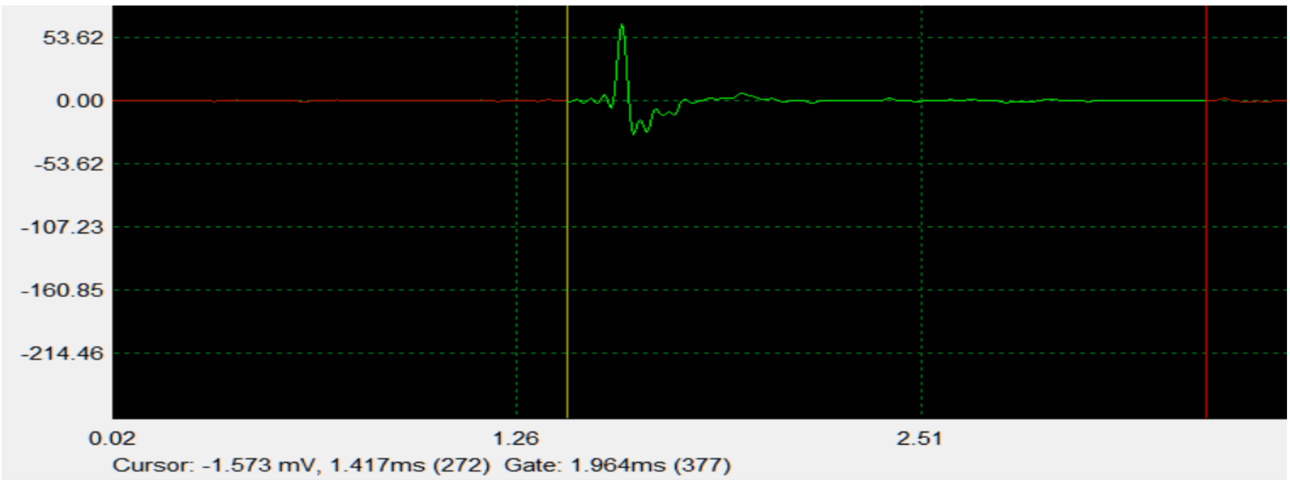

Here we can see that I've chosen a start point of 1.417ms with a total gate length of 1.946ms. This produces the following completely un-smoothed graph.

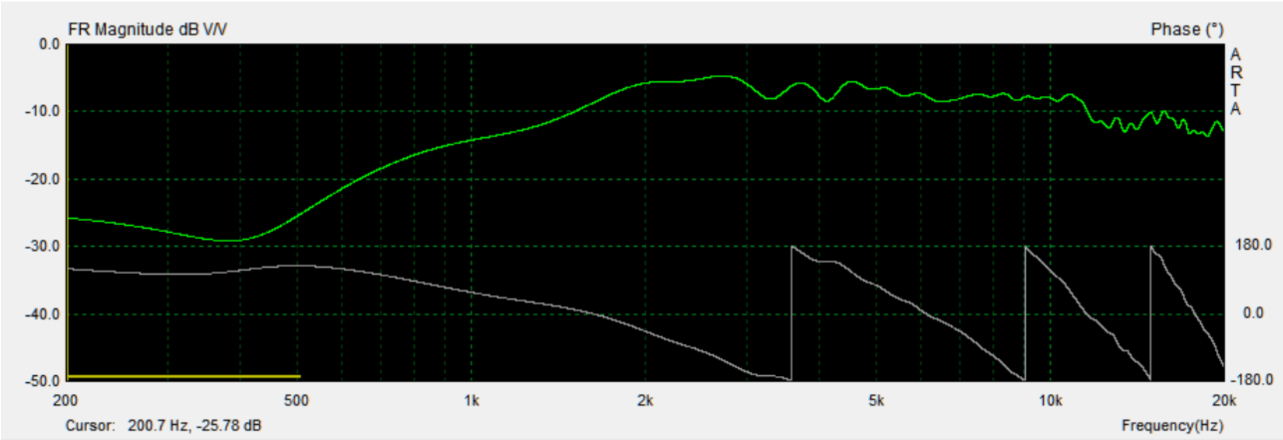

As we can see the relatively short gate only allows us accuracy down to around 500Hz, but as this is miles away from the proposed crossover of around 2khz it is more then acceptable. The next stage is to export the graph and save it. Don't forget to turn on the phase if it isn't on as standard. Next up is measuring the mid/bass and this is what we get.

