Understanding Filter Specification Sheets

Thomas P. Gerber
Douglas R. Firth
Alan R. Szary

Precision Filters, Inc.

Ithaca, New York
(607) 277-3550
E-mail: sales@pfinc.com

PRECISION FILTERS, INC.

Transducer Conditioning Systems ■ Filter/Amplifier Systems ■ Signal Switching Systems

1. Introduction

Analog filters are a key component of any signal conditioning system. Precision Filters (PFI) offers a complete line of programmable filters to meet the exacting demands of our customers. When compared to conventional industry types, our proprietary FLAT and PULSE designs provide superior response characteristics that can be adjusted for a given set of measurement criteria. These features are detailed in our standard filter specification sheets, which are intended to guide customers to the right design for their intended application. To facilitate the use of this information for product evaluation and setup, we present here an introduction to the mathematical models used in filter analysis, with a particular focus on how important measures are derived and displayed. Key terms introduced in the text are italicized where they first appear and are listed, along with definitions, in the appendix (Section A.2). We address all the specifications listed in our product documentation, and the examples provided are based on real PFI filters.

2. Filter Definitions

An analog filter is a frequency selective circuit designed to modify the spectral content of electrical waveforms. Filter design is guided by a mathematical *transfer function* that relates filter output to input in terms of frequency. The form of the transfer function determines the topology of the circuit and the properties of its elements, which include both passive and active components.

The transfer function represents a continuous-time linear system that operates on an input to produce a response. The input-response relationship can be analyzed in either the time or frequency domain. In the time domain, the relationship is defined by convolution with the filter's *impulse response*, which is just the response of the filter to a unit impulse. In the frequency domain, the response is defined by multiplication with the transfer function, which is a rational function of the complex Laplace[1] variable s. These relationships are depicted schematically in Figure 1. In the s domain, a filter's *frequency response* relates the frequency of a periodic input signal to both the amplitude and phase of the output. In the time domain, a filter's *transient response* relates the shape of a reference input – usually an impulse or its integral, a step function – to the output waveform. Though attention is often focused on the frequency response, our filter specifications include information on both.

[1] The Laplace transform can be thought of as a generalization of the Fourier transform for causal functions (i.e., $x(t) = 0$ for $t < 0$), and provides a similar mapping from the time to frequency domain. The imaginary part of the complex Laplace variable s is the frequency: $s = \sigma + j\omega$.

Phone: 607-277-3550
E-mail: sales@pfinc.com

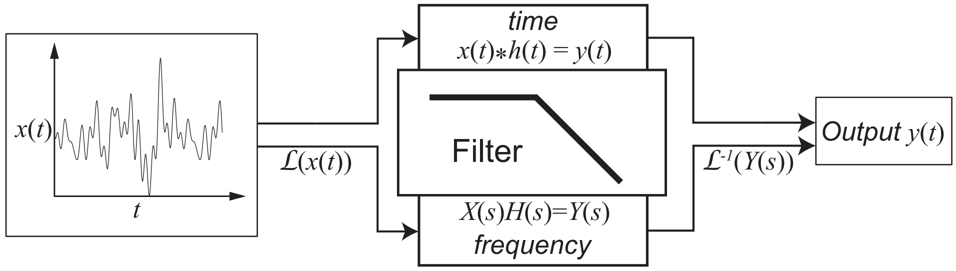

Figure 1. Graphical representation of filtering as an input-output system. Filtering can be modeled in the time domain via convolution and in the frequency domain by multiplication. Here $h(t)$ represents the filter's impulse response, and $H(s)$ represents the filter's transfer function.

2.1 Transfer Function

All analog filters are defined by a transfer function of the complex Laplace variable s with the form

$$H(s) = \frac{N(s)}{D(s)} \tag{1}$$

where N and D are polynomials. The *order* of a filter is just the order of the polynomial D. Reference is often made to a filter's *poles* and *zeros*. The poles and zeros are simply the roots of D (i.e. $D(s) = 0$) and N (i.e. $N(s) = 0$), respectively; hence the number of poles is equal to the order of the filter. For example, PFI's LP4F filter is defined by a 4th order transfer function with 4 poles. (The filter also has 4 zeros, but not all designs have equal numbers of poles and zeros.) Higher order filters offer some desirable features, but they come at a cost: the amount of hardware required to implement a filter in a circuit scales proportionally with its order.

2.2 Frequency Response

The argument of the transfer function is the complex Laplace domain coordinate s. Of importance here is that only the imaginary part of s depends on frequency. The frequency response of the filter is determined by evaluating Equation (1) along this frequency axis. The response can be written in terms of the poles and zeros defined above:

$$H(j\omega) = K\frac{(j\omega - z_1)(j\omega - z_2)\ldots(j\omega - z_m)}{(j\omega - p_1)(j\omega - p_2)\ldots(j\omega - p_n)} \tag{2}$$

Note that the poles and zeros of the n^{th} order filter appear in the factors of the transfer function, and ω is angular frequency ($= 2\pi f$). As both (1) and (2) define the ratio of output to input for the filter as a function of frequency, the coefficient K is called a *gain factor*. The frequency response $H(j\omega)$ is complex-valued: it defines both the magnitude and phase of the response as a function of frequency.

The complete frequency response therefore consists of an *amplitude* and *phase response*.

Amplitude Response

As a ratio, the amplitude response is commonly expressed on the decibel scale:

$$H_{dB} = 20\log_{10}\left(|H|\right) \tag{3}$$

If a filter has an amplitude response of 1 for a given frequency (i.e. output amplitude = input amplitude), the response in decibels is 0: the filter is said to pass energy at that frequency. Correspondingly, negative values imply *attenuation* (output amplitude < input amplitude) of a frequency component while positive values imply amplification (output amplitude > input amplitude). By convention, a filter's *gain* is reported as H_{dB} and its attenuation as $-H_{dB}$. In practice, analog filters are designed to pass frequencies over one portion of the input spectrum while attenuating the remaining portion (i.e. gain ≤ 0). The relative position of this *passband* with respect to the *stopband* defines the four main *filter classes*: *Low-pass*, *high-pass*, *band-pass*, and *band-stop*[2].

A filter's amplitude response is most easily seen by plotting Equation (3) over a range of frequencies. From a design perspective, it's most useful to consider the *normalized response*. The normalized amplitude response is a scaled version of Equation (2): the frequencies are normalized by the *cutoff frequency*, F_c, defined as the frequency at which a given attenuation is reached (conventionally H_{dB} = 3.01, or $H_{dB} = 1/\sqrt{2}$). Working with normalized frequencies facilitates the comparison of different filter types and their response characteristics (both amplitude and phase). Once a filter is chosen for implementation, the actual frequency response can be obtained by rescaling the normalized frequencies by the desired cutoff frequency.

The amplitude response of PFI's LP4F low-pass filter is shown in Figure 2A. Overlain on the response are a number of commonly used descriptors. The response is segmented into three domains by two reference frequencies: the cutoff frequency F_c (= 1 after normalization) and the stopband frequency (= 5.95F_c for LP4F). The passband (attenuation ≤ 3.01 dB; Figure 2B) and stopband (attenuation ≥ 80 dB for LP4F) are separated by the *transition band*, whose width along the frequency axis depends on the filter's *roll-off slope*: filters with higher roll-off slopes are said to be sharper. An alternative characterization for the sharpness of a filter is the *shape factor*, which is simply the ratio of the stopband frequency to cutoff frequency (= 5.95 for LP4F). (Expressed this way, a lower shape factor

2 The gain factor K determines the total transfer function gain. We report this with respect to the filter's passband: *DC gain* for our low-pass filters and *high-frequency gain* for our high-pass filters (both 0 dB for our designs).

Phone: 607-277-3550
E-mail: sales@pfinc.com

indicates a sharper low-pass filter.) The term *ripple* is used to describe deviations in the transfer function's amplitude response over a given frequency band and is given as an amplitude range. The LP4F filter has only stopband ripple – its passband response is flat – but some filter types also have passband ripple.

