

Application Note

7

Digital Audio FIR Crossover

Highlights

Importing Transducer Response Data

FIR Window Functions

FIR Approximation Methods


■ Design Objective

3-Way Active Crossover

200Hz/2kHz Crossover Points

FIR Optimized Response

48kHz Sampling Rate




This design will provide an example of how to integrate many of the powerful features in the software, to produce a custom digital FIR crossover.

The audio spectrum has a Log frequency characteristic, while FIR filters are best suited to Linear frequency work. Using FIR filters across the entire audio frequency range can be difficult. For the case of IIR filters, the order of the filter is not so dependent on the location of the filter's corner frequency. FIR filters will generally require much higher order to produce equivalent response at low frequencies.


■ Prototype Data

We will assume that a prototype 3-way enclosure has been constructed, similar to that shown above, and that SPL data has been obtained for each of the three sections: Woofer, Midrange, and Tweeter. This data is measured prior to the use of any crossover, and reflects the response of each section measured separately.


The three graphs on the following page show the SPL Magnitude and Phase for each of the three transducers, measured using a 10Hz to 40kHz Log frequency range. Measured phase data may or may not contain the absolute position (delay) information. In this case it does not. The SPL phase measurements for each of the three transducers are relative to the acoustic origin (voice coil) location. Consequently, we will need to handle the true acoustic time position separately.



Woofer SPL Response



Midrange SPL Response




Tweeter SPL Response

■ **Relative Transducer Delay**

To determine the actual time positions of each transducer, a *design point* in space must be selected. Unless all of the transducers in a multiway acoustic system are coaxial mounted, an ideal crossover can only be designed for a single point in space. The drawing below illustrates the situation.

We know the acoustic response for each transducer, relative to their acoustic origins, but we must also determine any differences in delay.



At high frequencies, the acoustic origin of each transducer is very near the voice coil. This is where the conversion from electrical to acoustic wave propagation takes place.

The Tweeter is used as the reference since its voice coil is nearly in the plane of the baffle board. We will also assume the Design Point is on-axis with the Tweeter, and that the original SPL response for each driver was measured at the Design Point.

The voice coil of the Midrange driver is back set 0.033M and the Woofer is back set 0.058M.

However the total path difference between the drivers to the Design Point in space depends on the chosen reference distance, and the vertical spacing between the drivers.

At very far distance, the vertical driver spacing is relatively unimportant, and the path difference between the drivers becomes equal to the voice coil off-sets. At closer distances the vertical driver spacing must be included in path computations.

Typical reference distances are usually given as either 1 or 2 Meters for most loudspeaker products. For this example we shall use 1 Meter as the reference distance. We shall also assume that the vertical spacing between each driver is about 6 Inches, or 0.15 Meters.

Using basic trig the acoustic path distances are:

$$\begin{aligned} D_{\text{TWEETER}} &= 1 \text{ Meter} \\ D_{\text{MIDRANGE}} &= \text{sqrt}((1+0.033)^2 + 0.15^2) = 1.044 \text{ Meter} \\ D_{\text{WOOFER}} &= \text{sqrt}((1+0.058)^2 + 0.30^2) = 1.100 \text{ Meter} \end{aligned}$$

Therefore the Midrange and Woofer path delay differences are:

$$\begin{aligned} \Delta D_{\text{MIDRANGE}} &= 0.044 \text{ Meter} & T_{\text{MIDRANGE}} &= 126\mu \text{ Sec} \\ \Delta D_{\text{WOOFER}} &= 0.100 \text{ Meter} & T_{\text{WOOFER}} &= 286\mu \text{ Sec} \end{aligned}$$

(Note: Speed of Sound constant 350M/Sec)

The Midrange acoustic response will be delayed by 126u Sec, and the Woofer will be delayed by 286u Sec. These delays must be included as the complete acoustic response for each section.

■ Crossover FIR Transfer Functions

There are many different kinds of FIR filters, and methods used to create them. For this example, one of the key characteristics is how the Lowpass/Highpass sections sum together. Ideally we desire that LP/HP pairs of filters sum to a flat magnitude response.

