

Crossover Filter Shape Comparisons

A White Paper from Linea Research

Paul Williams July 2013

Overview

In this paper, we consider the various active crossover filter types that are in general use and compare the properties of these which are relevant to professional audio. We also look at a new crossover filter type and consider how this compares with the more traditional types.

Why do we need crossover filtering?


Professional loudspeaker systems invariably require two or more drivers with differing properties to enable the system to cover the wide audio frequency range effectively. Low frequencies require large volumes of air to be moved, requiring diaphragms with a large surface area. Such large diaphragms would be unsuitable for high frequencies because the large mass of the diaphragm could not be efficiently accelerated to the required velocities. It is clear then that we need to use drivers with different properties in combination to cover the frequency range, and to split the audio signal into bands appropriate for each driver, using a filter bank. In its simplest form, this filtering would comprise a high-pass filter for the high frequency driver and a low-pass filter for the low frequency driver. This is our crossover filter bank (or crossover network, or just crossover).

Common Requirements

When choosing an appropriate crossover filter type, there are several factors to be considered.

Magnitude response

Commonly referred to as "frequency response", this is a measure of to what degree the combined acoustic response of the system (assuming ideal drivers) remains at a constant level with respect to changing frequency. We would ideally like there to be no variation, but small deviations can usually be tolerated, particularly if other more desirable properties can be attained. Although perhaps the most obvious parameter, it is not necessarily the most important. It is after all the most readily corrected using equalisation on the 'input' to the system (or alternatively identical equalisation applied to each of the outputs from the crossover). However, one needs to be aware that applying such equalisation will have an effect on the Phase Response. Figure 1 shows the Magnitude Response of a well behaved crossover filter pair.



Phase response

This is a measure of how much the phase shift suffered by a signal passing from input to the summed acoustic output (again assuming ideal drivers) varies with changing frequency. The group delay is a parameter derived from phase response which describes how different frequencies are delayed. Since the group delay is proportional to the gradient of the phase response curve, it follows that a linear phase response produces a flat group delay, meaning that all frequencies are delayed by the same amount. We would therefore like the phase response to be as linear as possible. However, lack of phase linearity is not necessarily audible.


Polar response

In most loudspeaker systems, the drivers are not normally in the same vertical and horizontal position (in other words they are not coincident), the exception being the Dual Concentric type. Assuming the drivers are separated in the vertical plane, then it is clear that the distance between the listener and each driver can only be the same in one position. If the listener is above or below the centreline of the driver pair, then there will be a difference in distance, and therefore a difference in delay which the signal arriving at the listener from the drivers will suffer. These delay differences will cause attenuation, and ultimately cancellation at certain frequencies. On the centreline however will be the main lobe where there is no such attenuation. Any differential

phase shift between crossover filter bands will cause this main lobe of the acoustic beam from the loudspeaker system to tilt. If such phase shift varies with frequency, this tilting will be frequency dependant, resulting in colouration of the sound which will be strongly dependant on the listening position. The ideal filter bank will therefore maintain zero degrees phase difference between adjacent outputs across the crossover region, and ideally at all frequencies.

A Polar Response plot is essentially the magnitude response at different observation angles. The Polar Response is often referred to as the *Dispersion Pattern*, which can be different between the vertical plane and the horizontal plane for a given system. Here, we are interested in the *vertical dispersion*. Figure 2 shows the Polar Response of a system with vertically separated drivers, using a crossover filter pair in which the phase response is identical. This is plotted at the crossover frequency with driver separation equal to the wavelength of the crossover frequency. The scale lines are at 10 degree and 10dB intervals. It does not take individual driver directionality into account. In this case, we can see that the centre of the main lobe is at 0dB at all frequencies. We have plotted the Polar Response at three different frequencies. Note that the cancellation that is evident on the 1kHz plot has nothing to do with crossover filters, and would indeed be the same if the filters were taken out. This is purely a function of the dissimilar distance between the drivers and the observer, which will clearly change significantly for differing distances between drivers.

Fig.2 - Well behaved Polar Response



Separation/Band-Stop effectiveness

This is a measure of how effectively a filter output attenuates frequencies which are not intended for its driver. There are two issues: The rate at which the attenuation increases as the frequency gets further away from the intended band edge, and the ultimate

attenuation. The latter is rarely an issue, but the rate of attenuation (often referred to as the cut-off slope, or cut-off rate) is usually seen as an important consideration. Driver design, like many aspects of engineering, involve many carefully balanced design decisions. Problems such as cone break-up and narrowing of the acoustic beam at high frequencies present the designer with challenges, while over-excitation distortion or damage can become an issue at low frequencies, particularly below the effective loading frequency of horn loaded drivers. More effective stop-band attenuation in the crossover filters will allow the driver to be better protected from the problematic frequencies.

Frequently, some of these requirements are mutually exclusive, forcing designers to choose the best compromise.

Active or Passive

A crossover filter bank (network) can of course be implemented with purely passive components. In professional high-power systems which employ high voltages and high currents, the size and cost of the passive components can become prohibitive, and can lead to distortion, losses and degradation of the damping factor. These problems tend to be exacerbated when high filter orders are required.

It is often desirable instead to employ an active crossover filter bank, where the band-splitting is done before the amplification, requiring amplification to be provided for each frequency band (conceptually on each driver). Such systems are sometimes referred to as "Bi-Amping" (for two-way active crossovers), or "Tri-Amping" (for three-way active crossovers). Such systems also benefit from the ability to match the power rating of the amplifier to the power requirements of a driver; lower frequency drivers usually requiring an amplifier with a higher power rating than high frequency drivers. Furthermore in such a system, one crossover filter bank might be delivering a driver signal to several amplifiers, thus requiring fewer crossovers. Sometimes, passive and active crossover filters are used together in different parts of a multi-driver loudspeaker system, the passive filters being used on the highest frequency crossover, where voltages and currents are lower.

Only active crossover filters will be considered in any detail in this paper, although much of the discussion applies to both active and passive implementations.

