Part V: Technical Series on Data Acquisition – Filtering

Filtering is used in data acquisition for several reasons:

- To limit the input signal bandwidth to prevent signal aliasing
- To eliminate noise, particularly high frequency
- To eliminate stray pickup, e.g., 50 or 60 Hz

A common reason for filtering is to eliminate signals with frequency components above the Nyquist frequency that would be aliased to lower frequencies. In the case of antialiasing, the filter is a lowpass or bandpass filter, and the filtering *must* occur before the signal is sampled or multiplexed. Thus, filtering *is always* performed on a per-channel basis. Also, if a low-level signal is involved, amplifying the signal prior to filtering is generally necessary to minimize the effects of noise. These two considerations both tend to increase the cost of solving the aliasing problem, since taking advantage of common circuitry after multiplexing several different signals is not possible.

Filtering to limit high frequency noise is another common requirement, and like antialias filtering, it must be implemented on a per-channel basis prior to any multiplexing or sampling. Noise filtering is usually less stringent than antialiasing since the noise floor is typically low, and the filter simply limits the high frequencies.

Filters are not without their own issues. They introduce their own kinds of distortion, and matching their responses across a number of different channels can be expensive, especially when high roll-off filters are needed to minimize the sampling rate.

Antialias Filter Considerations

Designing a system with antialiasing filters involves a number of very critical considerations. These include:

- The system cost (exclusive of filters) for a given sampling rate.
- The filter cost, particularly as sharper filters are used to allow a lower sampling rate.
- The channel-to-channel phase and gain matching required.
- The phase and gain response of the "type" filter needed.

- The frequency characteristics of the signal being acquired.

It is important to realize that *all the above considerations have a direct impact on the cost of the overall solution* and *that they are tightly coupled to each other.*

System Costs

System cost, exclusive of the antialiasing filters, is typically driven by sampling rate or throughput and accuracy requirements. When considering system cost, higher sampling rates mean larger memories, larger disks, higher bandwidth buses, increased processor power, and higher analysis costs (more data to process). A frequently overlooked consideration is that more data may involve more staff time to process, analyze, and interpret the data.

Overall system accuracy is the other cost driver. Assuming a "16-bit" ADC gives you 16 bits of accuracy is not sufficient. Overall system accuracy starts with the sensor, and involves *every component,* up to and including the ADC. Also, system accuracy is typically different at low frequencies than at frequencies near the top of the pass band. If one wants *true 16-bit accuracy,* one must be prepared for a sizable investment and look very carefully at the specifications.

Filter Costs

Generally filters can be categorized by:

- Filter Type (Bessel, Butterworth, Chebyshev, Elliptical,...).
- Number of Poles (typically 2, 4, 6, or 8).
- Channel-to-channel gain match.
- Channel-to-channel phase match.
- Passband ripple.
- Stopband floor.

In general, filter cost is driven primarily by the number of poles, which govern the sharpness of a given filter type, and by the degree of gain and phase match near the cutoff frequency. The number of poles drives the amplifier stages needed to implement the filter and the gain-bandwidth of the operational amplifiers as well as the required precision resistors and capacitors. The gain and phase match, as well as the number of poles, drive the precision of the components needed to implement the filter.

Types of Filters

Lowpass filters provide a means for limiting the bandwidth of signals to be sampled and digitized. Depending on the nature of the application and type of analysis one plans to perform on the data, different filters offer various advantages. In general, one needs to select a filter that preserves the characteristics of the information that is of interest in the application.

This section reviews characteristics of the classical filter designs.

Bessel Filters

The Bessel filter provides a linear phase response with a ripple-free passband and monotonic rolloff. It delays signals at passband frequencies by a constant amount of time and has an overshoot-free step response. The Bessel filter produces a delayed (but accurate) replica of the input signal. These characteristics make the Bessel filter ideal for time domain applications.

The primary limitation of the Bessel filter is its relatively slow roll-off compared to other filters with the same number of poles.