The location where the filter summing action is of most concern, is of course in the transition region of the filters. Optimal filters are designed based on criteria specified for the passband and stopband. The transition band between these is basically undefined.

There is generally no direct means of controlling the response in the transition bands of optimal filters, which is the primary area of concern for crossover work. Therefore optimal FIR filters are not so well suited to this application.

However, FIR filters designed by the window method are quite different. By their nature the center of the transition band is always -6dB, and LP/HP pairs do indeed sum to a flat response. This class of FIR filter is very well suited for crossovers.

There are many different window functions. Each offers different characteristics, and for this application we need a high performance window which provides very high stopband attenuation. We will assume that a stopband attenuation of 100dB is specified.

Some windows offer only fixed stopband attenuation, while others are adjustable. The Nuttall-4T30 would be a good fixed window choice, while the Kaiser, Dolph, VanDerMaas, and Taylor Series would also be good adjustable windows. Any of these would be suitable and provide similar results. The Kaiser window will be selected with a stopband attenuation of 100dB.

The next issue that arises is the specification for the width of the transition region. FIR transition width is usually specified on a Linear frequency basis. Audio crossovers are generally thought of as having the same width at each crossover point on a Log frequency basis.

To accomplish this same result using FIR filters, the transition width of each crossover point must be determined relative to its frequency location. The significance of this will become clear, and results in very different order requirements for the 200Hz and 2kHz filters. Essentially they will differ by 10:1.


Woofer FIR Design

For the Woofer section a 200Hz Lowpass is needed. We must now choose an Ft bandwidth for the transition region. A common specification for an equivalent analog crossover might be a 4th order Butterworth 6dB filter. To see how different FIR Ft values compare, we can first generate the analog filter and then run some comparisons with different FIR Ft widths.



The graph below shows the response of the 4th order BW6 Lowpass, along with several different FIR filters. As a starting point, an Ft value equal to the corner frequency is chosen and then increased. The parameters are:

- Kaiser, 100dB, Ft = 200Hz, N = 1534
- Kaiser, 100dB, Ft = 300Hz, N = 1023
- Kaiser, 100dB, Ft = 400Hz, N = 767
- Kaiser, 100dB, Ft = 500Hz, N = 614
- Kaiser, 100dB, Ft = 600Hz, N = 512

All have higher ultimate attenuation rates than the analog filter. The best fit to the analog filter passband knee lies in the Ft = 400-500Hz range. The best fit occurs with an Ft = 425Hz which requires an order of 720.



The analog filter and $F_t = 425\text{Hz}$ FIR filter are shown below. The knees of both are very similar, while the attenuation rate is much steeper in the FIR case. We could have probably gotten by with less attenuation if we needed a lower order.



■ Midrange FIR Design

For the Midrange section a 200Hz Highpass and 2kHz Lowpass is needed. The Highpass filter must be complementary to the previous Lowpass, so we must use the same order of $N = 720$.

Since we previously determined that an $F_t = 425\text{Hz}$ was suitable for the 200Hz filters, we can assume that an $F_t = 4.25\text{kHz}$ for the 2kHz filters will be required. However the 2kHz frequency is now much closer to the sampling frequency which causes some warping of the F_t/F_c ratio.

The Highpass FIR is first generated, and then moved/copied to TFB entry #2 in the Target Parameters. The Lowpass FIR is then generated which places it in TFB entry #1. When both are enabled in the M1 group, they are multiplied together to form the FIR Bandpass transfer function.



As before, an equivalent analog Lowpass/Highpass filter combination is created and shown for comparison. To maintain a similar match to the analog knee, it was found that a slight reduction in the F_t value was required for the 2kHz Lowpass. The two graphs on the following page show the comparison of the combined FIR filters and the analog model.

For this Midrange section the parameters are as follows:

- 200Hz Highpass, Kaiser, 100dB, $F_t = 425\text{Hz}$, $N = 720$
- 2kHz Lowpass, Kaiser, 100dB, $F_t = 3.38\text{kHz}$, $N = 90$

Note that the attenuation rate of the 200Hz Highpass is less than the analog model, and the rate of the 2kHz Lowpass is higher than the analog. This comes as a direct result of the Linear frequency behavior of FIR filters when evaluated on a Log frequency axis.