Digital Or Analogue

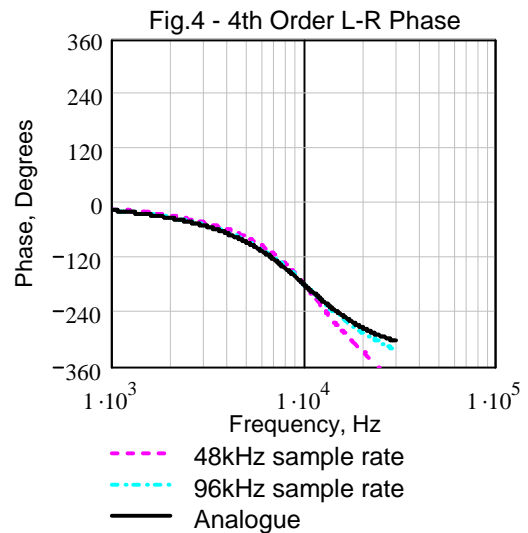

It is a commonly held misconception that digital implementation of filters in Digital Signal Processors (DSPs) behave in a fundamentally different way to their analogue counterparts. In fact, if a digital filter is designed properly to imitate an analogue filter design, then it will duplicate the magnitude, phase and group

delay characteristics perfectly; there will be no difference.

Digital filters can however start to depart from their analogue cousins when the filter frequency is set at very high values - particularly when this frequency approaches half the sample rate. Towards this region, the magnitude response and phase response of the filter will fail to perfectly match the analogue filter. There are various ways of mitigating this problem, the most effective being to simply use a high sample rate. A 96kHz sample rate almost completely eliminates the problem.

Of course, the digital version will be superior in terms of accuracy, and in terms of consistency from one instance of a product to another because it does not suffer the inconsistencies of component value tolerance that will plague the analogue filter. This inconsistency can often be experienced as shifts in the stereo image, where component value differences between the channels can cause small, but acoustically significant phase differences.

Figure 3 shows the magnitude responses of three different filter designs; a digital filter at 48kHz sample rate, a digital filter at 96kHz sample rate, and an analogue filter - all designed for a 4th order (24dB/octave) Linkwitz-Riley low-pass function with a crossover frequency of 10kHz. Even at this high filter frequency, one can see that the digital filters remain quite faithful to their analogue counterpart, particularly for the 96kHz design. At lower filter frequencies, the differences become very trivial. Interestingly, these minor departures in phase and magnitude for the digital filter are usually duplicated in the corresponding high-pass filter, resulting in the same summed and inter-band characteristics as an analogue design. We compare the phase responses in Figure 4.



Digital Filter Classes

When implementing filters digitally, there are two primary classes of digital filter which can be used: Infinite Impulse Response (IIR), and Finite Impulse Response (FIR).

IIR

The IIR filter class is often used in digital crossover filtering because this is not only efficient in terms of DSP resource, but is also perfectly suited to reproducing the characteristics of tried and trusted analogue filter designs. IIR filters, just like their analogue equivalents are almost without exception *Minimum Phase*. This means that the phase shift suffered by a signal passing through the filter will not necessarily be linearly related to the frequency of that signal; rather, the phase shift is the minimum amount that is required to produce the filtering action at a given frequency, and will thus vary with frequency.

FIR

A filter design using an FIR would not normally be based on an analogue filter 'prototype', but would instead implement a filter which could not easily be implemented by analogue means, or indeed not exist in the natural world at all. The impulse response of an FIR is finite because it reconstructs the required impulse response by using a set of delay taps from which a weighted sum is produced. There can only be a finite number of taps, so the length of the impulse response is similarly finite. If an FIR filter is designed to have an impulse response which is symmetrical either side of the impulse peak, it will have a *Linear Phase* characteristic, that is where the amount by which the signal is phase shifted through the filter is linearly related to the signal frequency, which results in all frequencies being delayed through the filter by an equal amount. It must be stressed that the terms "Linear Phase" and "FIR" are not synonymous; An FIR filter can easily be designed to produce a non-linear phase response, and there are also other ways of producing a linear phase response without using an FIR filter, as we shall see later in this paper.

Whilst this is a worthwhile characteristic to strive for, one has to understand the ramifications of achieving Linear Phase using an FIR filter.

In order to provide any filtering action, it is clear that a filter cannot 'know' how much to attenuate a signal until it has 'acquired' enough of its waveform to determine what the frequency is. All filters will thus delay the signal by an amount which is related to the wavelength of the signal. Generally speaking, a Minimum Phase filter will delay the low frequencies more than the high frequencies, thus blurring the arrival times of different frequencies.

To achieve Linear Phase operation, this delay does not depend on the frequency of the signal; instead it is a fixed constant. A Linear Phase filter is effectively 'padding out' the delay on the high frequencies to that required to be able to provide the necessary filtering action on the low frequencies.

So, a Linear Phase filter will introduce a delay, as required by the laws of physics. There is thus an inescapable relationship between the amount of low-frequency filtering detail attainable, and the propagation delay, which may be unacceptable in some live sound situations.

Of course, one has to put this delay in context. The distance of the loudspeaker from the listener might be several tens of feet (and thus introduce several tens of milliseconds of delay), so a few added milliseconds in the filters may be of little consequence.

LIR

The Linear Impulse Resonse (LIR) filter [10] is a relatively new technology which was developed to achieve Linear Phase operation with the very lowest latency delay possible. LIR crossover filters will be explored later in this paper.


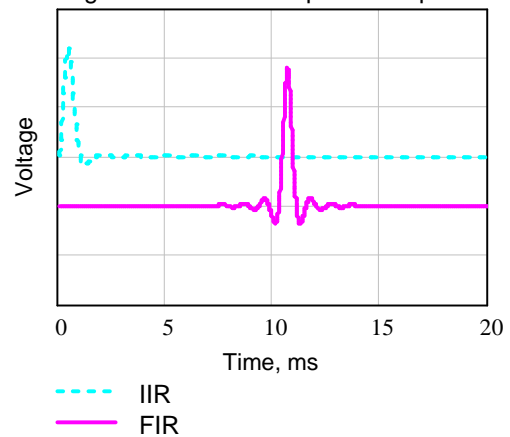



Fig.6 - IIR and FIR Impulse Response



Secondly, depending on how the FIR is designed, it can also exhibit a phenomenon known as 'pre-echo' which can be audible on percussive sounds, resulting in blurring which listeners can find very unnatural. The impulse response of an FIR showing excessive pre-echo is shown in Figure 7. There are decaying ripples either side of the main impulse peak at 11ms, which are a natural consequence of filtering, and are discussed in the next section. At about 6ms before and 6ms after the main impulse however, one can see a pre-echo and a post-echo. The post-echo will most likely be masked by the ear, but the pre-echo will almost certainly be audible. This is a particularly extreme example, but is an example of what can go wrong when an FIR is not designed with proper attention to detail.