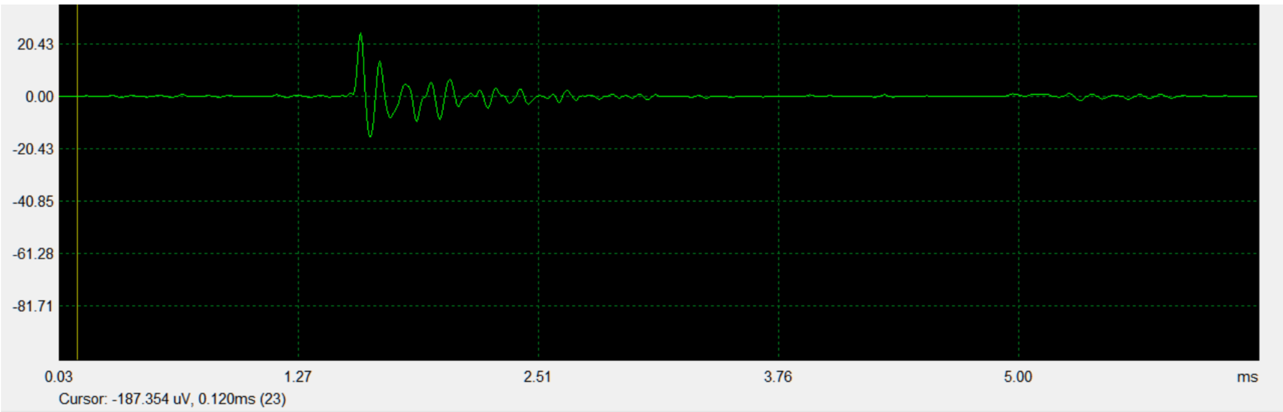

Once again we now have to set the gate, but we must remember to set the start of the gate to the exact same position we had it at when we set the gate for the tweeter, that is 1.417ms.

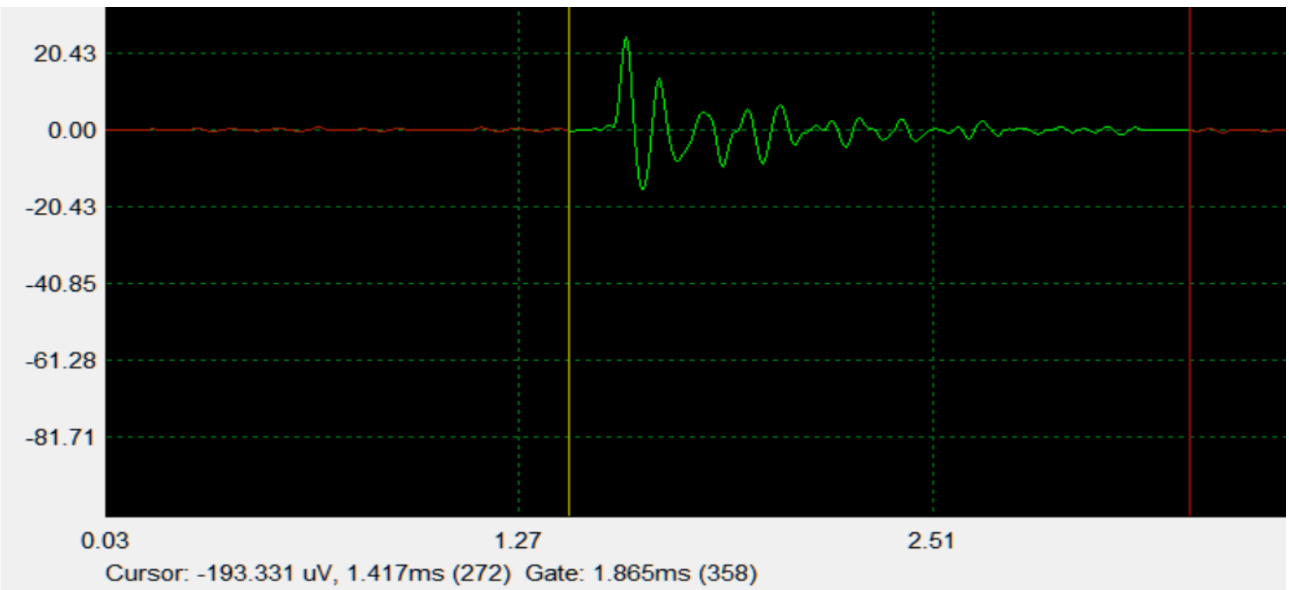

This then turns into the corresponding frequency response graph. It is worth pointing out here that the start point of the gate has to be in the same place, but the end point is free to move around. This is handy because the drivers are mounted in different locations in the room relative to the reflective surfaces and this obviously alters the timing of the reflections incident on the mic. Here I had to shorten the gate by a fraction to keep the first reflection out.

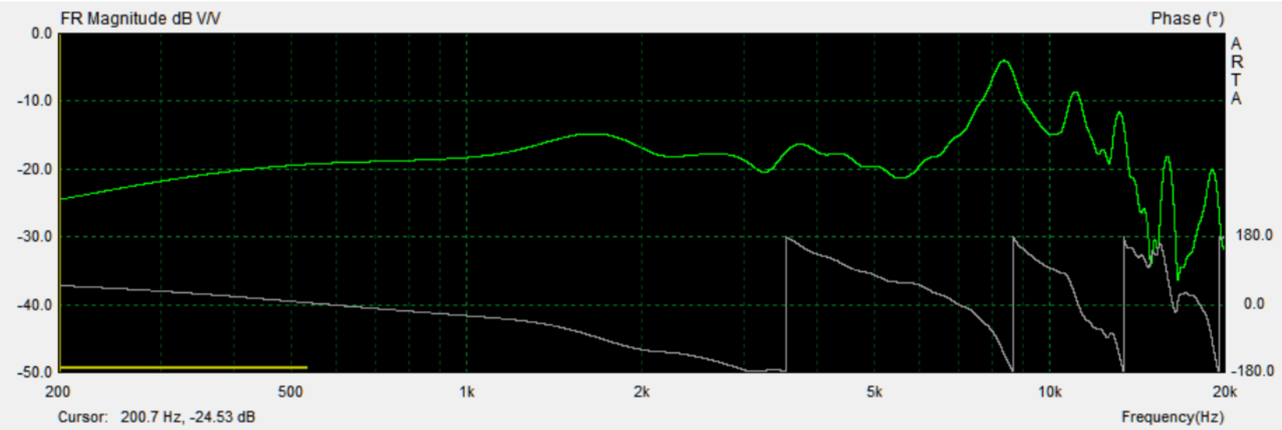

One thing to notice here is the hump at around 1700Hz, this is the standard diffraction bump caused by the narrow cabinet. What you can see however after the bump is the response starting to trail away, this is the start of baffle-step and will be important later on. Export the data.