Butterworth Filters

The Butterworth filter has a maximally flat frequency response in the passband with a monotonic rolloff that is much sharper than the Bessel filter. Its phase response varies non-linearly with frequency, the delay is no longer constant and the step response exhibits a moderate amount of overshoot (ringing). These characteristics present no problems for amplitude-based applications. The Butterworth filter is a good general purpose filter.

Chebyshev Filters

The Chebyshev filter provides an even faster rolloff than Butterworth, but exhibits an equal-amplitude ripple across the passband as well as a step response with more overshoot than the Butterworth filter. The phase response is nonlinear and the passband delay is not constant which results in serious ringing in the time domain. This frequency dependence may not present a problem when the primary focus is the attenuation characteristics needed to minimize oversampling and reducing sampling rates.

Elliptic or Cauer Filters

The elliptic or Cauer active lowpass filter has a wide and nearly flat passband response with an extremely sharp roll-off characteristic. It has equal-amplitude ripple in the passband (typically 0.1 dB) and a well-defined attenuation floor in the stopband. This filter also has a nonlinear phase response as well as an overshoot step response. The elliptic filter is ideally suited to amplitude-based antialiasing applications. When selecting an elliptic filter, the attenuation floor must be low enough to ensure that aliased stopband frequencies are sufficiently attenuated to not introduce significant errors into the measurement. This is generally not an issue with other filters since they have monotonic rolloffs.

Filter Implementations

With today's technology, most high-performance filters are implemented as continuous active filters—or in the case of one- and two-pole filters, passive RC filters may be used. Switched capacitor filter technology is improving and is used in some less demanding applications.

Continuous Filters

Continuous active filters introduce minimum noise into the signal; however, it is not easy to provide switchable or programmable bandedges. Components need to be very accurate, especially in the high Q stages used for Chebyshev and Elliptic filter designs.

An efficient way to provide flexibility in filtering is to have a relatively small number of programmable analog bandedges (e.g. decade boundaries) to manage noise and antialiasing considerations, followed by digital filtering to provide exact cutoff characteristics where needed.

Switched Capacitor Filters

The cutoff frequency of switched capacitor filters is determined by controlling the input clock frequency, allowing easy control of the cutoff or corner frequency. Unfortunately, the output of the switched capacitor filter must be filtered again with a continuous filter to reduce clock noise. If this post filter has a fixed bandedges, then obviously it can handle only a limited range of corner frequencies from the switched capacitor filter. The limiting factor in switched capacitor filters for high-performance applications is the noise floor that can be achieved in an environment with a very high frequency clock.

Choosing Filters and Sampling Rates

Clearly, the first decision is selecting the passband frequency range. This is determined by either the bandwidth of the transducer or the bandwidth of the information of interest. Related to selecting the filter is the frequency range of the signals that the filter will pass relatively unattenuated. For a lowpass filter this is characterized by the cutoff frequency f_c where the filter attenuates signals by 3dB. The choice of filter type may play a role as far as the amount of passband ripple (the variation in filter attenuation in the passband) that can be tolerated. Passband ripple affects system accuracy, particularly near the cut-off frequency.

The next decision involves choosing a filter having characteristics consistent with the kind of analysis one intends to perform. Generally, one would choose filters with the steepest rolloff for a given number of poles without compromising the data.

Finally, one must choose a sampling rate and a filter with an appropriate number of poles that will ensure that any aliased signal will be sufficiently attenuated as to have an acceptably small effect on the information content of the data. This is far easier said than done, and the decision can have dramatic effects on the system cost.

Understanding of the frequency spectrum of the input signals is an important consideration. An expensive filter is not needed to deal with a signal that has negligible frequency components above the frequency range of interest. In this case, a relatively simple lowpass filter with a cutoff frequency f_c just above the frequency range of interest will suffice to limit noise and aliasing. A sampling rate consistent with the highest frequency of interest will suffice—typically 3 to 5 times f_c.