However, the passband and upper transition regions of the FIR and analog transfer functions are quite similar. These are the areas of most concern for a crossover application.



■ Tweeter FIR Design



For the Tweeter section a 2kHz Highpass is needed. The Highpass filter must be complementary to the previous 2kHz Lowpass, so we must use the same order of $N = 90$.

As before, an equivalent 4th order analog Highpass filter is created and shown for comparison in the two graphs on the following page.

For this Tweeter section the FIR window design parameters are as follows:

- 2kHz Highpass, Kaiser, 100dB, $F_t = 3.38\text{kHz}$, $N = 90$

Note that the attenuation rate of this 2kHz Highpass is greater than the analog model. The knee is also stiffer than the analog version.



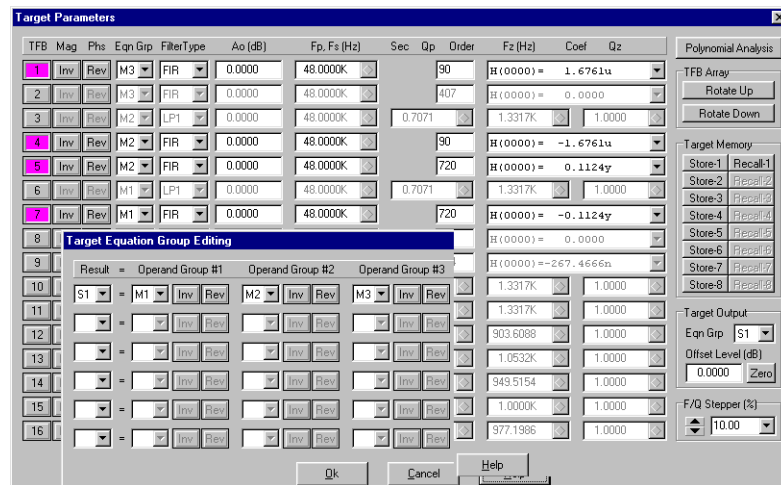
■ **Combined FIR Response**



We should now verify that the combined (summed) response of all three FIR bands yields a flat magnitude. To begin with we plot each of the three Woofer, Midrange, Tweeter FIR sections on the following page.


To illustrate how difficult it is working with Linear frequency scales for audio spectrum applications, note the second graph on the following page. The Woofer band is completely compressed and barely visible at the low end.

To produce the summed response of all three bands, the Target Parameters dialog is setup with all of the previously designed FIR filters. Since the three bands must be summed together, an equation group S1 is used with each of the three bands assigned to different multiply groups: M1, M2, and M3.

The main output group is now S1, which represents the sum of the three FIR filter bands. The Target Parameters setup is shown below, along with an insert showing the setup for the S1 equation group.

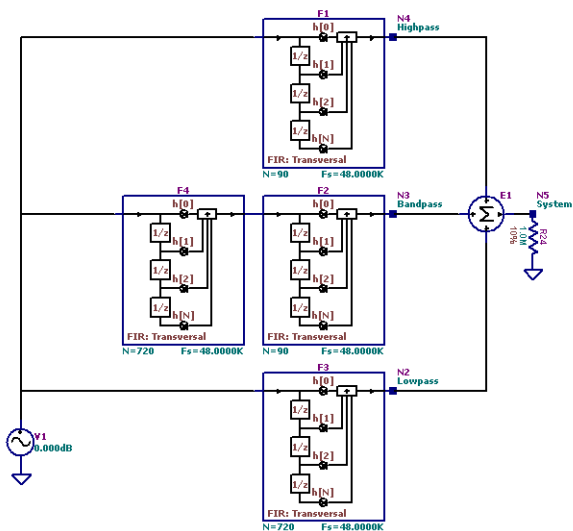






The graph above shows the summed response. There seems to be a problem. The crossover regions are not flat as we had expected. What happened?