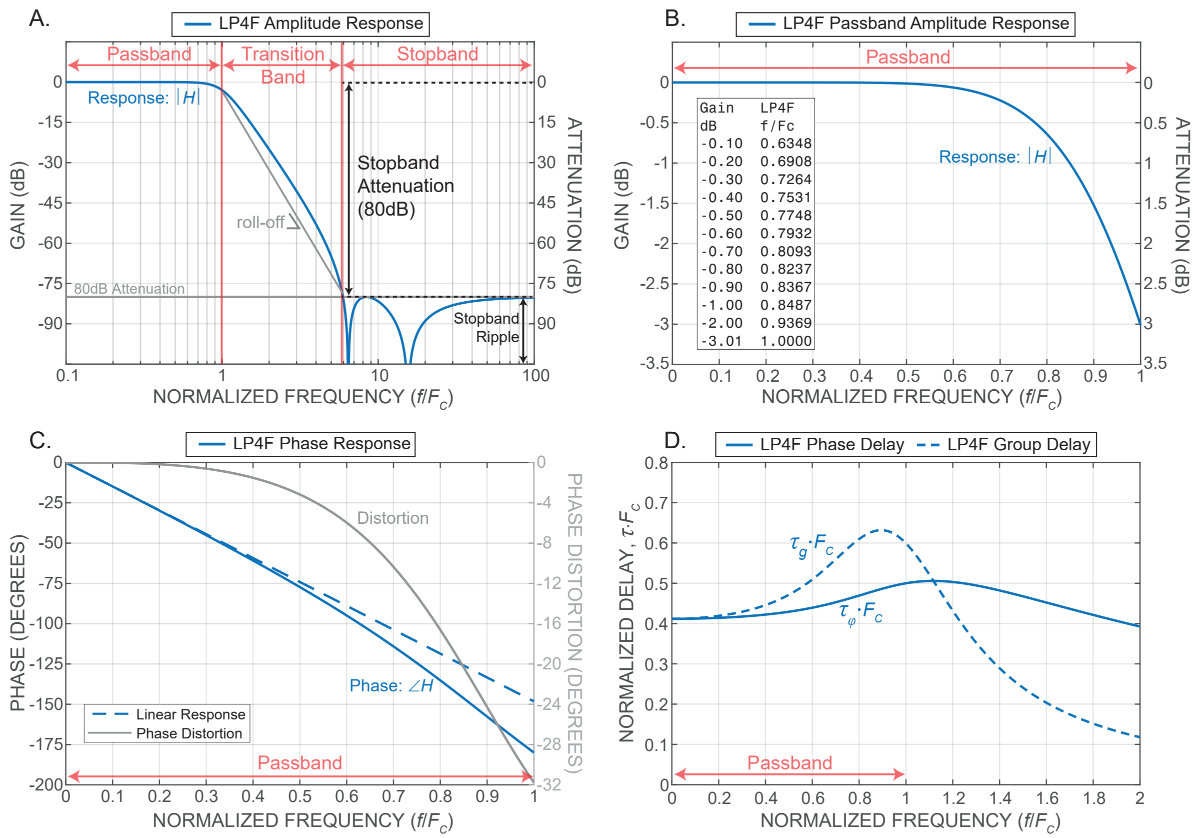

Figure 2. Summary of the low-pass LP4F filter's frequency response: (A) Full amplitude response, (B) Passband amplitude response, (C) Phase response, and (D) Phase and group delays. The passband amplitude response in (B) includes tabulated data, as is standard on PFI specification sheets. Note that the gain and attenuation in (A) and (B) are equivalent except for a sign change.

Increasing the filter order changes the amplitude response, as can be seen between Figure 2A-B and Figure 3A-B. The latter presents the amplitude response of PFI's LP8F filter, which is an 8-pole version of the same *filter prototype* (PFI's FLAT Mode) used for the LP4F. The higher-order LP8F provides a flatter passband and sharper roll-off than the LP4F: the passband attenuation stays below 0.5 dB out to > 90% of the cutoff frequency, and the shape factor is 1.75.

The amplitude response defines how the filter will modify the amplitude of any spectral component passed through it. However, the filter will also affect the phase of those components and, consequently, the shape of the output waveform.

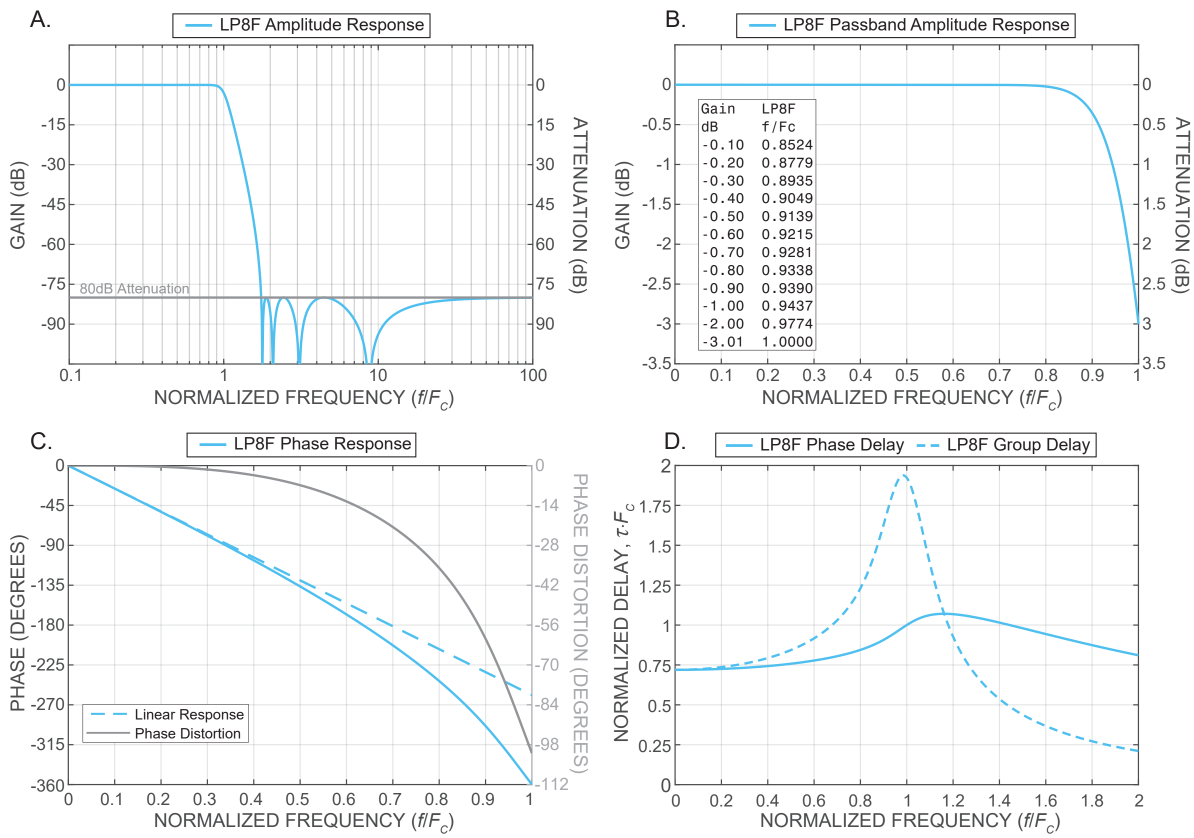

Figure 3. Summary of the low-pass LP8F filter's frequency response: (A) Full amplitude response, (B) Passband amplitude response, (C) Phase response, and (D) Phase and group delays. The passband amplitude response in (B) includes tabulated data, as is standard on PFI specification sheets. Note that the gain and attenuation in (A) and (B) are equivalent except for a sign change.

Phase Response

As noted above, the amplitude and phase responses are obtained from the total frequency response as $|H(j\omega)|$ and $\angle H(j\omega)$, respectively. The phase angles can also be plotted as a function of frequency, again normalizing by F_c. The LP4F phase response is shown in Figure 2C, with the angles plotted in degrees. Note that the phase response is plotted over the passband only (DC to F_c), as we are mainly interested in the phase of in-band signals[3].

An important property of the phase response is its departure from linearity. This deviation is termed *phase distortion* because a filter with a nonlinear phase response in the passband will modify the shape of a spectrally rich in-band waveform. The distortion is measured in units of phase (degrees). We include the phase

3 Note that here and elsewhere, signals of interest that should be preserved by the filter are referred to as *in-band* signals. Spectral components that lie outside the band of interest are described as *out-of-band*.

distortion whenever we present the phase response of a low-pass filter (Figure 2C, Figure 3C).

Phase distortion is most easily understood in terms of time delays. From the phase response of the filter, we can define two delays in terms of the phase $\varphi = \angle H(j\omega)$:

$$\tau_\varphi = -\frac{\varphi}{\omega} \tag{4}$$

$$\tau_g = -\frac{d\varphi}{d\omega} \tag{5}$$

The *phase delay*, τ_φ, defines the delay imparted by the filter to a spectral component with frequency ω. A filter with linear phase will impart the same time delay to all frequencies, implying a phase shift proportional to frequency (i.e. $\varphi = 2\pi f \tau_\varphi$). Phase distortion arises when the delay changes with frequency and is quantified using the *group delay*, τ_g. Between (4) and (5), it's clear that a constant group delay implies a constant phase delay (i.e. $\tau_g = \tau_\varphi =$ constant). The group and phase delays for the LP4F and LP8F are given in Figure 2D and Figure 3D, respectively.

The effect of filter order on the phase response can be seen by comparing the LP4F (Figure 2C-D) to the LP8F (Figure 3C-D). The passband phase distortion, phase delay, and group delay are all greater for the 8-pole filter. Comparing the total frequency response of the LP4F to the LP8F (Figure 2 vs. Figure 3) highlights an important point: for some prototypes, increasing the filter order to obtain a flatter passband response and sharper roll-off comes at the expense of phase linearity. In other words, it is generally not possible to simultaneously maximize both. Even for designs that overcome this limitation[4], there is still a tradeoff in cost: as noted earlier, higher order filters require more hardware to implement.