In the more general case, the object is to choose a filter with a sharp enough roll-off to ensure that all frequency components of the signal above the Nyquist frequency (one-half the sampling frequency) are attenuated sufficiently so that their contribution to the overall system error is within acceptable limits. Further, the filter should not introduce artifacts (such as those caused by passband ripple) into the signals in the passband that result in an overly large system error.

ADC Quantization Noise

Another consideration is *how small must a signal be before it does not significantly degrade the overall accuracy of the measurement, or, more importantly, does not mask the information of interest?* Certainly the ADC resolution imposes one limitation. For example an N-bit ADC *(N* - 1 bits plus sign) has 2^{N-1} possible "discrete values." Assuming a perfect ADC, the maximum uncertainty of any measured value is ±1/2 least significant bit (lsb) or 1 part in 2^N. When such an ADC is used to digitize a continuous signal, the difference between the quantized—or digitized—value and the signal can be viewed as an error or noise (quantization noise). A more in-depth analysis is provided in the book by Oppenheim[6] which gives a signal to noise ratio (SNR) in decibels (dB) for an ideal ADC of

$$SNR \approx 6B - 1.25dB,$$

where *B* is the number of bits of resolution of the ADC. Attenuating the signal frequency components above the Nyquist frequency significantly below this value provides minimum improvement in overall system accuracy. The table below provides theoretical SNR values for some *typical* ADC resolutions.

Signal Characteristics

In choosing a filter/sampling rate combination, considering the frequency characteristics of the signal is important, specifically the frequency components above the filter cutoff frequency f_c, and particularly those above the Nyquist frequency, f_n,. The *standard operating assumption* in the industry for choosing antialias filters is to *assume the signal*

Table 1: Theoretical Signal to Noise Ratios for Typical ADCs

ADC Resolution	SNR	Limiting Resolution
12-bits	-70.75 dB	0.0290%
14-bils	-82.75 dB	0.0073%
16-bits	-94.75 dB	0.0018%

components above f_n are of equal strength as in the passband. This assumption is probably better for system suppliers than users, as it generally leads to an overly conservative system configuration.

More frequently than not, the major portion of the signal strength is in the passband. If not, then it is possible that the interesting information may lie at higher frequencies and the f_c should have been set higher. For example, if all of the high-frequency signal strengths fall below 10% of the passband frequency strength, then there is 20dB less attenuation needed from the filter to attenuate frequencies above f_n. This translates into a much less expensive filter—or a lower sampling rate.

Selecting the Filter and the Sampling Rate

The various issues for establishing a sampling rate and filter have been discussed. The most practical procedure for reaching a final selection is unfortunately iterative.

1. The desired accuracy of the measurement must be established, but *not lower than the limiting value for the ADC given in the table above.* Making this error percentage artificially small will increase the system cost.

2. The type of filter (Bessel, Butterworth,…) that best meets your needs must be established. One may need to compromise here to get the roll-off one needs to keep sampling rates down. Also, one will need to consider passband ripple and stopband attenuation as well as phase response and step or impulse response of the filter type selected.

3. A cutoff frequency f_c for the filter must be established. The larger one makes this, the higher the sampling rate the system will need.

4. The frequency characteristics of your signal must be determined. How much is the signal strength down at two or three times f_c and above? Some reasonable assumptions should be made here. The better the ratio of passband to stopband signal strengths, the less the filter will need to do.

5. The amount of filter attenuation one needs to attenuate the stopband signals (Item 4) to a value that is below the measurement accuracy (Item 1) must be computed. For example, if the ratio of passband to stopband signal strength is 10:1 or -20dB (Item 4), and the desired accuracy is 0.1% or -60dB, then the filter needs to give 40dB attenuation for frequencies above the Nyquist frequency f_n.

6. The attenuation curves for the filter that one has chosen must be examined to determine the frequency above which the filter attenuation meets or exceeds the attenuation established in item 5. For a 6-pole Butterworth filter, an attenuation of -47dB is achieved at $2.5\,f_c = f_n$. Thus, in this example, a sampling frequency of $2\,f_n = 5\,f_c$ is required, and a 12-bit ADC will provide the required accuracy.