The equivalent schematic for the system transfer function being created is shown here on the left. There are two FIR sections used in the Midrange band, while both the Tweeter and Woofer bands use only one. Now consider the total FIR order of each band:



- Woofer: $N = 720$
- Midrange: $N = 720 + 90$
- Tweeter: $N = 90$

They are all different. This means that there is a different signal delay present in the outputs of all three bands. For the summed response to be flat, the delay through each band must be the same.

The delay through any linear phase FIR filter is equal to 1/2 the order of the filter, multiplied by the delay of a sample period (1/Fs). The total delays are therefore:

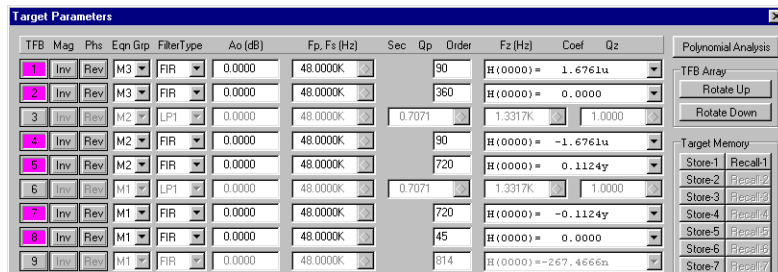
- Woofer: $N = 720/2/48k = 7.5m$ Sec
- Midrange: $N = (720 + 90)/2/48k = 8.4375m$ Sec
- Tweeter: $N = 90/2/48k = 0.9375m$ Sec

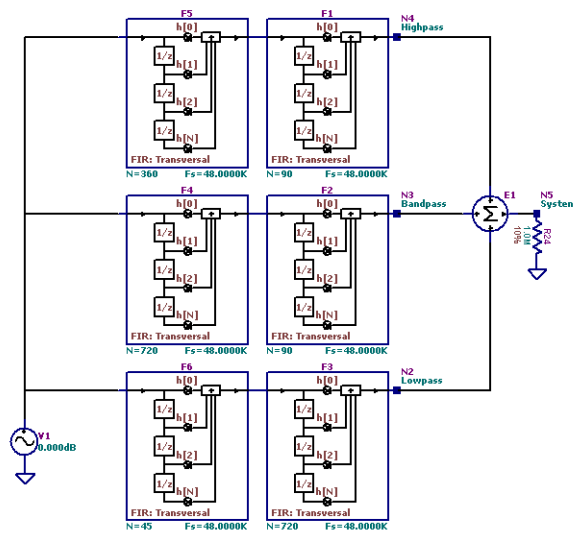

To equalize the paths, the Woofer and Tweeter bands require additional delay to yield the same total delay as the Midrange band. The additional delays are:

- Woofer: $8.4375m - 7.5m = 0.9375m$ Sec
- Tweeter: $8.4375m - 0.9375m = 7.5m$ Sec

Obviously these are complementary delays from each opposite band. To represent these new delays in the Target Parameters, we can add two TFBs setup as FIR type with 1/2 the order of the originals. To use these as pure delay, all of their coefficients are zero except for the N order coefficient which has a value of unity. This target configuration is shown below.

A schematic representation is also shown on the following page, along with the new summed response. It is now perfectly flat.





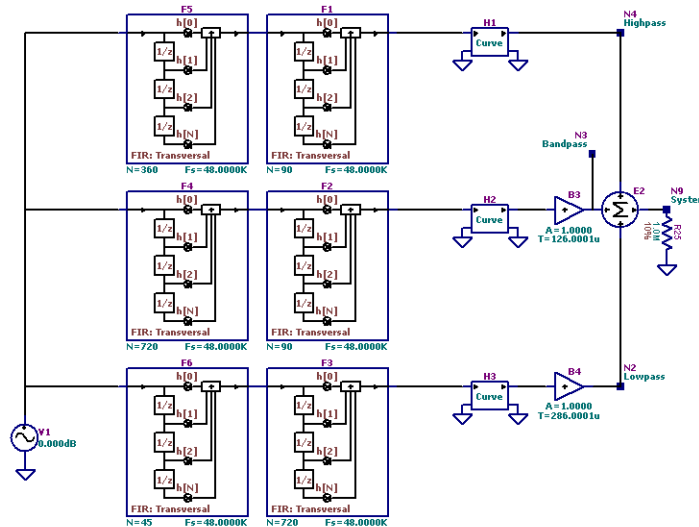
Acoustic System Response

If the response from each of the three transducers were perfectly flat, the FIR configuration just obtained for each band would be complete. However, the acoustic response of each transducer is not flat and corrective transfer functions must be added to produce the best possible system response.