The next important measurement to make is one with *both* drivers connected up in parallel. By now I don't think I need to show you the impulse responses etc, but remember once again to keep the start point of the gate identical.

The combined response ends up looking like this.

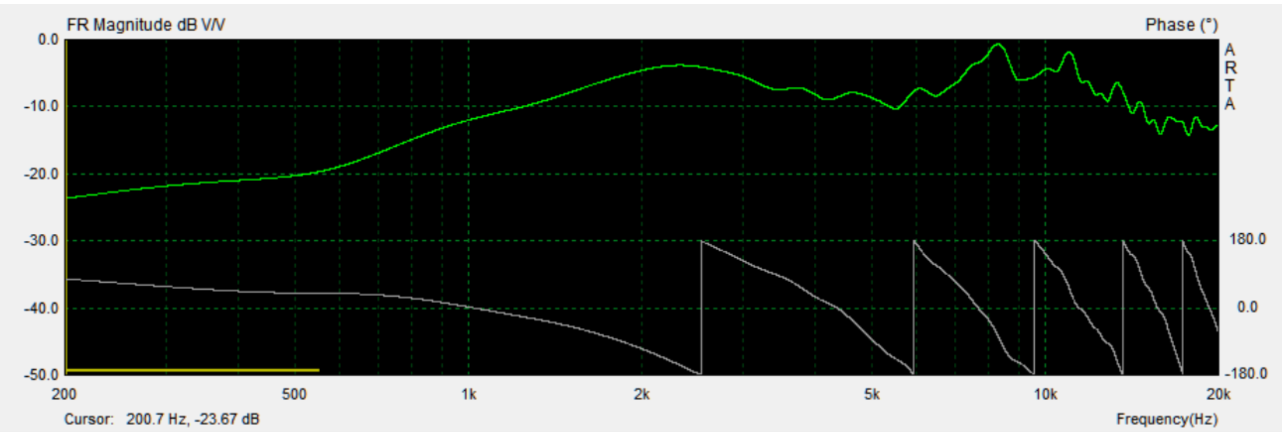

Now to a certain degree this is all we need to do decent crossover work. The far-field measurements have given us accurate data down to a low enough frequency to allow us to simulate well enough around the proposed crossover frequency of 2kHz. The far-field measurements have also given us accurate measurements of cabinet diffraction too so we can attempt to account for it in the crossover. The only thing we haven't yet got an accurate representation of is the lower end of the mid/basses response and the complete effect of baffle-step.

If one were designing an active two way system they could stop here. Baffle-step compensation is a relatively easy thing to compensate for and is accomplished by using a simple shelving network and this can be accurately implemented using mathematics alone – ie we don't actually need to measure it. Indeed when combining a near-field response with a far-field one, we don't actually measure baffle-step we simply tack on a best guess simulation. As mentioned before the far-field measurement contains all the useful phase data necessary for accurate crossover simulation and we do not wish to alter this, so if we want to add in the low frequency data we simply replace the low frequency portion of the far-field measurement with the low frequency portion of a near-field measurement/simulation. This leaves the high frequency phase data in tact so we can still simulate accurately.

Lets assume that for whatever reason you do not need to measure the low frequency response of the driver, or you don't need to have the effects of baffle-step included in your measurements. Now what? This is a very good question. You've taken various measurements but what do you do with them?

The first thing you need to do is import them into your loudspeaker design program. For this I use LspCAD and once imported the program will end up displaying something like this.

