TECH LIBRARY

FIR Filters Explained – Part 1

How they work & how they compare with IIR (standard) filters

Introduction

FIR filters have become increasingly popular in pro-audio in recent years, driven by the availability of more powerful DSPs which are able to process multiple FIR filters in real-time. However, whilst FIR filters are a powerful tool they are also frequently misunderstood.

Back to Basics – The Standard Analogue Crossover

Firstly consider an analogue 24dB/Octave Linkwitz low-pass filter with the following frequency response :

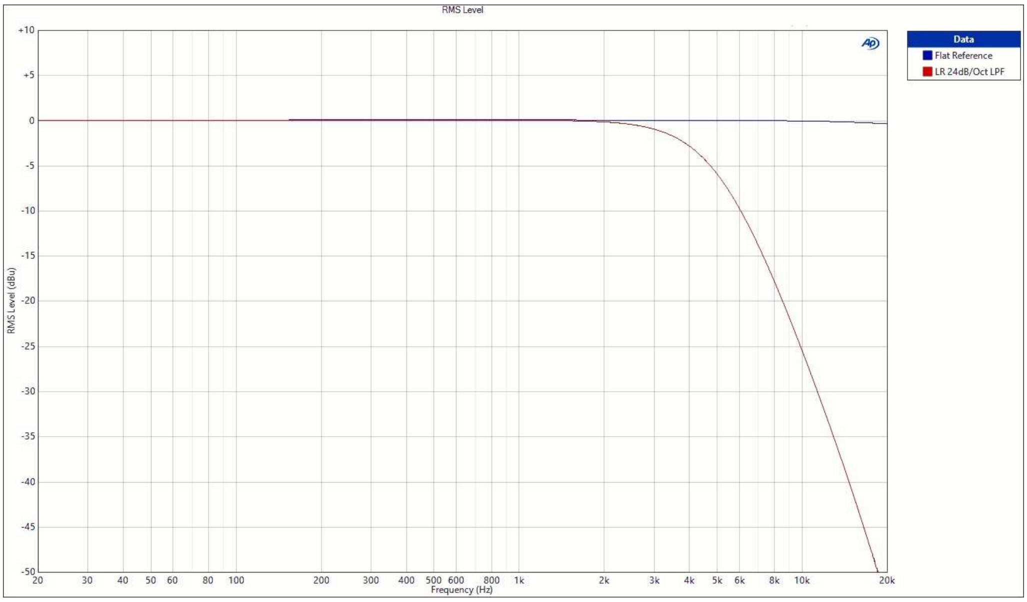

As well as the frequency response of this filter, the other important behaviour in audio systems is the phase response:

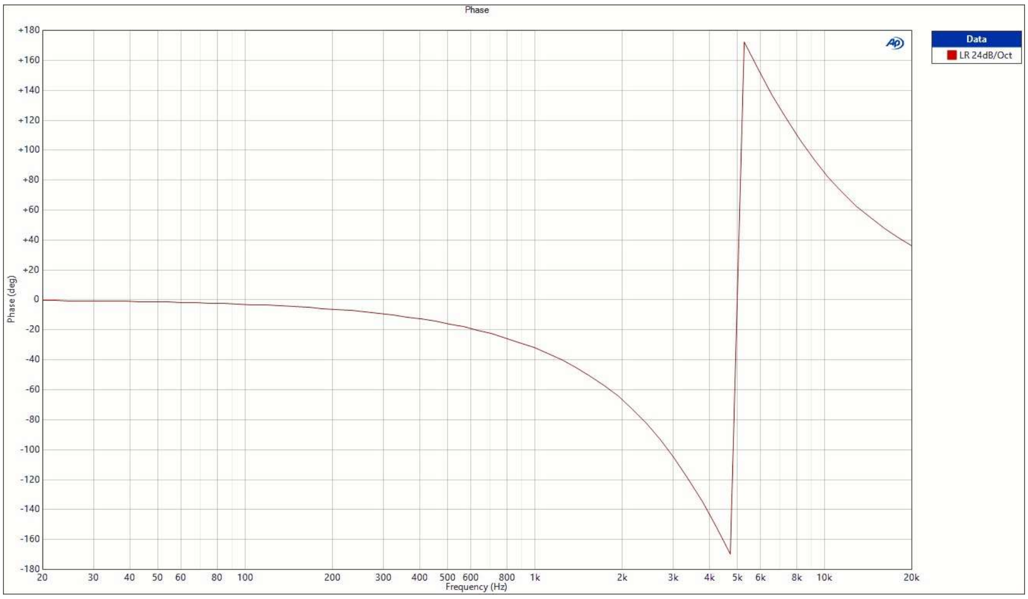

The phase response curve of the filter describes a frequency-dependent delay that the filter adds to the signal. So we can see that as well as altering the response of the system in the frequency domain, the filter also alters the system in the time-domain. Normally this isn't a problem, but it's something that loudspeaker designers spend a great deal of time optimising in multi-way systems using phase plugs and careful positioning to mechanically correct for phase summing and phase cancelling.

These phase shifts can have significant effects on the interaction of loudspeaker drivers in different frequency bands near the crossover points as well as the polar response of the loudspeaker.

The ability to manipulate the phase response of a system electronically is of great benefit to speaker designers as they can now optimise out any phase related anomalies in their multi-way systems, and even improve on existing designs by using new crossover settings with no change to hardware.

Digital Equivalents

In a DSP we can produce digital filters which behave almost identically to their analogue equivalents. The most common way to do this is by using a DSP algorithm known as a "biquad" to design IIR filters. We won't go into detail here about the details of IIR filter design, as the maths gets complex quite quickly. To summarise their operation simply, IIR filters feed part of the previous input and output samples back in to the filter at different gain levels to create the next sample output, creating a feedback cycle which will theoretically go on forever. This continuous feedback system is why this type of digital filter is called Infinite Impulse Response (IIR).

IIR Advantages

IIR filters have been the standard filter method in digital audio for over 20 years and have some distinct advantages :

– Almost identical behaviour to the analogue equivalents (which eased the transition from analogue to digital)
– Computationally efficient using "biquads" (important in the early days of DSP when there was little processing power available)

IIR Limitations

Despite their widespread use, the design of IIR filters is quite mathematically complex. However, in pro-audio DSP systems this complexity is hidden from the user – just enter the standard parameters to define the filter shape and frequency and the software calculates the coefficients needed to ensure the filter behaves as expected and stay stable. Despite their advantages and widespread use, IIR filters have some limitations :

– Design of asymmetric filters, or filters with a complex frequency response is difficult
– Difficult to control phase response of filters

Can FIR Filters Overcome These Limitations?

Yes they can – but they come with some important considerations of their own, which we explore in part 2…