Figure 4 illustrates the relationship between phase delay and distortion using the LP8F as an example. A test signal is constructed from the first two (odd) harmonics of a square wave with frequencies f_1 and f_3, where by definition $f_3 = 3f_1$ (Figure 4A). This test signal is then filtered using the LP8F with a cutoff frequency set so that $f_3 = .75F_c$. From the LP8F frequency response (Figure 3), we see that both components are within the flat portion of the LP8F's passband (i.e. no attenuation), but the f_3 component will lag the f_1 component in the filter's output. The unequal phase delays are evident in Figure 4B. This phase distortion translates to an output waveform that is both delayed in time and modified in shape (Figure 4C). Note that since we are working with normalized frequencies, units of time are scaled by the cutoff frequency F_c.

4 For example, Bessel filters are optimized for phase linearity; increasing the order of a Bessel filter improves the phase linearity with almost no effect on the amplitude response. It is also possible to phase equalize an elliptic filter – a design optimized for passband flatness and sharp roll-off – so that increasing the order simultaneously improves both the amplitude and phase response.

Phone: 607-277-3550
Web Site: www.pfinc.com

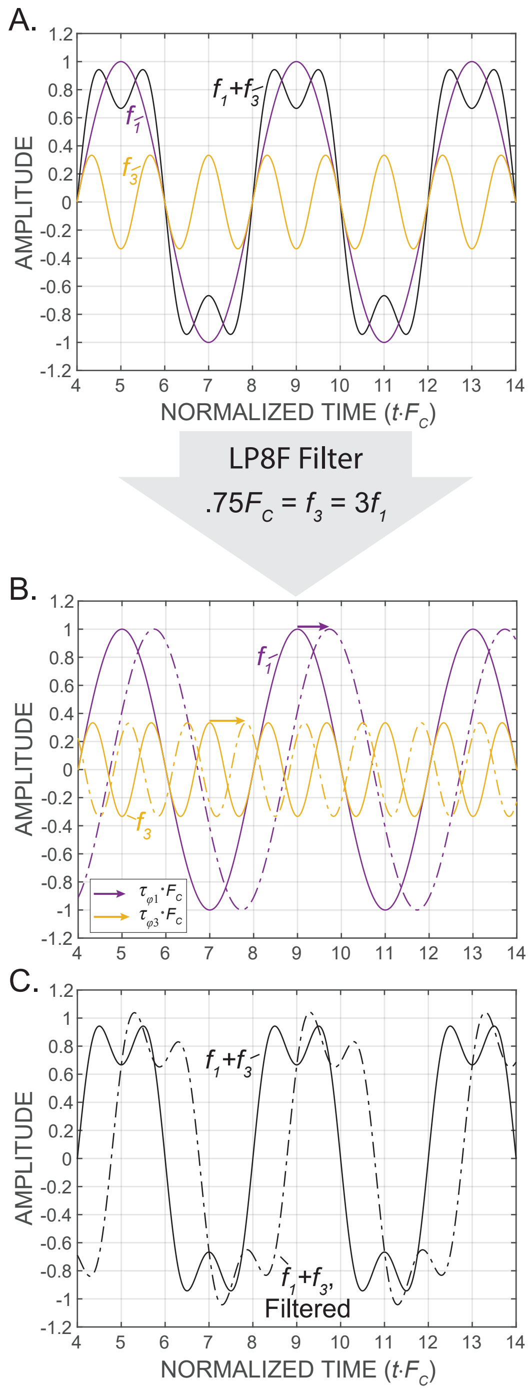

Figure 4. Example illustrating the effect of phase delays on a filtered signal. (A) The test signal is a square wave composed of the first two harmonics, f_1 and f_3, where $f_3 = 3f_1$. (B) After filtering with the LP8F (dashed lines), the two components are delayed by different amounts. (C) As a consequence, the composite signal output from the filter is a delayed and distorted version of the input signal. Note that in all plots, only a steady-state section of the signal is shown; the initial transient response is omitted for clarity.

Phone: 607-277-3550
E-mail: sales@pfinc.com

The previous example showed how the phase delay translates to measurable time shifts in the filter output. The group delay also has physical meaning, which is most easily understood by considering an amplitude modulated (AM) signal. In the appendix (Section A.3), we show how a filter's phase response can separate the phase and group delays of AM inputs, and in doing so shed light on the distinction between the two.

2.3 Transient Response

Though the frequency response guides most decisions regarding filter design and selection, it is also useful to consider the time-domain response to a reference waveform. Conventionally, the reference waveform is a hypothetical, physically unrealistic transient: it is chosen as an extreme case to challenge the filter. The impulse response is defined by passing a unit impulse through the filter and examining the output waveform in time[5]. Similarly, the *step response* is defined by passing a unit step function through the filter (Figure 5A). Here we address the latter, as it is the only transient response we include with our filter specification sheets.

Step Response

The step response of the LP4F filter is summarized in Figure 5B. As with the amplitude response, a number of standard descriptors are used for characterization and comparison. There is an initial delay in the output waveform relative to the unit step. We quantify this delay using two measures: the 50% delay is the time it takes for the output waveform to reach half the amplitude of the unit step; the 10-90% rise is the time it takes for the output waveform to increase from 10% of the unit step amplitude to 90%. The maximum difference (> 0) between the output waveform and the unit step is termed the *overshoot* and is expressed as a percent deviation. Subsequent oscillation about the ideal response is termed *ringing*.

The step response can also be examined by determining the *settling time* (Figure 5C). The settling time is computed in two parts. For the initial rise, it is simply the absolute deviation from the ideal step response. This deviation is then extended by plotting the amplitude of the ringing (relative to the ideal response) as it decays over time. We call this deviation the *settling error* and express it as a percentage. Notice that the settling time plot for LP4F is extended beyond the limits of the step response, where the settling error diminishes to <<1%. The settling time provides an indication of how 'responsive' the filter is to rapidly changing input signals, such as pulses, shocks, or other waveforms with sharply rising or falling edges. To extend the LP4F and LP8F comparison to the time domain, we also

5 Recall that filtering in the time domain is represented as convolution between an input signal and the filter's impulse response (Figure 1): the impulse response defines the convolution kernel.

Phone: 607-277-3550
Web Site: www.pfinc.com

include the LP8F settling time in Figure 5C. The higher-order LP8F is less responsive than the LP4F: it takes approximately half as long for the LP4F to reach settling errors below 10% as the LP8F.

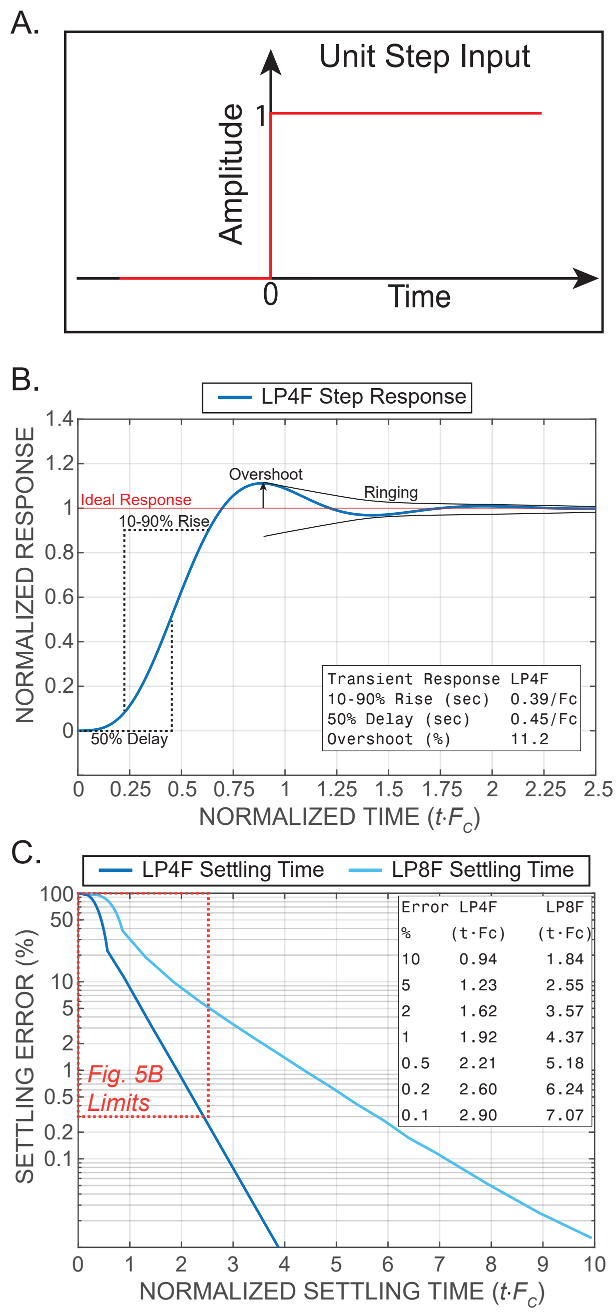

Figure 5. Example of the step response. (A) Sketch depicting the reference input used to define a filter's step response. (B) The step response of the LP4F, with standard measures used in PFI's specification sheets labeled. (C) Settling time representation of the step response for both the LP4F and LP8F. Both (B) and (C) include tabulated data, as is standard on PFI specification sheets.