Linear Phase filters most definitely have their place, but before using an FIR for this purpose, one has to understand their potential problems.



Brick-Wall Filtering

We will see later in this paper that some of the problems with crossover filters are associated with the crossover region; that is the range of frequencies where both bands are contributing to the summed output. This, together with the fact that the loudspeaker designer might well wish the frequency range presented to each driver to be tightly constrained, might lead one to assume that the steeper the cut-off slope the better. In


fact, why not have an infinitely fast slope so that there is no crossover region, and absolutely no signal outside the prescribed frequency range for a driver? In other words, a Brick-Wall filter.

Very high order filters (see the section on Filter Order) are capable of very high attenuation slopes, but will usually, with analogue and IIR digital filters, be associated with severe phase problems as we will see later. FIR digital filters of sufficient order (length) can conceivably achieve close to Brick-Wall operation, but there is of course a down-side..

When we filter a signal and remove a part of the spectrum, we are removing some of the Fourier series (the infinitely large number of sine waves of differing frequencies that would be needed to represent a complex signal). A low-pass filter at 1kHz is removing the Fourier components above 1kHz, and a high-pass would remove the Fourier components below 1kHz. This truncation of the Fourier series gives rise to a ripple in the impulse response due to the Gibbs phenomenon [1]. This is not the result of any imperfection in the filter, but an inescapable result of physics. It is easy to see why when you consider that, to adequately represent a square wave using a series of sine waves, there must be many high-frequency sine waves adding together to produce the vertical sides of the square-wave. If we were to remove the high frequency sine waves (by low-pass filtering), then the few low frequency sine waves that remain would not be capable of representing anything like a square-wave but rather a wobbly approximation to one, along with much ripple. Figure 8 shows the Impulse Responses of two different filter designs with different cut-off slopes. It can be seen that the design with the faster cut-off slope has Gibbs ripples which last for a longer period of time. Note that the responses have been separated vertically for clarity. Also shown in Figure 8 is the impulse responses for a complementary high-pass brick-wall FIR. If we construct a crossover filter bank with such a complementary pair of filters, the Gibbs ripples are also complementary (since we will expect the filters to sum to a flat response with linear phase, producing a perfect impulse). The summed result will thus be free of any Gibbs ripple, so what's the problem? The problem is *off-axis*. The complementary ripples will only cancel if the delay suffered by the signal from each driver is identical. Off-axis, where the path lengths differ, the ripples will not cancel, leading to the possibility that Gibbs ripple might become audible (just like a high-Q ringing filter). Such summing errors will be more pronounced at higher crossover frequencies because the ripples are more closely spaced. Lower frequency crossovers will have wider ripples which will more easily cancel in the presence of off-axis induced delays.

It is evident that steeper cut-off slopes give rise to Gibbs ripples of greater duration. It makes sense therefore to restrict the cut-off slope to be no more than is necessary for the application.

Fig.8 - Brick Wall Impulse Response



Filter Order

For much of the remainder of this paper, we will concentrate on analogue crossover filters, or their digital equivalents (including Linear Phase).

Each 'edge' of a crossover is realised using a filter. This filter may be a very simple affair, consisting of nothing more than some capacitance and some resistance, or it may involve the equivalent of a series of cascaded second-order filter sections. In the former case, this is said to be a first-order filter, whose magnitude response will fall off at 6dB/Octave. In the latter case, we might have for example an eighth order filter, whose magnitude response might fall off at 48dB/Octave. The Order of the filter is effectively the number of reactive elements it has in it, or mathematically, how many 'poles' it constitutes. As a general rule of thumb, the magnitude response will ultimately fall off at (Order X 6dB/Octave). This is not always the case however. Some filters have a more complex shape, making this rule difficult to apply. So, although the term "24dB/Octave" for example is often used to refer to a rank of filter, "4th Order" would be more accurate. A high-order filter can be thought of as comprising a number of second-order filter sections in cascade. Each of these sections may well have a different cut-off frequency and Q value (a measure of the amount of resonance the filter exhibits). When connected together however, they produce the required low-pass (or high-pass) response. There are many ways of implementing high-order filters, which differ in the frequency and Q values used in the individual second-order sections, and each has a name (such as Butterworth or Linkwitz-Riley). Each of these filter shapes, or *alignments* has its own merits as we will see later.

Before we leave the topic of filter Order however, it is worth discussing the region where the low-pass and the high-pass filters combine; the Crossover region. Since the filters have a finite cut-off slope, the acoustic output from the drivers connected to each filter output will clearly combine in some way when the signal frequency is within the crossover region. Most filter types will cause phase shifting of (Order *times* 45 degrees) at the

crossover point. A low-pass filter will create a negative phase shift, and a high-pass filter a positive phase shift. A 4th order crossover filter pair will thus usually produce a -180 degree shift on the low-pass output, and a +180 degree shift on the high-pass output. The difference is 360 degrees, which is the same as 0 degrees, so the drivers will be in phase. A second order filter pair however will cause a phase difference of 180 degrees between the drivers. For this reason, it is necessary to apply a phase inversion to one of the drivers to correct for this. Odd-order alignments will result in a phase difference between drivers of 90 degrees which cannot be resolved using polarity inversion. Odd-order alignments will also exhibit non-ideal polar behaviour. The 3rd order Butterworth Polar Response is shown in Figure 9. It can be seen that the main lobe is not on-axis, but rather points downward, to different degrees depending on frequency. This tilting of the polar response will cause colouration of the signal, and this colouration will depend on the listening position. It is interesting to note that the on-axis magnitude response is entirely flat and is at 0dB. On-axis therefore, the response will be colourless. However, more energy is being pumped into the environment off axis than is being produced on-axis. This energy has to be delivered by the expensive plant you would wish to deliver useful dB rather than just exciting the reverberant field. This off-axis component, which varies with frequency, may well reflect to listening positions and thus be audible. The polar response of the 1st order alignment in Figure 10 shows similar asymmetry.

