

FIR vs IIR Crossovers in Active Loudspeaker Design

Revision 1

Contents

Introduction	3
IIR filtering	3
FIR filtering	4
Phase behaviour and time-domain response	4
Combining IIR and FIR	6
Crossover implementation	7
Latency considerations	7
Computational requirements	8
Practical workflow differences	8
Application considerations	8
Summary	8
Further reading	9

This note explains the practical differences between Infinite Impulse Response (IIR) and Finite Impulse Response (FIR) filters when used for crossovers and correction in active loudspeaker systems, and when each approach is appropriate. Both are widely used and well understood; the right choice depends on the requirements of the system rather than any single “best” method.

Introduction

Digital signal processing plays a central role in modern active loudspeaker systems. At its simplest, DSP enables crossover implementation and frequency response shaping. As system requirements increase, attention extends beyond magnitude response to include phase behaviour, time alignment, and overall system integration.

IIR and FIR filters are the two primary filter types used in this context. They offer different capabilities and trade-offs when applied to crossover design and system correction, which the following sections set out in turn.




Figure 1: Signal path comparison: sequential IIR stages versus a unified FIR processing block versus a hybrid IIR-plus-FIR approach

IIR filtering

IIR filters are the most commonly used approach for crossover implementation in both analogue and digital systems. Their typical characteristics are that they are computationally efficient, introduce very low latency, use well-established filter topologies such as Butterworth and Linkwitz–Riley, and are straightforward to implement.

In an active loudspeaker system, IIR filters are typically used to define crossover slopes, shape individual driver response, and implement parametric equalisation. Time alignment is usually achieved by applying delay to individual drivers, allowing the acoustic centres to be aligned at the listening position. This approach is widely used and can produce excellent results when properly implemented.

The graph below shows a typical Linkwitz–Riley bandpass, here the passband for a midrange driver in a three-way system, formed by a high-pass at 330 Hz and a low-pass at 3300 Hz, both 24 dB/octave. The solid trace is magnitude; the dashed trace is phase. The characteristic phase rotation through each crossover region is clearly visible.




Figure 2: Linkwitz-Riley 24 dB/octave bandpass crossover at 330 Hz and 3300 Hz, showing magnitude and the associated IIR phase rotation

FIR filtering

FIR filters differ fundamentally in that they allow independent control over both magnitude and phase response. Their key characteristics are that linear-phase or arbitrary-phase responses are possible, the impulse response can be precisely specified, computational requirements are higher, and latency is proportional to filter length.

In a loudspeaker system, FIR filters can be used to implement crossover functions, correct driver magnitude and phase response, and align acoustic output without relying solely on delay. Because a FIR filter can incorporate multiple functions within a single structure, crossover behaviour and correction can be treated as a unified problem rather than a sequence of independent steps.

The graph below shows a FIR filter generated to linearise the phase of the IIR bandpass above. The magnitude response is left unchanged (it is a phase-only correction) while the phase is adjusted by the equal and opposite amount to the IIR crossover. Applied together, the two produce the same bandpass magnitude with a flat phase response.




Figure 3: FIR phase correction filter: flat magnitude response with phase adjusted to cancel the IIR crossover phase rotation

Phase behaviour and time-domain response

The most significant difference between the two approaches lies in phase behaviour. IIR crossover filters introduce frequency-dependent phase rotation. In many systems this is perfectly acceptable and can be managed using established crossover topologies and delay alignment.

FIR filters can be designed to exhibit linear phase, that is, constant group delay, so that every frequency is delayed by the same amount of time. In practice this can give more consistent phase relationships through crossover regions, a more symmetrical impulse response, and improved time-domain coherence.

Linear-phase processing does carry one inherent trade-off. A linear-phase filter is symmetrical about its centre, so realising it requires delaying the output; the symmetric leading half of the response then appears as pre-ringing, low-level energy ahead of the main arrival. The two impulse responses below show this. The IIR bandpass (left) is causal, with a sharp leading edge and an asymmetric decaying tail. The same bandpass with FIR phase correction (right) is symmetrical around its peak, with pre-

ringing visible ahead of the main arrival. Whether this is audible depends on the system, the crossover frequencies, and the overall implementation.




Figure 4: Impulse response of the IIR bandpass crossover: causal, asymmetric, with a decaying tail




Figure 5: Impulse response of the IIR bandpass with FIR phase correction: symmetrical around the peak, with pre-ringing

Left: IIR bandpass impulse response. Right: the same response with FIR phase correction applied. Note the symmetry, and the pre-ringing ahead of the main arrival.

The step response shows the same behaviour from a complementary perspective. Because a bandpass passes no energy at DC, both responses settle back to zero rather than holding a level. The IIR step (left) rises sharply and rings asymmetrically as it decays; the FIR-corrected step (right) is symmetrical about its centre, with the characteristic pre- and post-arrival ringing of a linear-phase response.