2.4 Alias Attenuation

Alias attenuation, or *anti-aliasing*, is the primary function of an analog low-pass filter in a signal conditioning system. Aliasing is a consequence of sampling: the term *alias* refers to an apparent frequency component in a sampled signal whose true (analog) frequency is higher. In other words, the spectral component appears as a lower-frequency alias in the sampled signal. According to the sampling theorem, the choice of sample rate F_s defines the *baseband* as the interval from DC ($f = 0$) to $F_s/2$. Once a signal is sampled at this rate, any spectral component f that lies outside the baseband (i.e., $f > F_s/2$) will produce an alias f_a in the baseband according to

$$f_a = \left| f - nF_s \right| \tag{6}$$

where n is an integer ≥ 1. Note that multiple spectral components can produce the same baseband alias and, importantly, cannot be separated from the true baseband signal once it's sampled.

A useful way to think about aliasing is in terms of *spectral folding*. In this view, sampling an analog signal is equivalent to folding the frequency axis at integer multiples of the *folding frequency*, $F_s/2$. This can be visualized by constructing a spectral folding chart (Figure 6A). The chart depicts a horizontal frequency axis that is repeatedly folded over itself: the vertex of each fold corresponds to a multiple of the folding frequency. When plotted this way, vertical lines connect all out-of-band frequencies (i.e. $f > F_s/2$) with their baseband aliases on the bottom of the chart (i.e. $f_a < F_s/2$). Note that for the sake of illustration, the folding chart in Figure 6A is arbitrarily limited to $f/F_s = 3.5$, but can be extended to the highest frequency present in a signal.

The relationships summarized in Figure 6A can be combined with a filter's amplitude response to quantify the attenuation of aliased components for a given sample rate. In Figure 6B, the LP4F's amplitude response is shown over the passband and transition band. As an example, a sample rate $F_s = 6F_c$ is selected, which equates to a folding frequency of $3F_c$. By folding the amplitude response about $F_s/2$, any frequency component above $F_s/2$ can be connected to its baseband alias. To illustrate, the first alias corresponding to the cutoff frequency is labeled on the plot (f_a); it would appear in the sampled baseband down 65 dB from its original amplitude.

In practice, alias attenuation is usually measured in terms of the highest frequency of interest. Again referring to Figure 6B, assume the LP4F is used to filter a signal in which the highest frequency of interest, f_h, is at the cutoff frequency F_c. At this frequency, the *minimum attenuation* of in-band aliases will be 65 dB (i.e. f_a), compared to a maximum attenuation of 3 dB (at f_h) for the in-band signal. For

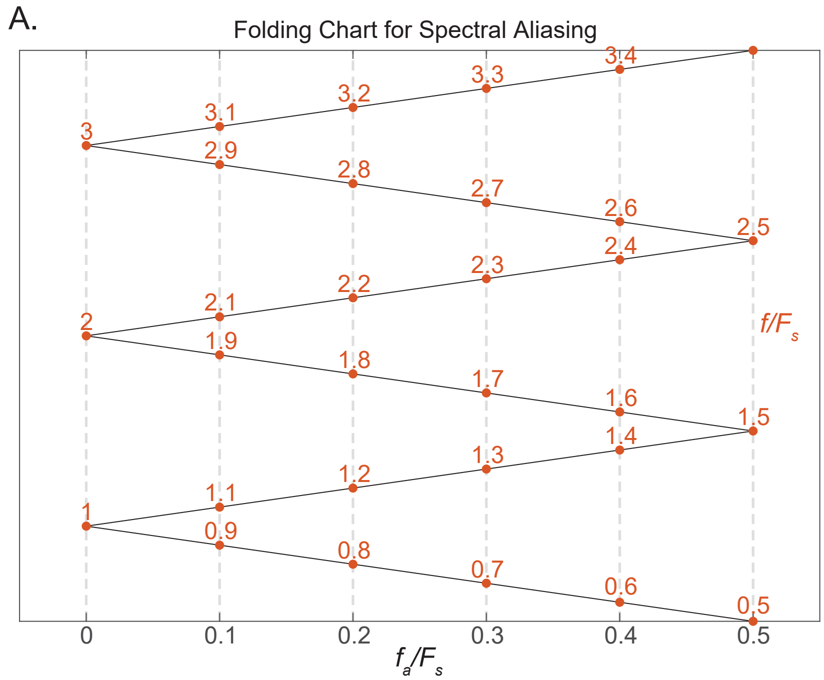

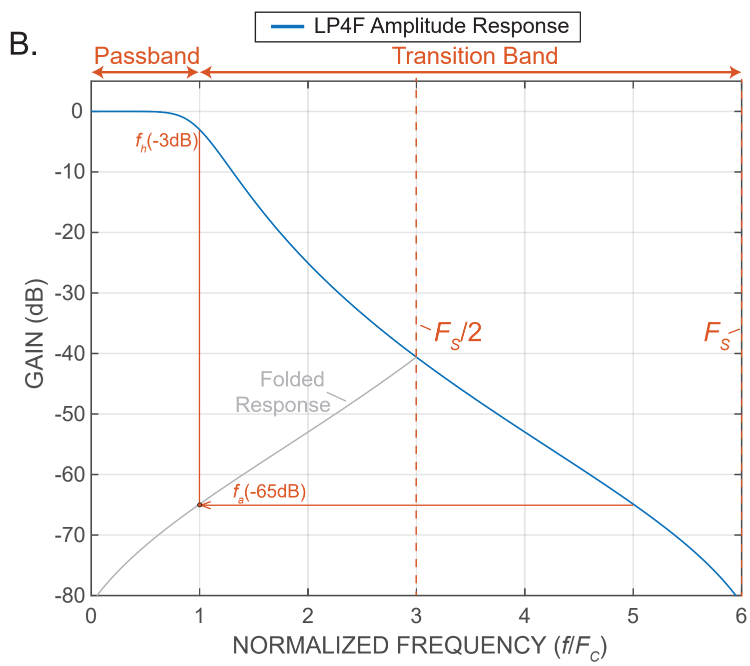

Figure 6. (A) Folding chart depicting the relationship between the folding frequency $F_S/2$ and the alias frequency f_a of out-of-band components. (B) Example showing the effect of spectral folding on the LP4F amplitude response. A sampling frequency of $6F_C$ is chosen, corresponding to a folding frequency of $3F_C$. The first alias, f_a, appearing at the cutoff frequency is shown for illustration.

Phone: 607-277-3550
E-mail: sales@pfinc.com

frequencies of interest less than f_h, the alias attenuation will be greater.

For a given F_s and f_h, the minimum attenuation will depend on the shape of a filter's amplitude response. This dependence is included on all our low-pass filter specification sheets; an example for the LP4F is shown in Figure 7. (The example alias f_a shown in Figure 6B is included to aid comparison.) In this plot, each line represents a different value of f_h (scaled by F_c). For a fixed f_h, the minimum attenuation increases as the sample rate increases relative to f_h. This is equivalent to moving F_s to the right in Figure 6B. When combined with the amplitude and phase response (e.g. Figure 2), the plot can be used to determine an appropriate sample rate and cutoff frequency for the desired level of in-band attenuation, phase linearity, and alias attenuation at the highest frequency of interest.

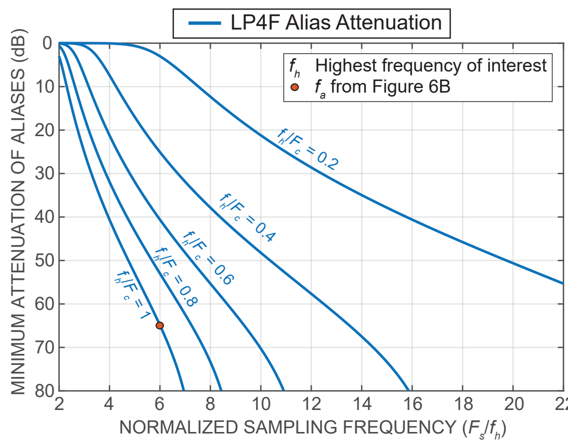

Figure 7. LP4F alias attenuation plot. The point corresponding to f_a in Figure 6B is included for reference. Note that the lower bound of the frequency axis equates to $F_s = 2f_h$, ensuring f_h is within the baseband defined by F_s.

2.5 High-pass and Band-pass Filters

While most of the concepts introduced above apply to all filter classes, our focus to this point has been on low-pass filtering. This is not without reason: low-pass filters are a critical component of most signal conditioning systems. However, our filter documentation also includes specifications for high-pass and band-pass classes, so we conclude with an example of a high-pass filter response, along with a brief overview of band-pass filter design.