Fig.9 - 3rd Order Butterworth Polar





Fig.10 - 1st Order Polar

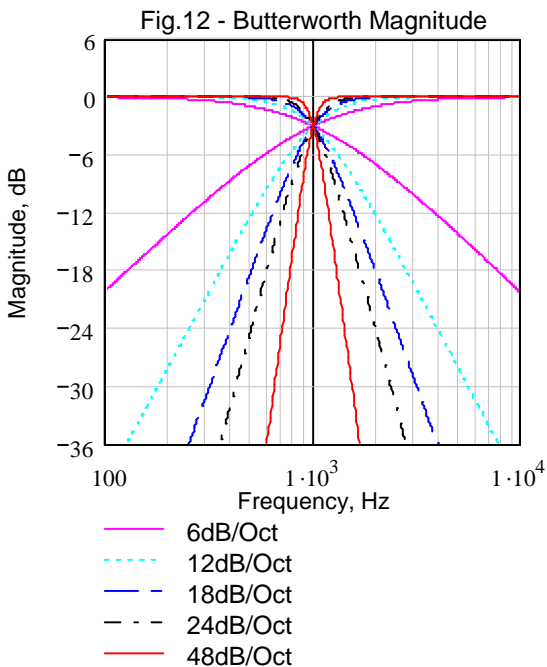
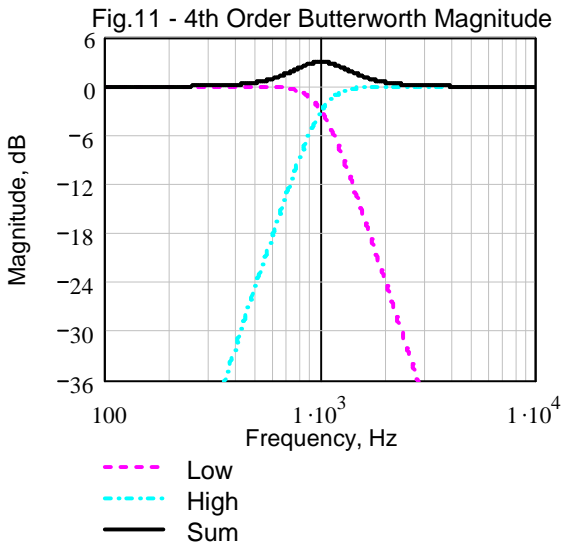


Common Filter Types

Various types of filter are used for crossover filtering in professional audio. A given filter type is usually referred to as an '*alignment*'. The most common crossover alignments in use in professional audio are Butterworth, Bessel, Linkwitz-Riley, and more recently Hardman. Here, we will make some comparisons between these alignments. Most of the examples will use a crossover frequency of 1kHz and will assume a 4th order alignment in each case. We will also introduce a new type of crossover filter, Linear Impulse Response (LIR), and illustrate this in the comparisons.

Butterworth

The Butterworth alignment offers what is termed a 'critically damped' response. That is, when a step is applied to its input, the output will attain the same voltage as the input in the minimum time possible, but without any overshoot. Figure 11 shows the magnitude response of a 4th order Butterworth crossover. You will notice that there is a peak in the summed magnitude response at the crossover point. Figure 12 shows the family of magnitude responses for filter order 1 thru 8. Order 1 (first order) is a special case in that, although technically it meets the criteria of a Butterworth filter, it is necessarily a single pole filter whose Q value cannot be altered. It is thus not possible for there to be any variants of a first order filter. Interestingly, the first order crossover is linear phase (in that the group delay is entirely flat). Such a crossover rarely finds practical application however due to the poor cut-off slope, and thus poor protection for HF drivers. Butterworth filters are not uncommon in passive crossover networks, particularly the 3rd order variant, which offers a reasonable compromise between complexity and HF driver protection. As discussed previously however, odd-order crossovers exhibit an asymmetrical polar response.



Bessel

The Bessel alignment is an interesting one, in that the intention of Thompson [2], who is usually credited for its development, was to create a filter which would produce a uniform delay across all frequencies within a certain bandwidth limit, rather than to perform a magnitude filtering function to attenuate certain frequencies. A side effect is that the Bessel filter does remove high frequencies (the classic Bessel filter being low-pass in nature, but may be transformed to high-pass). We can use this side effect to produce a crossover filter with very good phase characteristics, because the group delay is essentially constant across the bandwidth of the filter. A major drawback of the Bessel filter is that the cut-off slope is very slow, so it offers poor protection for tweeters when used as a crossover. It also droops approaching the crossover point, causing a dip in the summed magnitude response.

But, because the Bessel filter is intended for delay purposes, the 'design frequency' is not the crossover frequency, since the former is the frequency up to which the filter holds the delay constant. There is thus not a recognised standard way of 'scaling' the low-pass and high-pass partners of a Bessel crossover network. Various amounts of overlap/underlap are to be found in the audio industry, often based on the 3dB point of the filters (as shown in Figures 13 and 14), or an overlap which produces the best flatness of magnitude response (as shown in Figure 15 and 16). Both of these produce a poor polar response. A more satisfactory scaling [3] of Bessel filters however is that which gives the best polar response, but this is seldom seen in the industry. It is in fact possible to design a Bessel crossover filter pair which produces a perfect polar response, where the inter-driver phase is zero degrees at all frequencies, producing a main lobe in the polar response which does not move away from the on-axis position. Such a design is shown in Figures 17 and 18.




Fig.14 - Bessel - 3dB scaling




Fig.15 - Bessel - Best Magnitude scaling




Fig.18 - Bessel - Best Polar Scaling





Fig.16 - Bessel - Best Magnitude scaling



Linkwitz-Riley

The Linkwitz-Riley alignment [4] is doubtless the most widely used active crossover alignment in the professional audio industry. It uses cascaded pairs of Butterworth filters in combination to achieve -6dB attenuation at the crossover point in each of the High-pass and Low-pass filters which achieves a summed magnitude response which is completely flat. Not only that, the phase response is identical between the two filters so that the phase difference between adjacent drivers is identical, so the polar response is rock-steady, producing a main lobe which is precisely at 0 degrees.

Because of the way the constituent Butterworth filters are combined, Linkwitz-Riley alignments are all even-order. Figure 19 shows the Magnitude response of different Orders of Linkwitz-Riley alignment.