To examine the total system response by including the acoustic response of each transducer, the previous schematic circuit can be modified to include the acoustic transfer functions. The SPL data are imported into Guide Curves, and then loaded into H transfer function block components in the schematic. Buffer components are also added to represent the acoustic delay values for the Woofer and Midrange.

The schematic below shows the complete system model including the FIR crossover sections and the transducer acoustic response.


The top graph on the next page displays the response from each band along with the summed system response. Note that the high crossover point shows a serious dip. This indicates that the polarity of the Tweeter band should be reversed. There are several places where this can be done, in this case we change the sign of the single unity coefficient in the F5 delay filter. The bottom graph shows that result, and both crossover regions are now correct.



System Response, All Outputs Same Polarity



System Response, Tweeter Polarity Reversed



■ Acoustic Response Correction

As expected the overall response is far from optimal. Besides the different sensitivities of the transducers, there are many peaks and dips in the acoustic response that we would like to correct.

Correcting the response could be performed in two different ways: (1) correct the response in each band, or (2) correct the full system response before the bands.

We could attempt to fix each band individually, but there is a problem with this approach. If we include the FIR band response, we will end up having to perform a FIR approximation into stopband regions. That is problematic. We could add new FIR filters to each band, that only correct the acoustic problems, but that would require three additional FIR stages.

The most efficient way to correct the acoustic response is to add a single FIR filter in front of the crossover bands that fixes the entire system at once. This will automatically handle the shelving gain differences as well for each band.


Since the total system response is now up in the SPL range, we must make a copy of the system response in a Guide Curve, and then scale it down to a level more in the range of 0dB. A scaling level of -92dB was chosen. After this, the scaled system response is inverted to produce a correction curve. This is done by exponentiation using a value of -1. Both of these operations are performed using the *Processing | Unary Math Operations* dialog. The resulting scaled system response and inverted response is shown in the two graphs on the following page.

The method we will use for generating the FIR filter from the system correction curve will be *Digital-FIR: Frequency Sampling*. This method uses data from the *Amplitude* graph, which is the analytic magnitude response. We must check this data first before running the sampling dialog.


The amplitude response is shown on a following page. Note that the basic polarity of the curve is negative. However, there are also many small spikes in the curve. Now note the Phase graph below. The phase is rotating very fast, and this is causing aliasing errors due to the number of data points with Log frequency.

The large amount of phase rotation is due to the fact that the system response contains the full delay from the FIR filters and acoustic transfer functions. We don't need or want this for the correction curve.


System Response, Scaled -92dB




Scaled System Response, Inverted

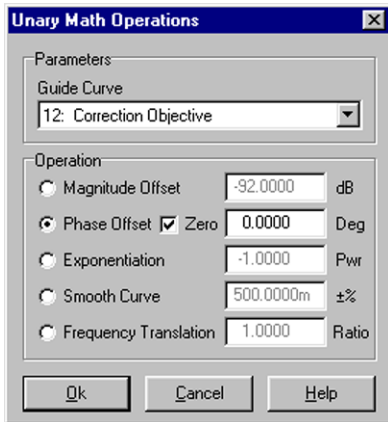


Amplitude Response for System Correction Curve



Phase Response for System Correction Curve









We can dump the phase/delay by using the *Processing / Unary Math Operations* dialog on the correction objective Guide Curve, this time with the Phase/Zero option and an offset of 0.0. This clears the phase to zero everywhere. The amplitude data is automatically updated.

The resulting amplitude response for the correction objective curve is shown below. This is what we want, a nice clean curve.