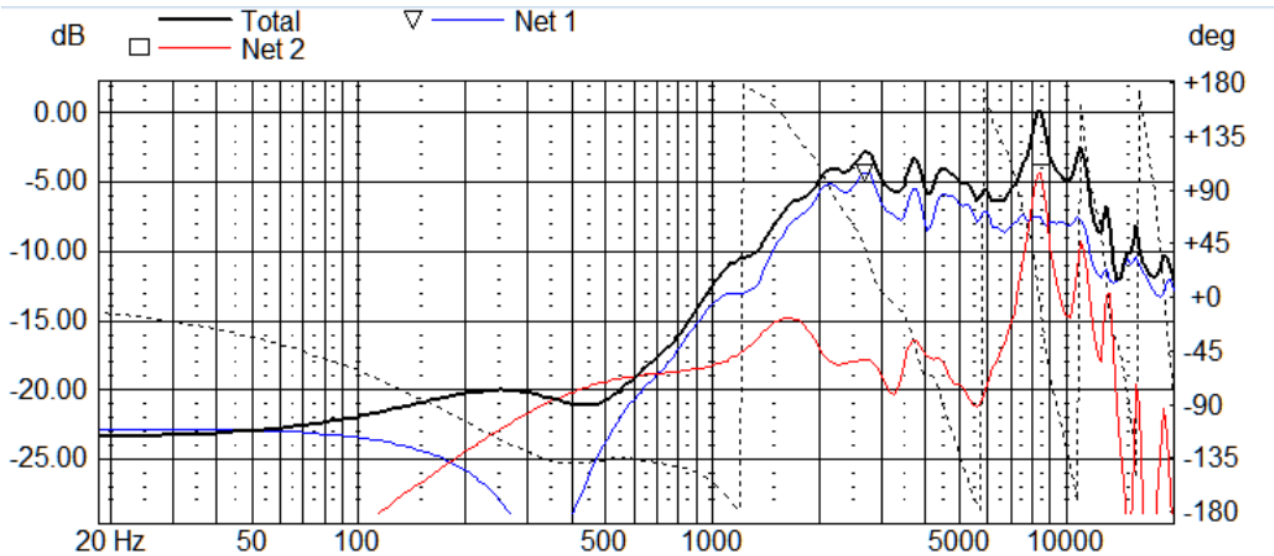

As you can see we've now got the tweeters response (blue) in combination with the mid/basses response (red). The black line however shows what the simulation program has calculated to represent how the two drivers sum together. Now as we actually measured how the two drivers summed before hand, we can now include that measurement and compare the two against one another. When done in LspCAD we end up with something like this.

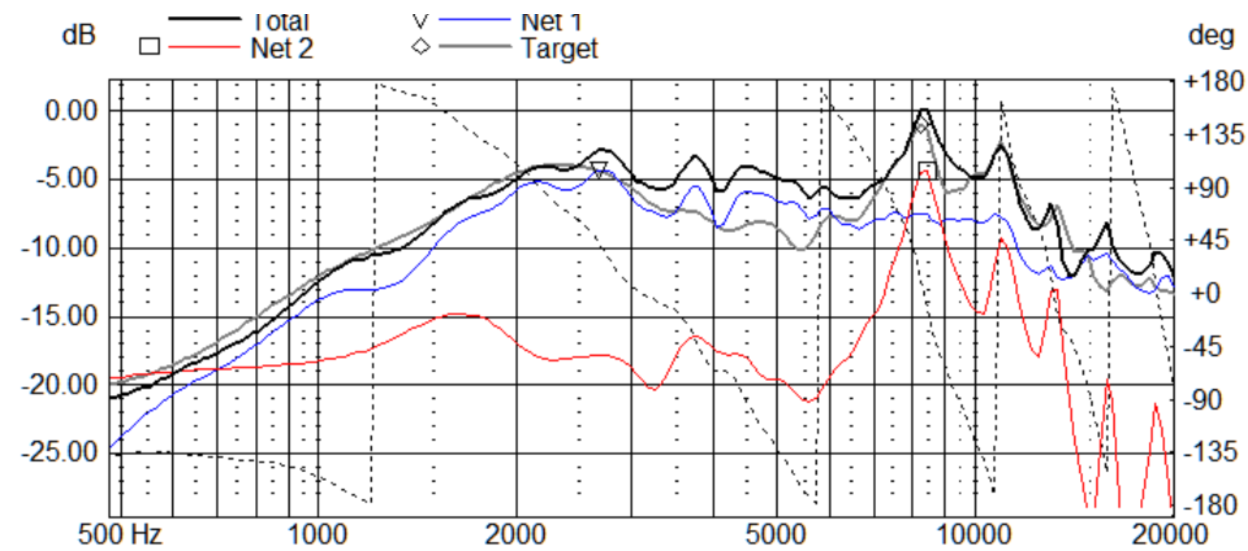

I have deliberately cut off the lower frequencies to make this easier to see. But what we're doing is comparing the grey line with the black line. The grey line represents what is, the black line represents what the simulator thinks is correct. Quite clearly the two aren't perfectly matched so what do we do?

The first thing is to make sure that the drivers positional data has been entered correctly. Earlier on, if you recall, I said the microphone was positioned on the mid/bass axis, but so far the measurement program does not know this. In LspCAD you have the option to enter data relevant to the drivers positions. This is defined as dX, dY and dZ. In this case both drivers were mounted directly above one another so dX for both = zero. The tweeter however was mounted 180mm above the mid/bass. As the microphone was on the mid/bass axis this means we need to enter +180mm for the dY axis, for the tweeter. For example.

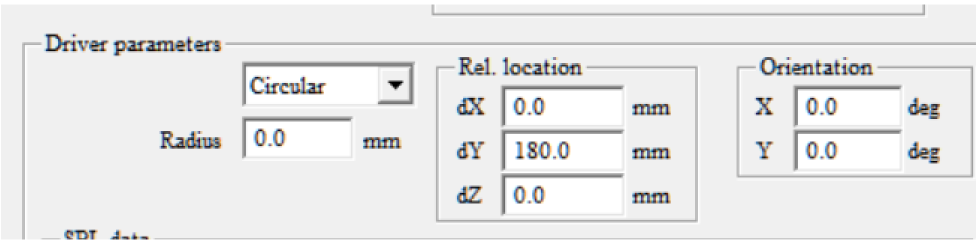

The next important thing to set is the measurement distance, in LspCAD this comes under the general tab and is usually set to infinite. In my case the measurement distance was 1 meter so that is what I have set.