Fig.17 - Bessel - Best Polar scaling

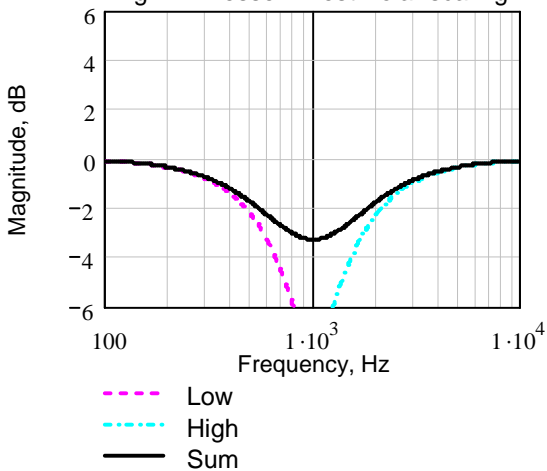
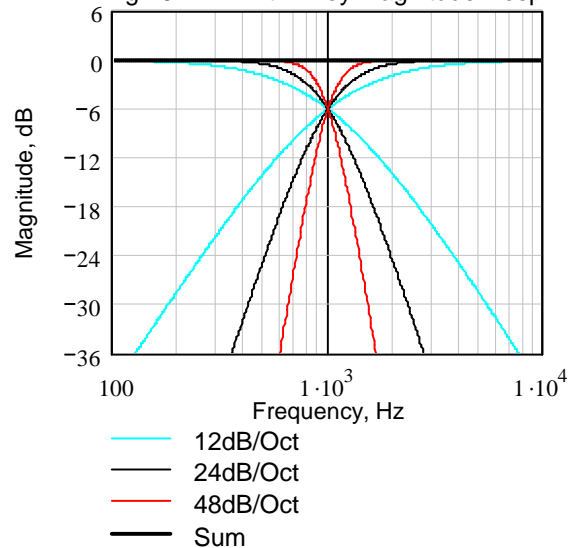



Fig.19 - Linkwitz-Riley Magnitude Resp.



Hardman

The Hardman crossover alignment [5] is a relatively new alignment which was introduced to the industry by Linear Research.. It has most of the benefits of the Linkwitz-Riley filter pair, but with significantly increased cut-off slope towards the stop-band, with a shape sometimes known as "Progressive Slope", because the cut-off does not resolve to a straight line, but rather it becomes progressively more steep until towards minus infinity, it tends towards infinitely steep. Figure 20 illustrates this for 4th order and 8th order Hardman alignments. Note that we have extended the vertical scaling on this plot to show the changing slope in the stop-bands.



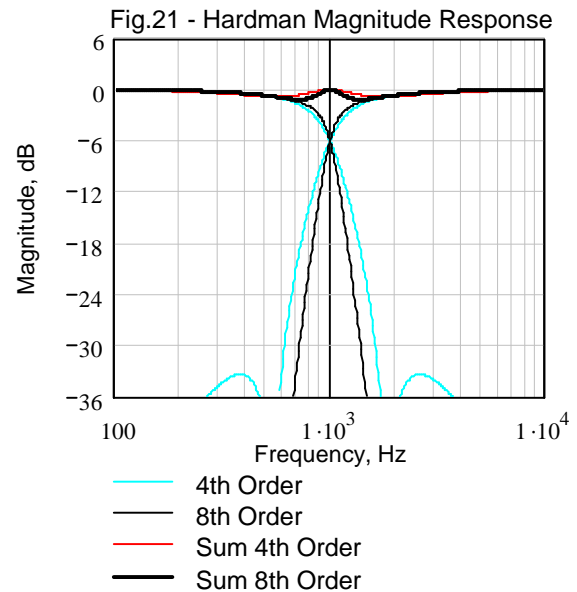
The steeper cut-off slope of the Hardman alignment is likely to be useful to driver and cabinet designers, often allowing the operating bandwidth to be increased since the rapid attenuation rate will allow operation 'closer to the edge', beyond which might be such perils as cone break-up, beaming, resonance or over-excursion. Since the Hardman filter achieves steeper cut-off slopes for a given filter order, another possibility is that a reduced filter order might be used, often thought to improve sonic transparency due to the lesser phase disturbances. An application normally requiring an 8th order crossover, or which would benefit from an 8th order crossover if it were not for the sonic drawbacks, could often be replaced by a 4th order Hardman crossover resulting in improved sonic qualities. Alternatively, an application requiring a cut-off slope exceeding that of a conventional 8th order crossover could be achieved with an 8th order Hardman filter. Even though the Hardman filter possesses the feature of steep cut-off slope, the phase response is identical to that of a Butterworth filter of the same order. This is because the pole positions are identical between Hardman and Butterworth; it is only the zeros which differ.

It is important to stress that the Hardman filter still preserves zero degrees phase shift between drivers at all frequencies, just like Linkwitz-Riley.

The small penalty paid for this functionality is a very

small ripple in the summed magnitude response (0.8dB p-p for 4th order and 1.1dB p-p for 8th order), and a slightly reduced cut-off slope far away from the crossover point ($(\text{Order}-1) \cdot 6 \text{ dB/Octave}$ rather than $(\text{Order} \cdot 6 \text{ dB/Octave})$). However, the transition of slope from rapid cut-off to this steeper slope occurs only when the attenuation has reached a useful degree (33dB for 4th order, and 62dB for 8th order).

Hardman alignments are necessarily even order, the minimum being 4th order. Figure 21 shows the individual and summed magnitude responses for 4th order and 8th order Hardman alignments.



Linear Phase – LIR


Phase linearity is a topic that is becoming ever more important to the industry for both practical and sonic reasons, some of which are:

- Absence of Group Delay Distortion
- Preservation of square-wave response
- Simpler integration of enclosures with differing crossover frequencies
- Simpler design of systems using overlapping crossovers

The LIR (Linear Impulse Response) crossover alignment [10] is a new alignment to the industry; a major development by Linear Research resulting from recognition of the shortcomings of other approaches. It has all of the benefits of the Linkwitz-Riley filter pair, but exhibits linear phase. Unlike many FIR based linear phase crossover implementations however, the LIR filter has ultra-low latency, and is as simple to use as a Linkwitz-Riley crossover.