Figure 6: Step response of the IIR bandpass crossover: sharp asymmetric rise and decaying ringing




Figure 7: Step response of the IIR bandpass with FIR phase correction: symmetrical about the centre with pre- and post-arrival ringing

Left: IIR bandpass step response. Right: the same response with FIR phase correction applied. Both return to zero, as a bandpass passes no energy at DC.

Combining IIR and FIR

The two approaches are not mutually exclusive. A practical and increasingly common method is to use IIR filters for the crossover slopes and equalisation, where they are efficient and well understood, and then apply a FIR filter purely to linearise the resulting phase. This is exactly the pairing illustrated above: the IIR bandpass defines the magnitude response, and the FIR filter corrects the phase without altering it.

This hybrid approach keeps the computational cost of the magnitude shaping low, while reserving FIR resources for the task where they add the most value. It also sidesteps a fundamental limitation of FIR filters: frequency resolution is tied to filter length, so controlling low frequencies with FIR alone requires very long filters, and correspondingly high latency. Using IIR for the low-frequency crossover slopes and FIR for phase correction avoids spending a large tap budget simply to reach down in frequency. The DSP-8C supports both filter types on every channel, so this combination can be implemented directly: IIR crossovers and EQ designed in DSPconfig, with a phase-correcting FIR filter (for example generated in rePhase) loaded alongside them.

Crossover implementation

In a conventional IIR system, crossover slopes are defined by the filter topology, drivers are equalised individually, and delay is used to align acoustic centres. This results in a system where magnitude response can be well controlled, but phase behaviour is determined by the chosen filters and alignment strategy.

In a FIR-based system, crossover behaviour, equalisation, and phase correction can be combined; filters can be designed to achieve a specific target response; and alignment can be integrated into the filter itself. This allows the system to be optimised as a whole, rather than as a set of independently corrected components.

The illustration below contrasts the two alignment strategies. On the left, IIR processing uses a bulk delay on the forward drivers to align arrival times at the listening axis; each driver still carries its own crossover phase rotation. On the right, a FIR filter per driver corrects the time offset and linearises phase in a single integrated step.




Figure 8: Driver alignment comparison: IIR bulk delay alignment versus FIR integrated time and phase correction in a three-way speaker

Latency considerations

Latency is an inherent consideration in FIR-based processing. IIR filters introduce negligible latency, only the small group delay of the filter itself. Linear-phase FIR filters introduce a fixed latency proportional to the number of taps: short filters introduce modest latency, while longer filters used for higher resolution or lower-frequency control increase it. (Minimum-phase FIR designs avoid this fixed delay, but give up the linear-phase property that is usually the reason for using FIR in the first place.)

For most playback applications this latency is not perceptually significant. It may be relevant in contexts requiring synchronisation with video, or real-time monitoring. The graph below shows the relationship between filter length and latency at 96 kHz and 192 kHz. Note that operating at the higher sample rate halves the latency for a given tap count, though it also roughly halves the number of taps available.




Figure 9: FIR filter length versus latency at 96 kHz and 192 kHz, with reference thresholds for live monitoring and video synchronisation

Computational requirements

IIR filters are computationally efficient and can be implemented with minimal processing resources. FIR filters require significantly more processing power, particularly as filter length increases. This has historically limited their use in real-time systems, although modern DSP platforms have expanded what is practical. The available processing capacity determines the maximum filter length, the number of channels that can be processed, and the sample rate at which filters can operate.

Practical workflow differences

An IIR-based workflow generally involves selecting a crossover topology, applying parametric equalisation, adjusting delay for alignment, and iterating based on measurements. It is typically fast to implement and well supported by integrated DSP platforms.

A FIR-based workflow generally involves measuring the driver and system response, defining the target magnitude and phase behaviour, designing filters (often in a dedicated tool), and implementing them within the DSP system. It provides greater flexibility but requires more detailed measurement and design work.

Application considerations

In practice, both approaches are valid and widely used. IIR-based systems are well suited to applications requiring low latency, rapid deployment and iteration, and systems where conventional crossover behaviour is sufficient. FIR-based systems are often used where control over phase response is a priority, where crossover and correction are treated as a unified problem, and where sufficient processing resources are available. The appropriate choice depends on system requirements rather than a single “best” method.

Summary

IIR and FIR filters represent two complementary approaches to active crossover design. IIR filters provide efficient, low-latency solutions that are widely used and well understood. FIR filters offer additional flexibility, particularly in controlling phase and time-domain behaviour, at the cost of increased computational demand and design complexity, along with the inherent pre-ringing of linear-phase processing.

Modern DSP platforms increasingly support both approaches, and the two can be combined to good effect: using IIR for efficient magnitude shaping and FIR for phase linearisation. This allows system designers to select, or blend, the most appropriate method for a given application.

Further reading

Getting Started with DSPconfig - configuring crossovers, EQ and FIR filters on the DSP-8C.

How does the DSP-8C compare to the miniDSP Flex Eight?