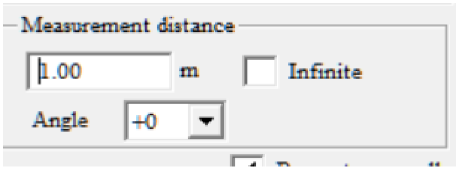

Now you will notice above that you can also enter the radius for the drive units diaphragm. In this case it is important to leave this at zero. As we had the microphone already placed off axis from the tweeter we actually measured the off axis response. If we entered a radius here, LspCAD would notice, from it's position above the mid/bass, that the tweeter is off axis and would now calculate how it's response should changed based off of a perfect model of off axis behaviour. As we actually measured the off axis response we do not need LspCAD to calculate this and add in something extra that is not require, so we leave the diaphragm radius at zero.

Having done all of this LspCAD now calculates that the combined response for the two drivers should look something like this.

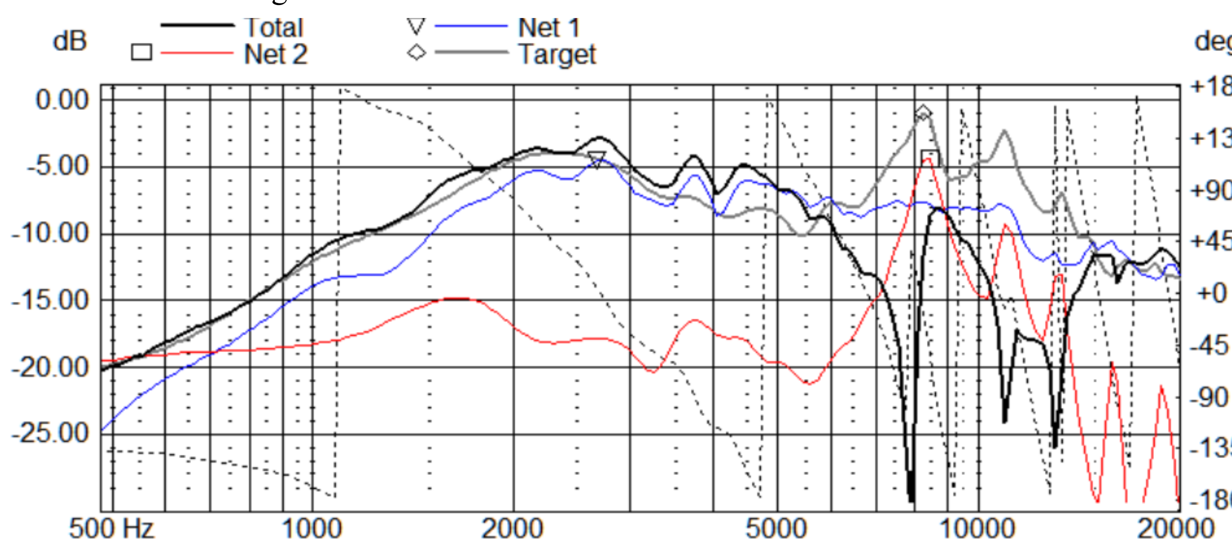

This is worse then the original you might be thinking, but there is one critical parameter that we have not yet set and that is the dZ parameter. The dZ parameter if you like represents where the acoustic centres are for the drivers and is difficult to measure. The measurement process outlined is done specifically so we can determine what this is by exclusion of every other measurable parameter. In other words, everything else that could have some affect on determining the position of the acoustic centres has been taken care of and all we're left with is trying to find out what the value of the offset between the two drivers is. In this case all we need to do is alter the dZ parameter until the measured combination matches the simulated combination. I always like to represent the dZ parameter as a negative value, ie the driver needs to move backwards to bring the Z plane of the drivers into alignment. In this case, as the tweeter is in wave-guide, its acoustic centre is pushed behind the mid/bass so I have to alter the dZ parameter of the mid/bass to a negative value to bring the two into alignment.

Bringing the dZ parameter of the mid/bass backwards by 15.8mm...

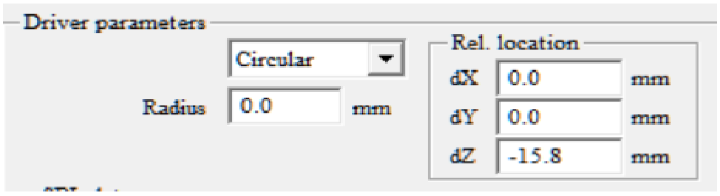

...arrives at this frequency response plot.