Figure 22 shows that the Magnitude Response looks very much like that of a Linkwitz-Riley, crossing over at -6dB as usual. Being linear phase, the phase responses in opposing bands are identical and so the polar response is always on-axis.

The LIR alignment will be compared against other alignments later in this paper. A deeper analysis is beyond the scope here, but may be found in the referenced paper [10].

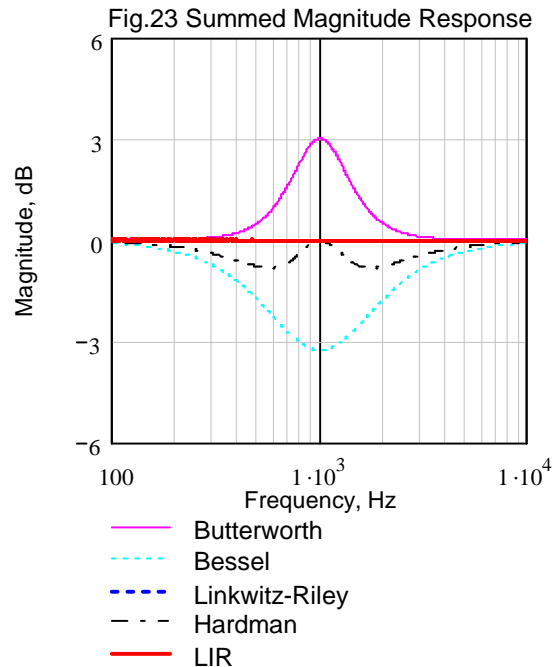


However, it is useful to compare each of the main features for each alignment so that it is possible to place each alignment in order of superiority.

We will make these comparisons for the 4th order case of each alignment, showing the traces for each alignment on a single graph for each feature, assuming a crossover frequency of 1kHz. In each case, the curve for alignment (or alignments) which is considered to be the best is shown in bold.

Summed Magnitude Response

We show the summed on-axis magnitude response for each of the alignments under consideration in Figure 23. We can see that the Linkwitz-Riley and LIR are the only alignments which have a completely flat response. The Hardman alignment is almost flat, while the other alignments suffer either a dip or a peak in the response. However, as discussed earlier, this anomaly may be easily rectified using equalisation.



Polar Response

When properly designed, all of the 4th order alignments discussed in this paper will exhibit polar behaviour which puts the main lobe precisely on-axis, which does not move with changing frequency.

If the Bessel filter pair are properly designed (to have identical phase characteristics between the high and low pass filters), then there is nothing to choose between the various even order alignments in terms of polar response. As discussed previously however (in the Bessel section), if Bessel crossover designed for a -3dB crossover point, or is designed for the flattest summed magnitude response, then the polar response will be poor.

As discussed in the Filter Order section above, odd

Summary

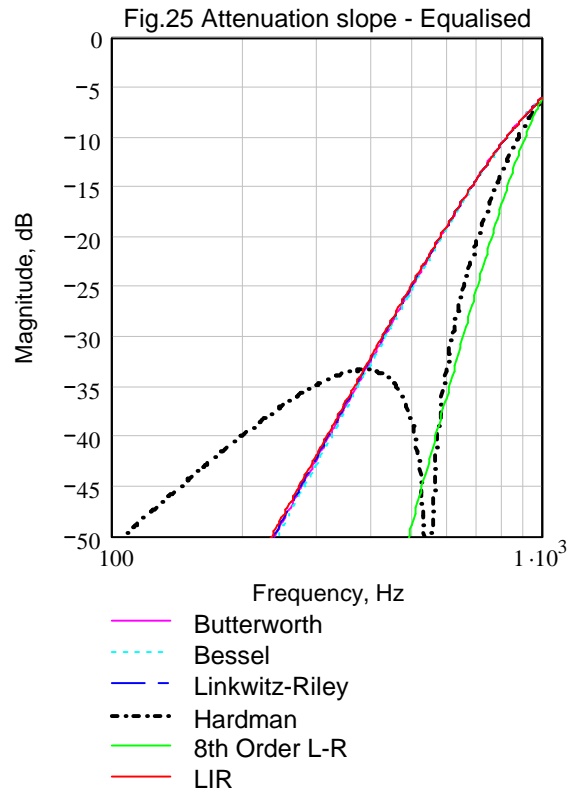
So, which crossover alignment is best? As you might have expected, there is no one alignment which is better than all others in all respects. They each have their merits and drawbacks, and the final selection will depend on which of these factors is of most significance to the application.

order alignments will exhibit non-ideal polar behaviour. In this paper the only odd-order alignments we have considered are 1st order and 3rd order Butterworth. These may therefore be thought of as the poorest in terms of polar response. The polar response of all the even-order Butterworth crossovers is perfectly symmetrical however.