High-pass Filters

Figure 8 shows the normalized frequency response of PFI's HP4F filter. The correspondence between the amplitude response of the LP4F (Figure 2A-B) and the HP4F (Figure 8A-B) is apparent: the high-pass response is an inverted version of the low-pass response. This is by design, as both transfer functions are based on the same filter prototype (in this case, PFI's FLAT mode). Consequently, the descriptors introduced in Figure 2 can be applied to the high-pass response. However, unlike the LP4F – whose passband is bounded at DC ($f = 0$) – there is no natural upper bound to the HP4F's passband. (Though there is a practical upper bound that depends on the filter's circuitry.)

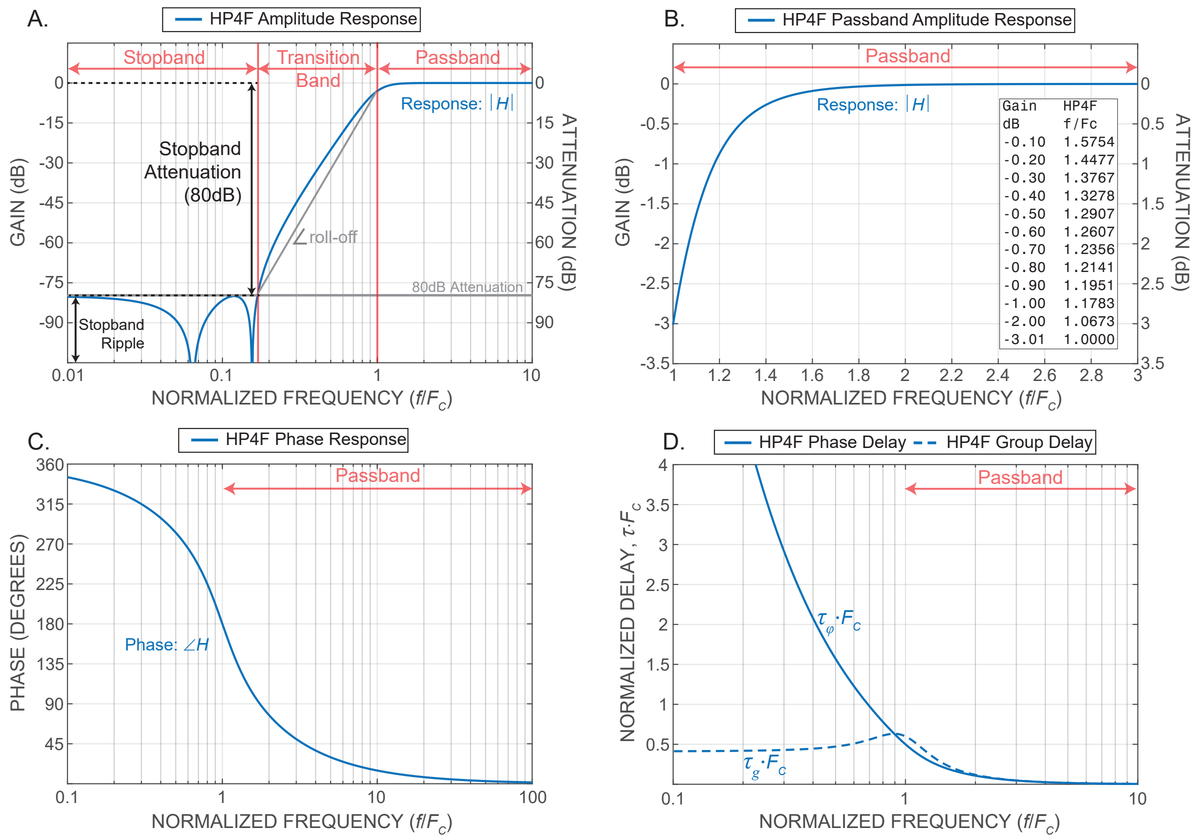

Figure 8. Summary of the high-pass HP4F filter's frequency response (compare with Figure 2): (A) Full amplitude response, (B) Passband amplitude response, (C) Phase response, and (D) Phase and group delays. The passband amplitude response in (B) includes tabulated data, as is standard on PFI specification sheets. Note that the gain and attenuation in (A) and (B) are equivalent except for a sign change.

For this reason, the HP4F's phase response is presented differently (Figure 8C). The HP4F phase is plotted on a logarithmic frequency scale over a wider range than the LP4F. There is no measure of phase distortion because the linear reference response is not well-defined[6]. So while not apparent between Figure 2C and Figure 8C, the two phase responses share the same correspondence as that noted for the amplitude response, with one important difference: the high-pass phase is positive. From Equation (4), this implies a negative phase delay. So for any spec-

6 In the low-pass case, the linear response is defined by the zero-phase slope at DC. In the high-pass case, the zero-phase slope is approached asymptotically as $f >> F_c$ and $\varphi \rightarrow 0$. In the high-frequency limit, the zero-phase slope $\rightarrow 0$, implying that if the high-pass phase distortion is defined in the same way as it is for a low-pass filter, the phase distortion and phase response are one and the same.

tral component passed through the HP4F, the output will lead the input.

The difference between a leading and lagging phase response can be illustrated by passing a pure sinusoidal tone through each filter (Figure 9). In Figure 9A, the phase response of the LP4F and HP4F are plotted together to show their inverse relationship. At the cutoff frequency F_c, the phases are equal but of opposite sign (± 180°). The response of each filter to a sinusoidal input with a frequency equal to F_c is shown in Figure 9B. The responses differ only in the early (transient) portion of the output signal; the steady-state responses are identical, since the difference in phase shift equates to a full period of the cycle (i.e. 360° = 2π radians). Observing the transient response, it's worth noting that a negative phase delay, or lead, does not violate causality – output does not precede input – but is simply a consequence of the frequency-domain description in terms of periodic inputs.

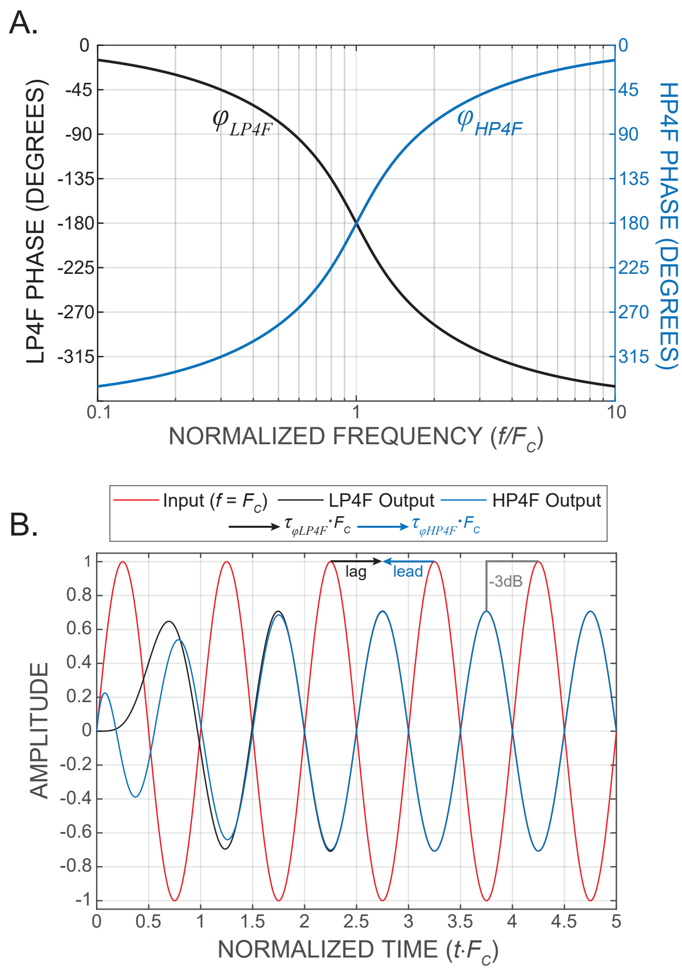

Figure 9. (A) Phase response of the HP4F and LP4F plotted on commensurate scales. (B) Example demonstrating the time-domain response of both filters to a pure sinusoidal tone at the cutoff frequency. The phase delay at the cutoff frequency for each filter is labeled on the plot: the delays are equal but of opposite sign, representing leading and lagging phase, respectively.