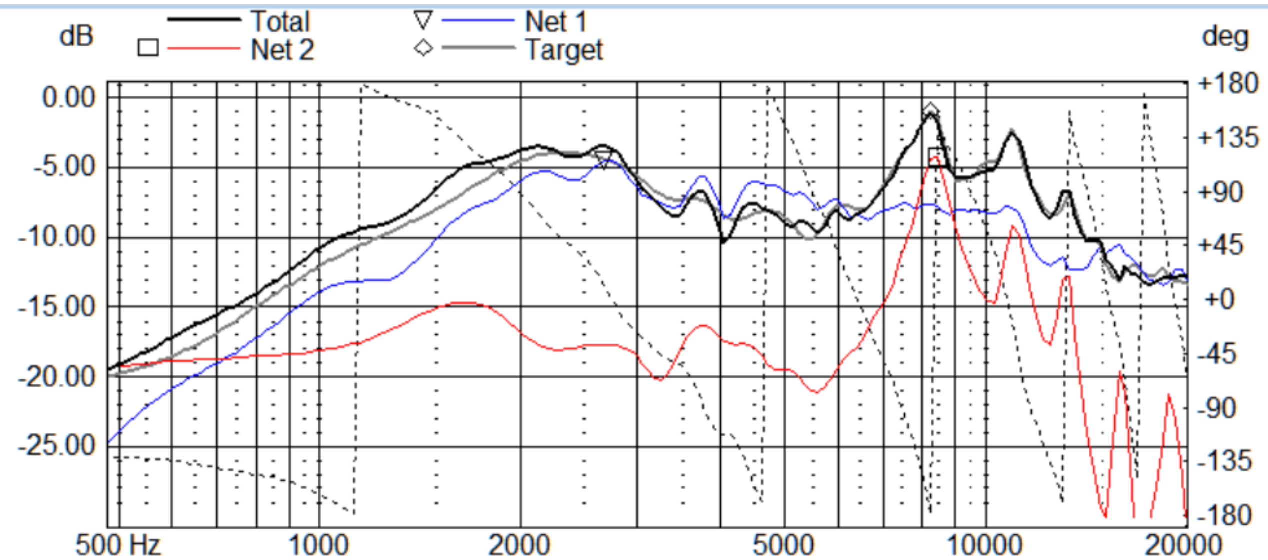

As you can see the two are now in pretty good agreement with one another. What this means is that LspCADs virtual world of how the two drive units are interacting is in good agreement with how the two drivers actually were interacting with one another in the real world. Now and only now are we in a position where we can hide the grey target response and are free to start simulating. It is worth mentioning here that once all the parameters have been set that you are free to alter the measurement distance now to = the listening distance if you so desire.

What do we do though if we want to combine a near-field response with a far-field response? This is a fairly simple process, but depending on the software you're using can be quite frustrating to actually do. JustMLS, as is provided with LspCAD, makes this swift and painless, doing it with ARTA on the other hand is a bit more of a pain. The first thing to do however is to pull up your far-field response and set it as an overlay. After having done so pull up your near-field response, you should end up with something like this.

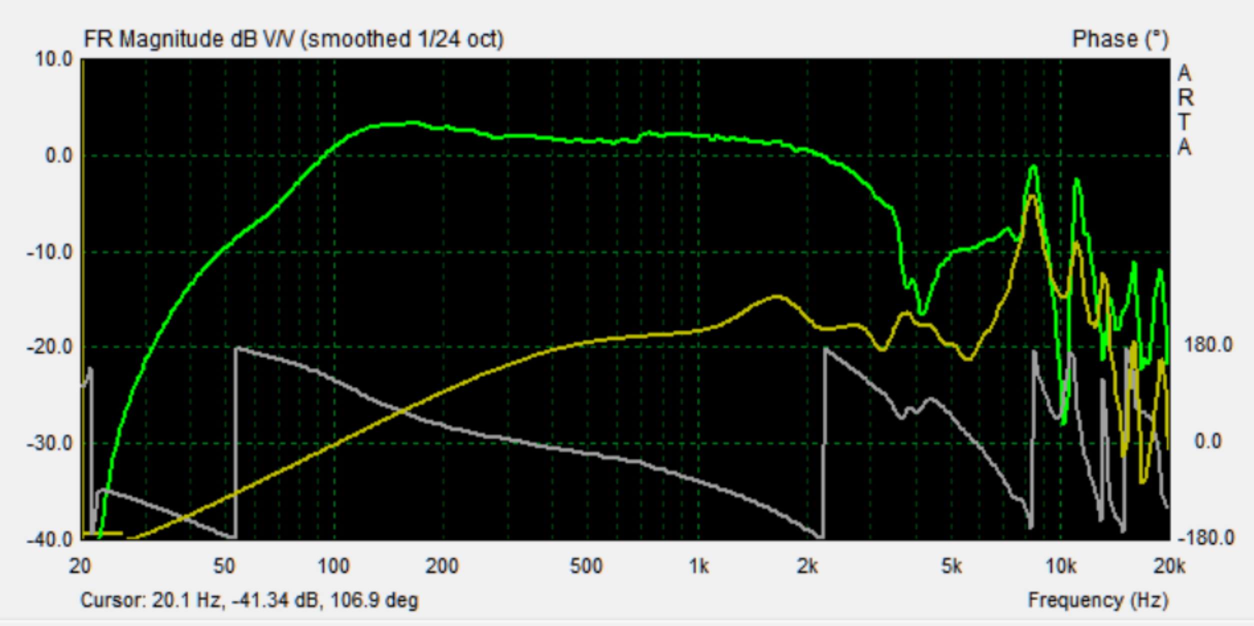

When setting the gate on the near-field response it is once again important to set the start of the gate to 1.417ms.

Now we've got both the far and near-field responses on the screen and we need to think about how's best to combine them. This is rather simple but needs a bit of thought. First of all pay attention to the region around 1200Hz. This represents the area just before baffle-step is about to take place, so has yet to affect either of the responses. If you compare the two you will notice that they are out by about 20dB. In this case we click on 'edit' in ARTA and alter the near-field response back by -20dB. The next thing we have to do is add in the effects of baffle-step to the near-field response. ARTA can do this for us by clicking on edit and selecting LF box diffraction, here you are prompted to enter a number that corresponds to the width of your enclosure. I am using a curved sided enclosure so I will enter the number that corresponds to the enclosures maximum width, in this case a little over 19cm. After having done so we end up with something that looks like this. It is important to alter the level in dB first, before adding in the diffraction. If you try altering the level after the diffraction it will for some reason reset the effect of the LF diffraction.