Attenuation Slope

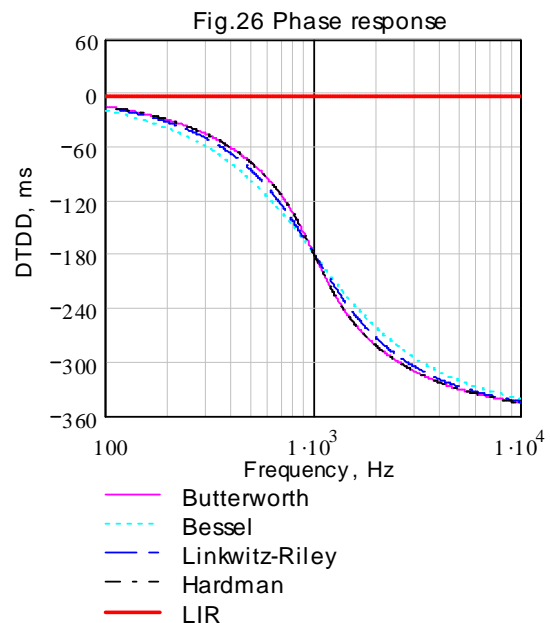
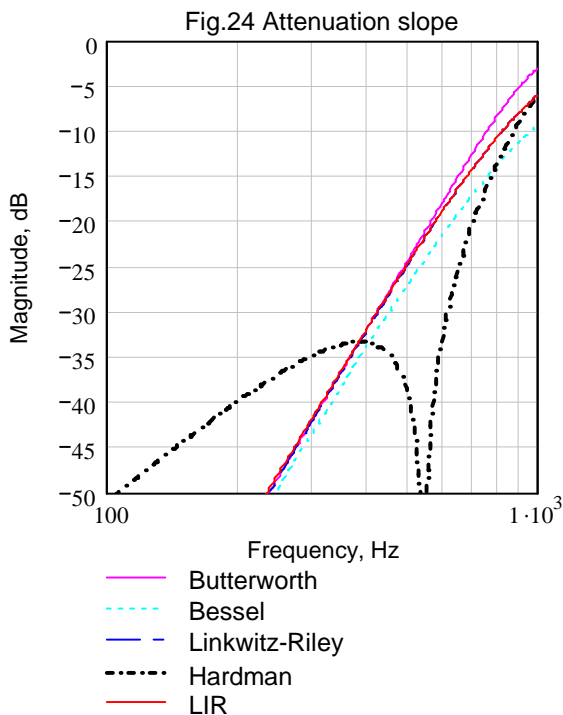
Figure 24 shows how rapidly the high-pass filter of each crossover alignment attenuates as the signal frequency falls below the crossover frequency. On the face of it, it would appear that The Linkwitz-Riley and LIR alignments offers a superior initial roll-off than the Butterworth, and that the Bessel alignment offers generally better attenuation characteristics than Linkwitz-Riley and Butterworth. However, if we apply some equalisation to flatten the summed magnitude response of the Butterworth and Bessel alignments (the others being largely flat anyway) we see in Figure 25 that Butterworth, Bessel, Linkwitz-Riley and LIR become identical in their attenuation characteristics. We also see clearly the superior initial attenuation slope offered by the Hardman alignment, which becomes less effective further away from the crossover point, but only where the degree of attenuation has reduced the signal to insignificant levels.

We also show the attenuation slope of an 8th order Linkwitz-Riley alignment to illustrate that the 4th order Hardman alignment offers almost the same attenuation performance.

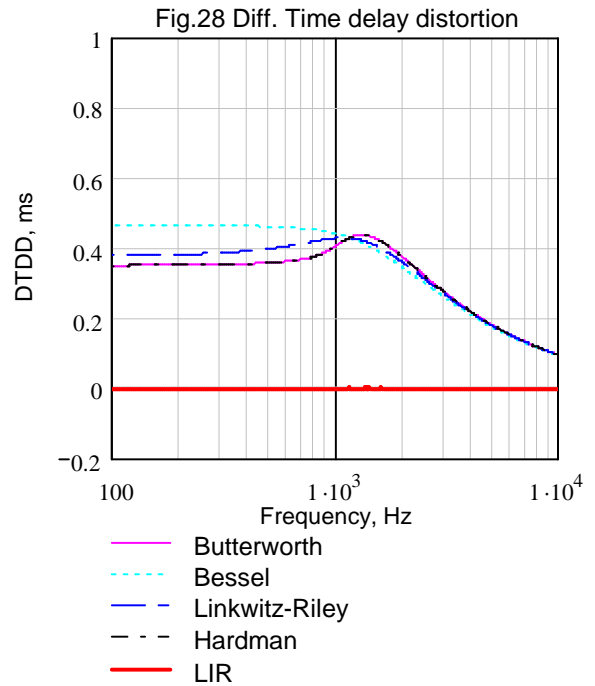
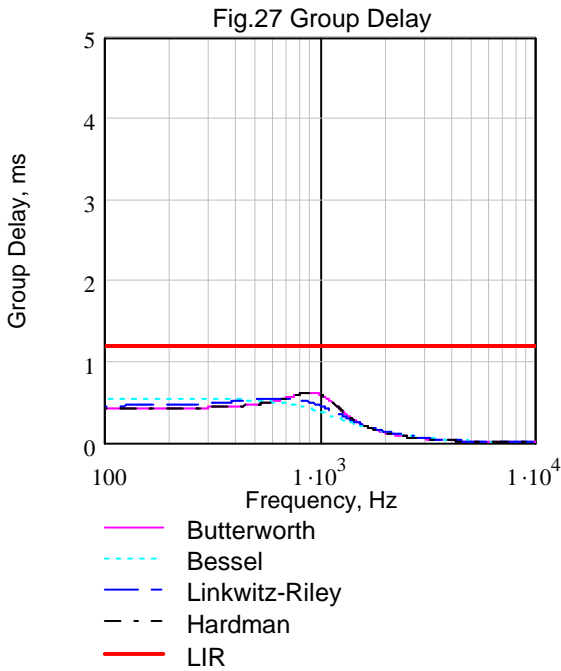


Phase Response

The phase response itself is not particularly revealing, as shown in Figure 26 for the various alignments. However, we can see that only the LIR alignment exhibits linear phase (shown here with the latency delay removed).



The problem is that much of the phase shift is due to linear phase shifting, which equates to pure delay, which we know is not audible. Figure 27 shows the Group Delay of the various alignments. This is a little more revealing, but does it truly give us an indication of the audibility of phase shifting artefacts?



Instead, we will compare the phase anomalies of the various alignments using Differential Time Delay Distortion (DTDD).

Time delay distortion [6] is defined as the difference between the phase delay and the group delay, and is an indication of how much phase shift the signal suffers beyond that of a linear phase shift, i.e. beyond that of a perfect delay constant at all frequencies. We use it here to try to bring some quantitative measure to the likely audibility of non-linear phase shifting. An ideal DTDD would be a flat line, which is what a linear phase process would produce. The more severe the kinks in the DTDD plot, the more audible the phase shift is likely to be. As regards the threshold of audibility, this remains in debate. Studies [7] and [8] have generally found the audibility threshold to be at least 1ms over a reasonable range of frequencies.

Figure 28. shows the DTDD for the various alignments at 1kHz crossover frequency.