The high-pass transient response can be characterized using the unit step as described in Section 2.3. Again, the descriptors used to characterize the low-pass step response (Figure 5) can be applied to the high-pass step response (Figure 10), with one key difference: the ideal low-pass response is full preservation of the unit step, whereas the ideal high-pass response is attenuation of the unit step

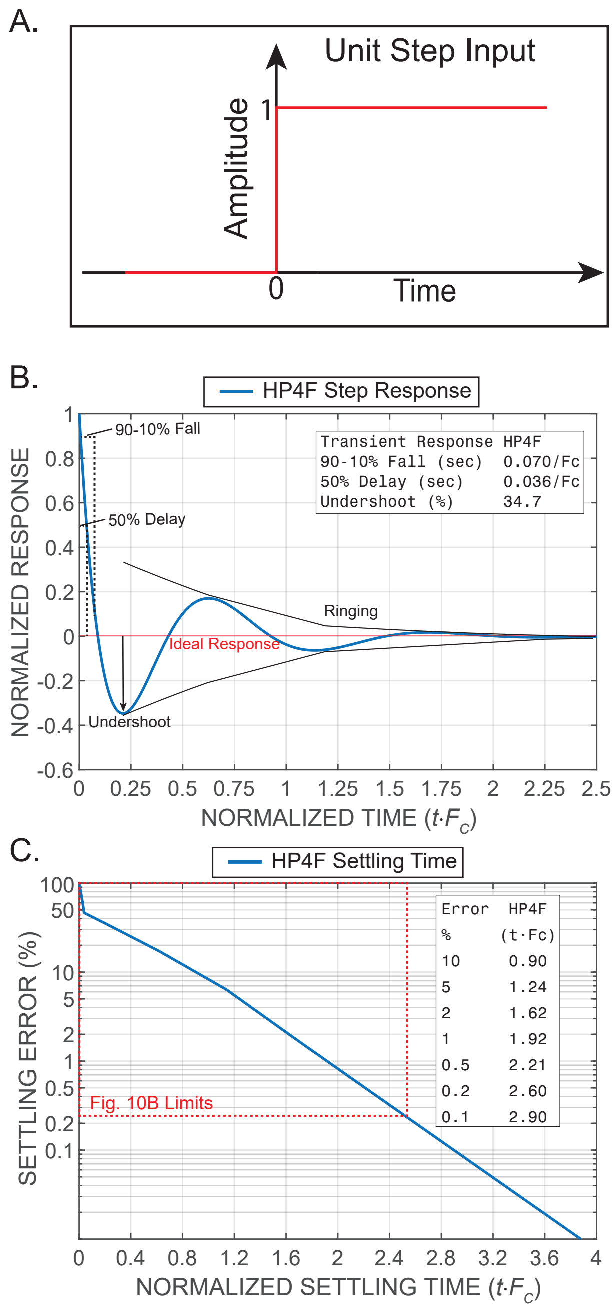

Figure 10. Summary of the HP4F step response (compare with Figure 5). (A) Sketch depicting the reference input used to define a filter's step response. (B) HP4F step response, with the standard measures used in PFI's specification sheets labeled. (C) Settling time representation of the step response. Both (B) and (C) include tabulated data, as is standard on PFI specification sheets.

Phone: 607-277-3550
E-mail: sales@pfinc.com

to the stopband level (80 dB for HP4F). Consequently, the 10-90% rise becomes the 90-10% fall, and the *undershoot* is below the ideal response level.

Band-pass Filters

Band-pass filters are designed to produce a passband bounded by two finite cutoff frequencies. A band-pass amplitude response can be realized by *cascading* a high-pass filter in series with a low-pass filter. In this case, the total amplitude response is the product of the high- and low-pass amplitude responses, and the total phase response is the sum of the high- and low-pass phase responses. The band-pass response can then be characterized in terms of the high-pass and low-pass cutoff frequencies, F_{HP} and F_{LP}. Extending the examples introduced above, consider the combined amplitude response of the HP4F and LP4F (Figure 11A). If both filters are set to the same cutoff frequency, $F_{HP} = F_{LP}$, then their passbands do not overlap: the minimum attenuation is twice the conventional cutoff attenuation (i.e. 6.02 dB, or $|H| = ½$). We call this minimum attenuation the *insertion loss*. A 3 dB bandwidth, B_{3dB}, can be defined relative to the level of insertion loss. To extend the bandwidth, the passband overlap must be increased. We quantify the overlap in terms of the ratio $\alpha = F_{LP} / F_{HP}$, with the design criterion $F_{LP} \geq F_{HP}$. A center frequency F_o corresponding to the insertion loss can also be defined as the geometric mean of the low-pass and high-pass cutoff frequencies,

$$F_o = \sqrt{F_{HP} F_{LP}} \tag{7}$$

In Figure 11B, the response shown in Figure 11A for $\alpha = 1$ is compared with responses for $\alpha = 2$ and $\alpha = 4$, respectively, where the frequency axis is now scaled by the center frequency F_o. Notice that as α increases, the bandwidth B_{3dB} increases and the insertion loss decreases, resulting in a broader passband. A convenient normalized measure of band-pass filter shape is the sharpness Q, which is just the ratio of the center frequency F_o to the 3 dB bandwidth B_{3dB}. A high Q value indicates a sharper band-pass filter with a narrower passband. For design purposes, the relations given in Figure 11B can be used to set the cutoff frequencies for a desired band-pass center frequency, bandwidth, and sharpness.

The band-pass phase response for the same three values of α are shown in Figure 11C. The phase changes sign at the center frequency – the reason for this should be apparent by inspection of Figure 9A – and, consequently, there is a change from leading to lagging phase across the passband. The symmetry of the HP4F-LP4F band-pass response means that if two spectral components with a geometric mean equal to F_o are passed through the filter, they will be attenuated and shifted by the same amount, but one will lead and the other will lag. Definitions of phase distortion for the band-pass response suffer from the same ambiguity noted for the high-pass phase response. Characterization of phase linearity should therefore be done with reference to the in-band limits.

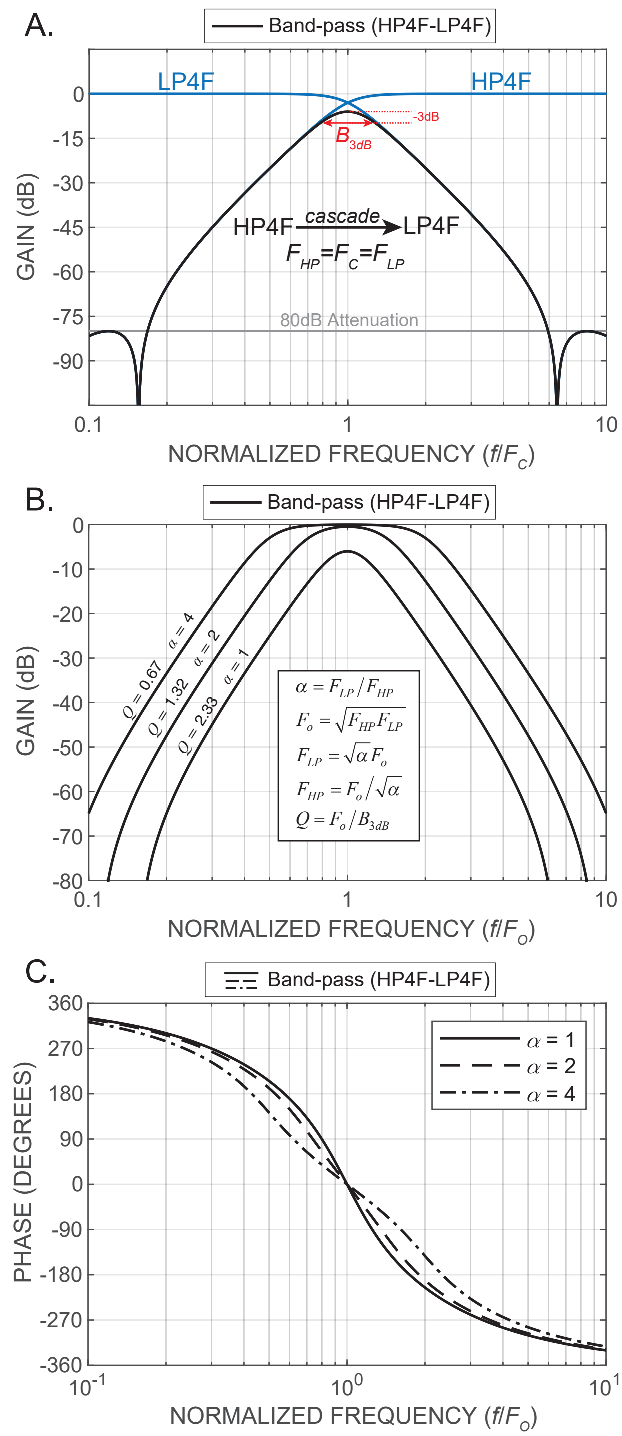

Figure 11. Example of a band-pass amplitude response. (A) Band-pass response generated by cascading HP4F with LP4F using the same cutoff frequency for both filters. (B) Band-pass response shown for different levels of passband overlap ($\alpha = 1, 2, 4$). Note the frequency axis is scaled by the center frequency F_O. (C) Band-pass phase response for the same values of α.

3. A Brief Word on Filter Selection

The preceding section introduced three responses – amplitude, phase, and transient step – that are used to analyze and compare filter characteristics. It's important to remember that the responses are not independent. Changing the order of a filter to achieve a desired amplitude response may cause unacceptable changes in the phase or transient response, and vice versa. As shown for PFI's FLAT mode filter, there is a tradeoff in increasing the order: passband flatness and roll-off sharpness are gained in exchange for phase linearity and time-domain responsiveness. (Not to mention cost of hardware implementation, which increases with filter order.) Similar trade-offs will occur when choosing between different filter prototypes. PFI's two proprietary filter designs are complementary in this regard. The FLAT mode filters introduced above are constructed to maximize performance in the amplitude response, while the PULSE mode filters are optimized for phase linearity and transient response. For this reason, PFI offers programmable filters that can operate in either mode, allowing users to change the response on the fly without modifying hardware.