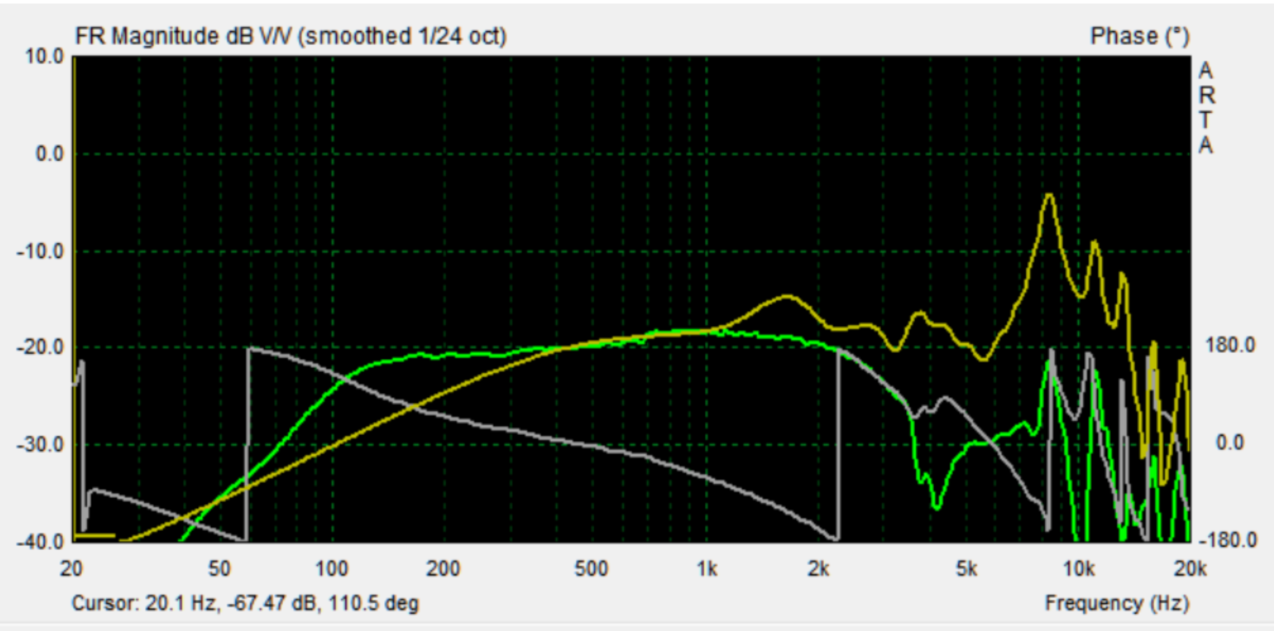

What you can see here is what you should hope to see when you do this for yourself. As mentioned before, as the far-field response extends down to 500Hz it contains the start of baffle-steps effect. We then added in a simulated version of baffle-step to the near-field response and over-layed the two, you can see here that both are in good agreement with one another over a large area of overlap. This means that everything is working well and working as it should do – ie the measured start of baffle-step is blending in nicely with the simulated baffle-step as applied to the near-field response. Now we have to merge the two together but first we have to decide where we want the merge to take place. In this case, because the area of overlap is so nice and broad, I think the best place to merge the two is at around 500Hz simply because we may as well take use of all the accurate part of the far-field measurement, so we put the cursor at 500Hz. After having done so you click 'merge overlay above cursor' and you end up with something that looks like this.