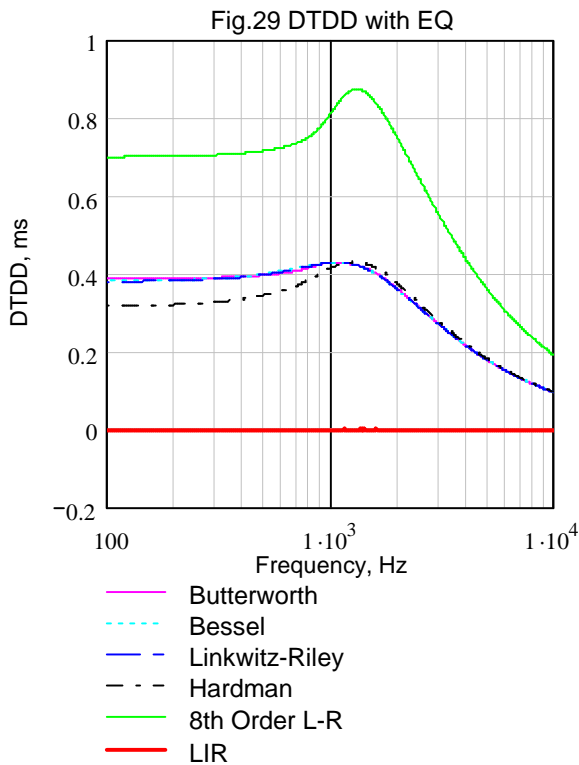
It would appear that the Butterworth and the Hardman alignments have the least flat DTDD, while the Bessel and LIR alignments have the flattest DTDD. We find however that if we apply minimum-phase equalisation to flatten the magnitude responses of the Butterworth, Bessel and Hardman alignments, as shown in Figure 29, the resulting DTDD responses change things somewhat. We now find that all alignments are almost identical (except LIR which remains entirely flat). If we were to use linear phase equalisation, then the results would remain as in 28. However, any 'system' equalisation applied to flatten the overall response is likely to be conventional minimum phase, so the results presented in Figure 29 are more likely.

Of importance here is that although we are showing that the DTDD does not in the end vary too much with the *type* of alignment, the *Order* does make a significant difference. We show the DTDD of the 8th Order Linkwitz-Riley alignment, which is well on its way towards the threshold of audibility, which one could easily imagine might be audible under some circumstances. Indeed, there is anecdotal evidence that the 8th Order Linkwitz-Riley alignment is avoided because of its 'hard' sound. Perhaps this is due to the higher DTDD?

Critically, it must be noted that the Group Delay and DTDD plots shown here are for a 1kHz crossover point. These figures are inversely proportional to the crossover frequency, so the 4th order Linkwitz-Riley alignment for example will exceed the 1ms audibility threshold at crossover frequencies below about 430Hz. However, studies such as [9] show that this threshold of phase audibility increases below about 800Hz, making phase linearity less important at very low crossover frequencies. This suggests that the audibility of DTDD in a 4th Order Linkwitz-Riley crossover with crossover

frequencies much below perhaps 200Hz might not be significant, so there may be no benefit from the additional latency delay that a Linear Phase crossover would introduce. However, an 8th Order Linkwitz-Riley crossover might have audible phase artefacts at crossover frequencies below 860Hz, and may remain audible even at much lower crossover frequencies.

It is certainly safe to assume that the lower the DTDD, the better, and that being the case, the best way towards that goal is to use LIR or otherwise a low Order alignment as possible.



Conclusions

We have discussed the differences between analogue filter designs (or their IIR digital equivalents) and Linear Phase LIR/FIR filters and concluded that each have their place and can be deployed successfully if sufficient care is applied to their design, and that their relative merits are properly considered for the application. We have seen that until recently, neither fast attenuation slopes nor linear phase filtering were the panacea they seemed to be, and that the traditional crossover types were still sometimes hard to beat. However, recent advances in Linear Phase technology have closed this gap, making the LIR alignment a natural choice since it combines the very desirable Linear Phase characteristic with low latency and ease of use.

Although all the alignments have their own unique advantages (otherwise they would not still be in use today), it has to be said that some are perhaps irrelevant. We have shown that if some kind of system equalisation is applied, many of the differences between the alignments disappear. Whilst it should be said that the Bessel alignment is only of any use whatsoever if it is phase matched, it appears to have no benefit if the magnitude response of the system is flattened since the appealing smooth group delay characteristics of the Bessel alignment disappears. The Butterworth alignment seems to have little to offer, and has a peak in the magnitude response which will usually have to be equalised out somewhere in the system. Linkwitz-Riley remains very attractive in all respects apart from not being linear phase. Where better discrimination between bands is required however, the Hardman alignment offers superior performance without sacrificing phase response, and will be preferred to a higher order Linkwitz-Riley alignment. If a small latency delay can be tolerated, then it is clear that the LIR alignment will be the preferred choice, especially for crossovers in the critical range of a few hundred Hertz.

References

- [1] J.W.Gibbs, "Fourier Series", Nature, vol 59, pp. 606, (1899).
- [2] W.E.Thompson, "Delay Networks Having Maximally Flat Frequency Characteristics", Proc IEEE, part 3, vol. 96, pp.487-490 (1949 Nov.).
- [3] R.Miller, "A Bessel Filter Crossover, and its Relation to Other Types" Audio Eng. Soc. Preprint 4776 (1998 August).
- [4] S.H.Linkwitz, "Active Crossover Networks For Non-Coincident Drivers", J.Audio Eng. Soc., vol. 24, pp 2-8 (1976 Jan./Feb.).
- [5] E.W.Hardman, "Precise Active x-over", Electronics World, (1999, August).
- [6] W.Marshall Leach, JR, "The Differential Time-Delay Distortion and Differential Phase-Shift Distortion as Measures of Phase Linearity" J.Audio Eng. Soc., vol. 37, pp. 709-715 (1989 Sept.).
- [7] S.Flanagan et al, "Discrimination of Group Delay in Clicklike Signals Presented via Headphones and Loudspeakers" J.Audio Eng. Soc., vol.53, pp. 593-611 (2005 July/August).
- [8] H. Moller et al, "On the Audibility of All-Pass Phase in Electroacoustical Transfer Functions" J.Audio Eng. Soc., vol55, pp115-134 (2007 March).
- [9] E. Jensen et al, "On the Audibility of Phase Distortion in Audio Systems" J.Audio Eng. Soc., vol.22, pp. 270 (1974 May).
- [10] P.Williams, "LIR Linear Phase Crossovers", Linea Research whitepaper (2013 May).

(C) Linea Research Ltd 2006-2013

Iss. C