Choosing the right filter for a particular application requires a clear assessment of measurement criteria. If the priority is amplitude preservation for spectral analysis, then optimizing the amplitude response in exchange for some phase nonlinearity may be acceptable. If waveform reproduction is important, then optimizing for phase linearity in exchange for a less discriminating amplitude response may be acceptable. As an example of the latter, a well-established rule[7] in the measurements community stipulates that for adequate waveform reproduction, in-band attenuation should deviate from a flat response by no more than 5%, and in-band phase distortion should be kept below 5°. With criteria like these in mind and an understanding of the material presented in a standard filter specification sheet, a judicious choice of filter for the application in question should be straightforward.

4. Summary

The content presented here was framed by the organization of our standard filter specification sheets. Our hope is that by structuring the paper this way, we facilitate careful scrutiny of our filter specifications and assist our users with product selection and configuration. The interested reader can find more detailed treatments in numerous textbooks and other curricular materials. Expanded discussions on some topics introduced here (e.g. anti-aliasing) are available in PFI's technical library.

7 Wright, C.P. 1995. Applied Measurement Engineering. Prentice Hall, Edgewood Cliffs, NJ. 402pp.

Phone: 607-277-3550
Web Site: www.pfinc.com

A. Appendix

A.1 Symbols

$N(s)$	Numerator polynomial of the filter transfer function
$D(s)$	Denominator polynomial of the filter transfer function
$H(s)$	Filter transfer function
$\|H(s = j\omega)\|$	Amplitude response
$\angle H(s = j\omega)$	Phase response (φ)
H_{dB}	Amplitude response in decibels
B_{3dB}	3 dB bandwidth of band-pass amplitude response
F_c	Filter cutoff frequency
F_{HP}	High-pass filter cutoff frequency for band-pass cascade
F_{LP}	Low-pass filter cutoff frequency for band-pass cascade
F_o	Band-pass center frequency
F_s	Sampling frequency
$\mathcal{L}$	Laplace transform operator
$\mathcal{L}^{-1}$	Inverse Laplace transform operator
$h(t)$	Impulse response
τ_φ	Phase delay
τ_g	Group delay
f	Frequency in Hz
f_a	Alias frequency
f_h	Highest frequency of interest (in-band upper bound)
s	Laplace domain coordinate
α	Band-pass overlap ratio (F_{LP}/F_{HP})
φ	Phase angle
ω	Frequency in radians per second
Q	Band-pass sharpness (F_o/B_{3dB})
z_i	Transfer function zero
p_i	Transfer function pole

A.2 Glossary of Terms

Alias – An apparent in-band frequency component in a sampled signal whose true (analog) frequency is higher. Can cause undetectable corruption to the in-band signal of interest.

Amplitude response – Magnitude of the complex frequency response. Defines the ratio of output to input amplitude. Along with the phase response, forms the total filter frequency response.

Anti-aliasing – Application of a low-pass filter to remove out-of-band spectral components so that their aliases do not corrupt sampled in-band signals.

Attenuation – Input-to-output reduction in amplitude measured in dB; equates to (-1)**Gain*.

Band-pass – Class of filters defined by an amplitude response with a passband of finite width, in which the lower frequency bound of the passband is above DC ($f > 0$).

Band-stop – Class of filters defined by an amplitude response with a stopband of finite width, in which the lower frequency bound of the stopband is above DC ($f > 0$).

Baseband – Any spectral band with a lower bound at DC ($f = 0$).

Cascading – Construction of a band-pass amplitude response by connecting high-pass and low-pass filters in series.

Cutoff frequency – The frequency defining the passband edge of a filter's amplitude response, conventionally defined as the 3.01 dB attenuation level ($1/\sqrt{2}$).

DC gain – Maximum gain for a low-pass filter.

Filter class – Filter classification defined by the relative spectral position of the passband(s), transition band(s), and stopband(s).

Filter prototype – Filter classification based on the mathematical form of the transfer function; can be adapted for different filter classes and orders.

Folding frequency – Frequency about which spectral folding occurs to produce aliases in sampled signals; equal to half the sampling frequency.

Frequency response – Relationship defined by the transfer function that relates filter output (amplitude and phase) to input in terms of frequency.

Gain – Input-to-output increase in amplitude measured in dB.

Gain factor – Coefficient of the factored transfer function (pole-zero form) that defines the frequency response; determines overall system gain between input and output.

Group delay – One of two time delays defined by the phase response. The group delay quantifies the change of phase with frequency; non-constant group delay provides a measure of input-to-output wave shape modification.

High-frequency gain – Maximum gain for a high-pass filter.

High-pass – Class of filters that attenuate frequencies from DC ($f = 0$) to F_c and pass frequencies above F_c.

Impulse response – The time-domain filter output resulting from a unit-impulse input.

Phone: 607-277-3550
Web Site: www.pfinc.com

In-band – The spectral region of interest in any measurement (i.e. the expected frequency range of the desired signal).

Insertion loss – For band-pass filters, the minimum attenuation in the amplitude response.

Low-pass – Class of filters that pass frequencies from DC ($f = 0$) to F_c and attenuate frequencies above F_c.

Minimum attenuation (aliasing) – Measure of anti-aliasing efficacy: the minimum attenuation of in-band aliases, defined as the attenuation of the first spectral component whose alias appears at the highest frequency of interest.

Normalized response – The frequency response over a dimensionless frequency domain, in which the frequency is scaled by the cutoff frequency.

Order (filter) – The order of the polynomial in the denominator of the transfer function, equal to the number of filter poles.

Out-of-band – The spectral region lying outside the region of interest (i.e. outside the expected frequency range of the desired signal).

Overshoot – Measure used to characterize the step response of a low-pass filter, quantified as the maximum response level above the ideal response.

Passband – The spectral range in the amplitude response over which the attenuation is less than a predefined threshold (conventionally 3.01 dB).

Phase delay – One of two time delays defined by the phase response. The phase delay converts the phase shift to a time delay for a given spectral component.

Phase distortion – Defined as the deviation of the phase response from a linear function of frequency. Provides a semi-quantitative measure of input-to-output wave shape modification.

Phase response – Phase angle of the complex frequency response. Defines the phase shift of spectral components as a function of frequency. Along with the amplitude response, forms the total filter frequency response.

Poles (filter) – Zeros of the factored polynomial in the denominator of the transfer function. Equates to the filter order.

Ringing – Describes oscillations in a filter's step response about the expected (ideal) value.

Ripple – Describes prescribed variation in the amplitude response of a filter prototype over the passband and stopband.

Roll-off – A measure of filter sharpness, defined as the average slope of the amplitude response over the transition band.

Settling error – The deviation of the step response from the ideal response, typically expressed as a percentage.

Settling time – The time required for the settling error to fall below a given threshold.

Shape factor – An alternative measure of filter sharpness, given as a ratio of the frequencies bounding the transition band.

Spectral folding – Term used to describe aliasing in sampled signals, which can be represented as repeated folding of the frequency axis at integer multiples of the folding frequency.

Step response – The time-domain filter output resulting from a unit-step input.

Stopband – The spectral range in the amplitude response over which the attenuation is greater than a predefined threshold (80 dB in the examples presented above).

Transfer function – A rational function in the complex Laplace variable s that defines the input-output relationship for the filter.

Transient response – The time-domain response of a filter to a reference waveform (see *impulse* and *step response*).

Transition Band – The spectral range in the amplitude response over which the attenuation lies between the passband and stopband levels.

Undershoot – Measure used to characterize the step response of a high-pass filter, quantified as the maximum response level below the ideal response.

Zeros (filter) – The zeros of the factored polynomial in the numerator of the transfer function.

A.3 Group Delay vs. Phase Delay: Amplitude Modulated Waveform

The phase response is commonly characterized in terms of two time delays, the phase and group delay. For some input signals, the two delays are distinct and measurable in the filter output. As an illustrative example, consider an amplitude modulated (AM) input signal with the following form:

$$x(t) = a(t)\sin(\omega_c t) \tag{A.1}$$

The carrier wave has angular frequency ω_c, while the modulation waveform,

$$a(t) = 1 + A\cos(\omega_m t) \tag{A.2}$$

has angular frequency ω_m, where by design $\omega_m << \omega_c$. Combining (A.1) and (A.2), we have

Phone: 607-277-3550
Web Site: www.pfinc.com

$$x(t) = A\cos(\omega_m t)\sin(\omega_c t) + \sin(\omega_c t) \tag{A.3}$$

Making use of a sum-to-product identity[8] for trigonometric functions, we can write (A.3) as

$$x(t) = \sin(\omega_c t) + \frac{A}{2}\sin\left((\omega_c + \omega_m)t\right) + \frac{A}{2}\sin\left((\omega_c - \omega_m)t\right) \tag{A.4}$$

Between (A.1) and (A.4), we see that the AM signal consists of three spectral components: the carrier signal at ω_c, and two adjacent sidebands at $\omega_c \pm \omega_m$. This is the origin of the term double sideband modulation.