We now run the *Digital-FIR: Frequency Sampling* dialog to create the new FIR correction filter. The order used will determine how low in frequency the correction objective curve can be approximated. The lower we want to go, the bigger the order.

As an initial guess we shall try N=500, and use the other dialog defaults. The magnitude result is shown in the top graph of the following page. The result is not very smooth. Note that the FIR curve passes through the objective curve at the lowest frequency of 20Hz. Because the level is being pulled up to +16dB, it causes the remaining approximation to contain heavy ripples. This is a direct result of the *Zero Freq Sample - Include* option being used. It forced the approximation through the curve at zero Hertz.






Clearly if we wish to hold the approximation all the way down to 20Hz, a much larger (higher order) FIR filter would be required. In this case we will assume that we are not that interested in equalizing the response down to these low frequencies, and will make a compromise.


We can give up holding the lowest frequencies by simply *excluding* the zero sample from the processing. In this case a different set of data points is chosen to approximate the curve, and it will not include zero frequency.

This result is shown in the second graph of the previous page. This looks much better. It appears that the first point used in the approximation is now around 50Hz. The next points are 150Hz, 250Hz, 350Hz, etc.



The FIR filters just designed were both Linear Phase. For this type of response, linear phase is unimportant since most of the curve is roughly flat, and the linear phase needlessly adds a considerable delay to the signal.

Linear Phase FIR filters use 1/2 their order just creating the linear phase behavior. If we choose a Minimum Phase filter, we can cut the order in half with little change on the approximation resolution.



This result is shown in the third graph of the previous page. The two graphs here on the left show the phase response of the linear vs. minimum phase FIR designs.

We will use the FIR N=250 Min phase filter for the correction.

■ **Completed System Circuit**

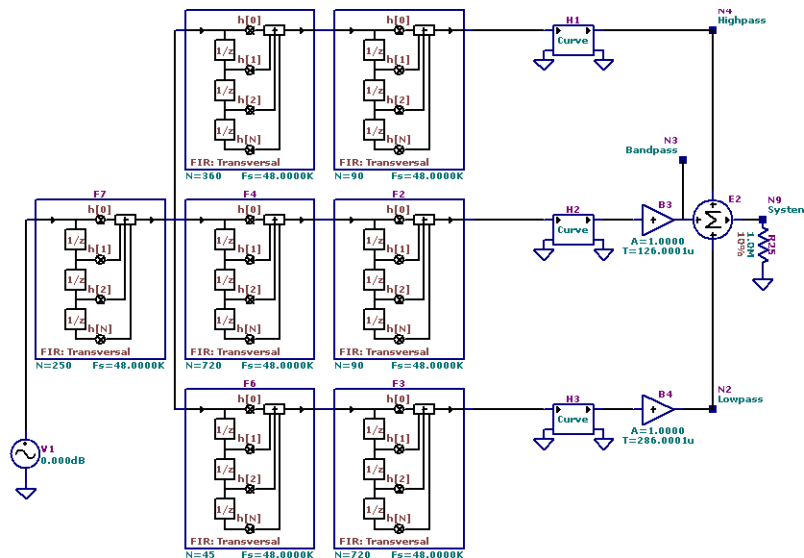
The circuit schematic is now modified to include another FIR component in front of the other band splitting FIR stages. The component is then loaded with the FIR data from the Target Parameters. The final schematic is shown below.

The two graphs on the following page show the final results for each band, and the combined system response. At mid and high frequencies the response becomes almost perfectly flat. The low end cutoff is -3dB at 30Hz.


■ **Summary**

Highly accurate digital FIR crossover designs can be easily obtained using the powerful FIR methods and extensive modeling features of the program. The seven FIR filters contain about 1060 actual calculation coefficients. At $F_s = 48\text{kHz}$ this is about 50M MACS. However, since the LP/HP blocks are complementary, one could be obtained from the other by subtracting from half delay. This would cut the coefficient calculation count to about 700 or 34M MACS.

This completes the Digital Audio FIR Crossover design.



Final Magnitude Response for each Band



Final Magnitude Response for System