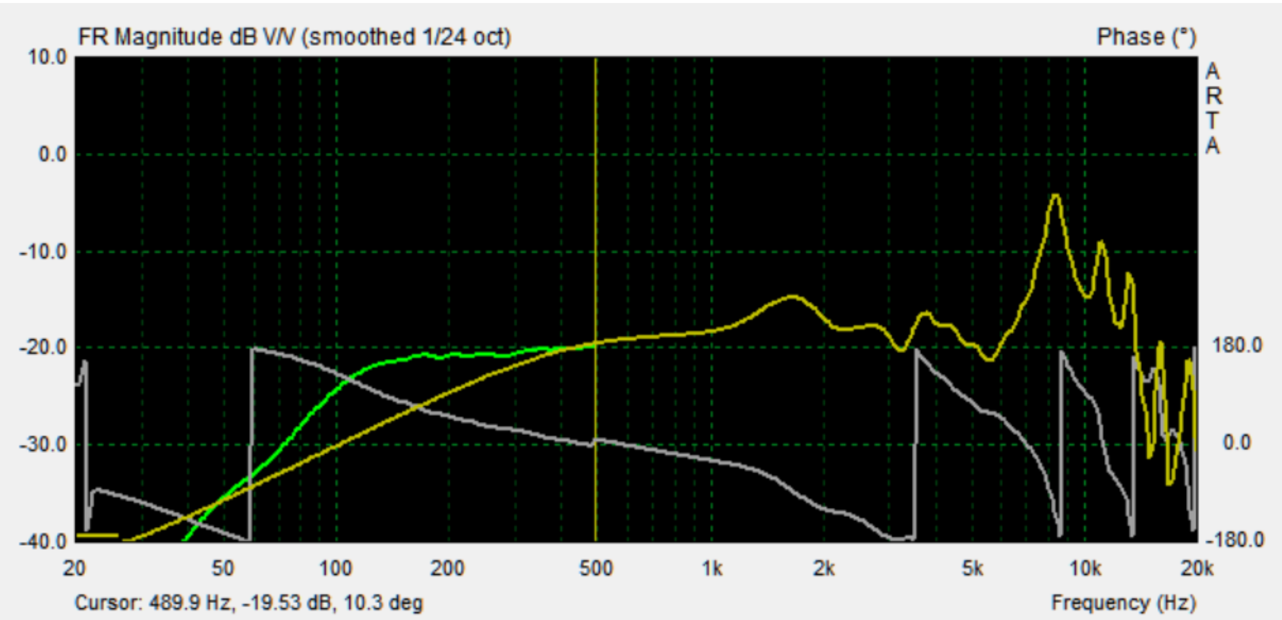

Next click on overlay and delete the orange trace or first overlay and you will end up with this.

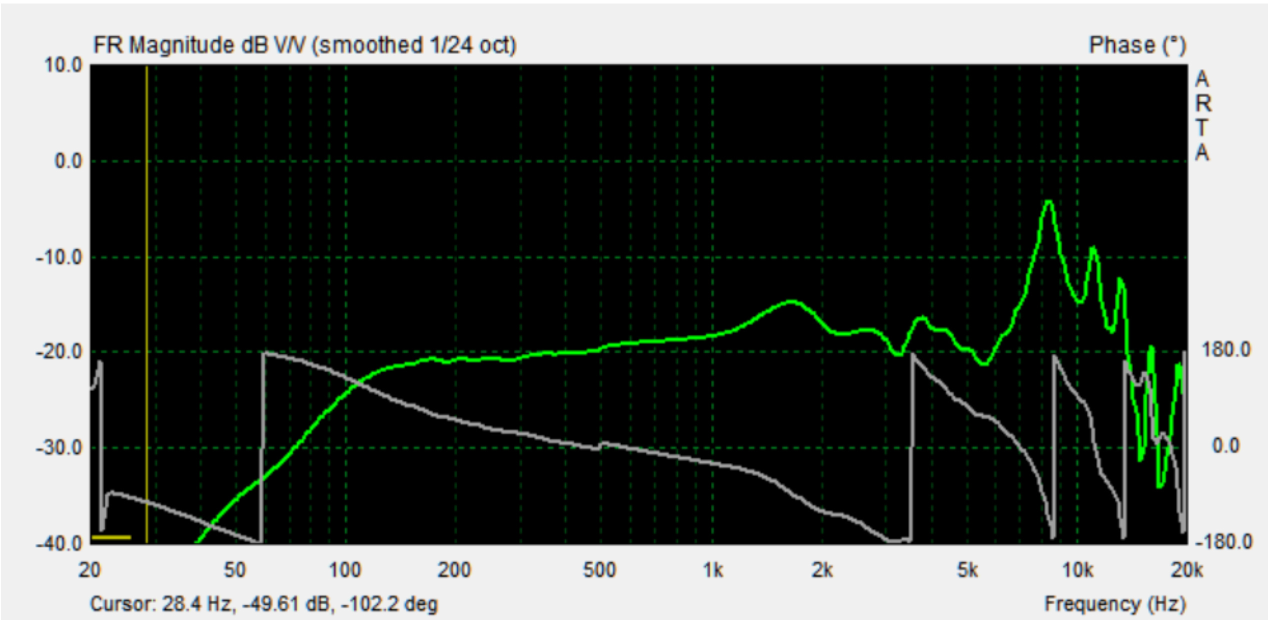

Export the file and you are done.

You will notice here that there is a little blip in the phase at 500Hz, this because the phase of the near-field wasn't matched perfectly with the phase of the far-field. If I had wanted to I could have altered the position of the 'start' of the gate in the near-field response to account for this. Given the situation though it is unnecessary as the phase below 500Hz isn't going to be used for anything. In this situation it is only the frequency response down low that is necessary for making sure that you've compensated correctly for baffle-step losses.

As you can appreciate, going back and altering the start position of the near-field data and then going through all that again just to see if they match up better is a bit of a pain. JustMLS actually allows you to alter the phase on the fly so you can match the two up perfectly.

Hopefully you will have noticed that we never once manipulated the far-field response. We *always* tailored the near-field response so that it matched the far-field, this is critical because its the far-field response that is necessary for accurately setting the parameters required for the accurate simulation of the driver integration.

A couple of things to note here is that I have a full version of ARTA and this allows me to save the impulse responses. This makes the process of merging the near and far-field responses a LOT easier. If working with an unregistered version you will need to make sure that you do all of your measurements in the correct order otherwise you will end up getting quite frustrated! Another thing to note is that the hardware I use is very reliable, that is if I take a measurement once and then take it again, the time delay and thus the start point of all of the impulse response remains the same. There are some situations where the delay in the hardware, between input and output, changes each time a measurement is performed. This is unsuitable for performing measurements, so make sure that your hardware is suitable, by repeating the same measurement and checking that your impulses are arriving at the same time.

If you are wanting to design passive crossovers then you will also need to measure the drivers impedance responses too. This is quite simple and the guide that is supplied with LIMP (this program comes with ARTA) is more then enough to get you making nice impedance measurements.