Now imagine passing (A.4) through an analog filter. We'll assume the three components lie entirely within the filter's passband, so that $|H(j\omega_c)| = |H(j\omega_c \pm j\omega_m)| \approx 1$. Passage through the filter then imparts only a phase shift to the components, resulting in an output signal $y(t)$ with the form:

$$y(t) = \sin(\omega_c t + \varphi_c) + \frac{A}{2}\sin\left((\omega_c + \omega_m)t + \varphi_c^+\right) + \frac{A}{2}\sin\left((\omega_c - \omega_m)t + \varphi_c^-\right) \tag{A.5}$$

The phase shifts can be determined by the phase response of the filter, i.e.

$$\varphi(\omega) = \angle H(j\omega) \tag{A.6}$$

Noting again the condition $\omega_m << \omega_c$, the phase shifts of the sidebands can be approximated in terms of $\varphi_c = \varphi(\omega_c)$ via a first-order Taylor series expansion:

$$\varphi_c^+ = \varphi_c + \frac{d\varphi_c}{d\omega}\omega_m \tag{A.7}$$

$$\varphi_c^- = \varphi_c - \frac{d\varphi_c}{d\omega}\omega_m \tag{A.8}$$

Substituting (A.7) and (A.8) into (A.5) and rearranging, we have:

$$\begin{aligned} &y(t) = \sin(\omega_c t + \varphi_c) + \ldots \\ &\frac{A}{2}\sin\left((\omega_c t + \varphi_c) + \left(\omega_m t + \frac{d\varphi_c}{d\omega}\omega_m\right)\right) + \frac{A}{2}\sin\left((\omega_c t + \varphi_c) - \left(\omega_m t + \frac{d\varphi_c}{d\omega}\omega_m\right)\right) \end{aligned} \tag{A.9}$$

We can now recast (A.9) in the form of (A.3) by writing the last two terms as a product[8]:

$$y(t) = \sin(\omega_c t + \varphi_c) + A\cos\left(\omega_m t + \frac{d\varphi_c}{d\omega}\omega_m\right)\sin(\omega_c t + \varphi_c) \tag{A.10}$$

---

[8] $2\sin(\alpha)\cos(\beta) = \sin(\alpha + \beta) + \sin(\alpha - \beta)$

or in AM signal form,

$$y(t) = \left(1 + A\cos\left(\omega_m t + \frac{d\varphi_c}{d\omega}\omega_m\right)\right)\sin(\omega_c t + \varphi_c) \tag{A.11}$$

Notice that the phase shifts of the modulation envelope and the carrier are not equal. To see this clearly, we can express the shifts in terms of the phase and group delays, τ_φ and τ_g, defined in Equations (4) and (5):

$$y(t) = \left(1 + A\cos\left(\omega_m(t-\tau_g)\right)\right)\sin\left(\omega_c(t-\tau_\varphi)\right) \tag{A.12}$$

The modulation waveform, or envelope, of the signal is delayed by the group delay, τ_g, while the carrier wave is delayed by the phase delay, τ_φ. Note that both delays are determined by the phase response at the carrier frequency, ω_c.

To demonstrate this effect, consider the following example. A test signal defined by Equation (A.3) is passed through the LP8F. The carrier frequency is set to $.75F_c$; at this frequency, the LP8F has a flat amplitude response ($|H(j\omega)| \approx 1$), significant phase distortion ($\approx 30°$), and measurable separation between the group and phase delays (Figure 3). The modulation frequency is set to $.05F_c$.

The test signal at the input of the filter is shown in Figure 12A. In Figure 12B, we overlay the unfiltered and filtered signals over a single cycle of the modulated waveform. Since $\tau_g \neq \tau_\varphi$, the output signal is not an exact shifted replica of the input signal, but the phase and group delays are easily recognized from the carrier and envelope, respectively. Upon close inspection, the delays are consistent with those predicted by the phase response at ω_c ($\tau_g = 1.2/F_c$, $\tau_\varphi = .83/F_c$: see Figure 3D). If we evaluate (A.12) using these delays, the result matches the filtered waveform.

The example presented here explains the origin of the term 'group delay'. The AM signal in Figure 12A is equivalent to a train of tone bursts at the carrier frequency, ω_c. If the duration of each burst is long such that the condition $\omega_c >> \omega_m$ holds, then each burst – or group – is delayed by τ_g.

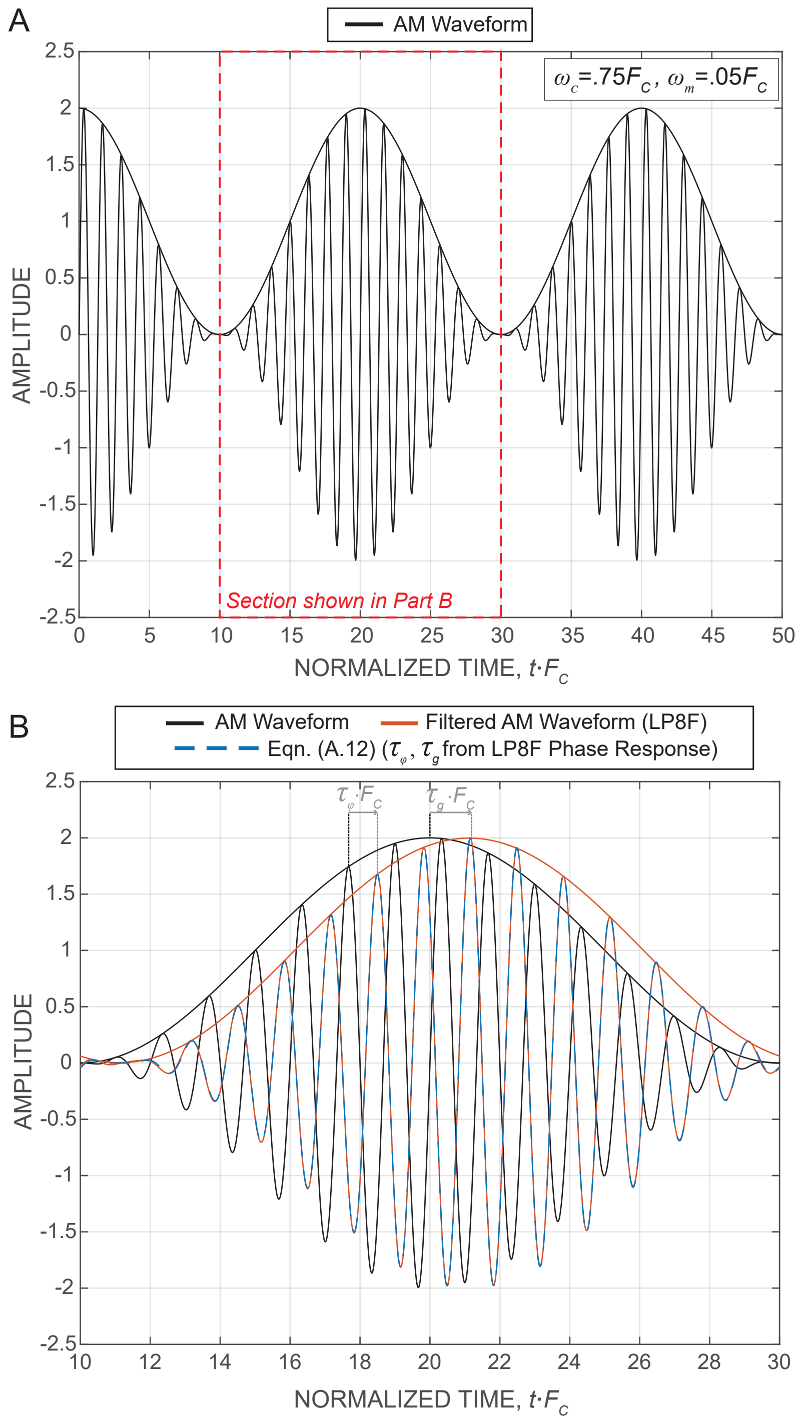

Figure 12. Example of filtering an amplitude modulated waveform. (A) Input AM waveform, as defined by Equation (A.3). The carrier and modulation frequencies are given in terms of the LP8F cutoff frequency F_C. (B) Comparison of AM input and output. The phase and group delays correspond to the carrier and modulation (envelope) delays, respectively. Also shown is $y(t)$ from Equation (A.12), where τ_φ and τ_g are taken from the LP8F phase response at ω_c.

Phone: 607-277-3550
E-mail: sales@pfinc.com

Precision Filters, Inc.
Ithaca, New York
(607) 277-3550
E-mail: sales@pfinc.com

PRECISION FILTERS, INC.

Transducer Conditioning Systems ■ Filter/Amplifier Systems ■ Signal Switching Systems