Note that we could repeat items 5-6 for filters of the same type but having a different number of poles to optimize system costs. For example, an 8-pole Butterworth filter would reduce the sampling rate to $4\,f_c$. One could change filter types and repeat items 2-6. For example, an 8-pole Elliptic filter would reduce the sampling rate to $3\,f_c$.

In some applications the data analysis is performed in the frequency domain, i.e., the time domain data is first processed using a Fast Fourier Transform (FFT) or equivalent. In these cases, compromising further with the selection of the filter and sampling rate is possible. Rather than requiring the filter to attenuate signals at the Nyquist frequency f_n, and above to the limiting accuracy, it is sufficient to set this requirement at the *highest frequency of interest* -typically near the filter cutoff f_c. This is illustrated in the figure below with an 8-Pole Butterworth Filter. Using this technique and the example in the figure, a sampling frequency of 3.8 f_c is needed, as opposed to twice the frequency where the attenuation drops below the desired *accuracy floor,* or 5.6 f_c.

Digital Filtering

Solving the aliasing problem must be done **before** *one samples!*

Initially, it may seem that digital filtering doesn't help significantly, but in some applications with critical filtering needs—particularly where there is a need for tight gain and phase matching requirements—digital filtering can be a cost-effective solution. The technique is to use lower-cost, slower-rolloff filters with relaxed gain and phase matching, and use a high speed ADC to oversample. Since the signals of interest are well below the filter cutoff frequency, there is considerably less sensitivity to the gain and phase matching. Also, a tradeoff exists in choosing the sampling rate, since allowing signals to alias into the front-end filter's passband is acceptable if the front-end analog filter removes any aliased signal between zero and the bandpass of the digital filter.

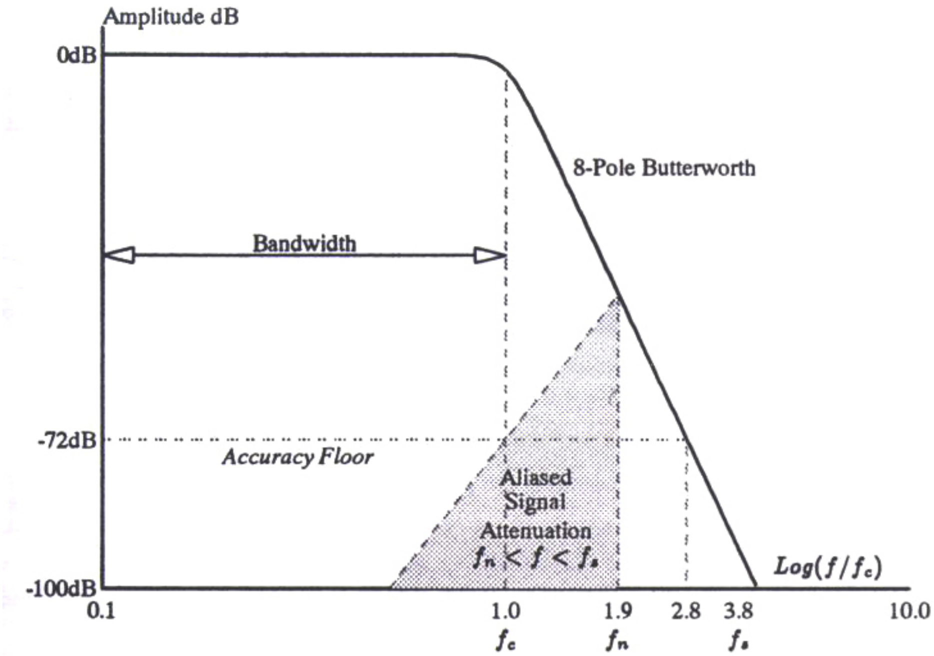

The above figure illustrates an 8-Pole Butterworth filter with the sampling rate set so aliased signals above an accuracy threshold fall back to f_c for frequency domain analysis applications.

Finally, the digitized signal is passed to a DSP to perform the digital filtering and decimation (downsampling with antialiasing), resulting in a reduced output data rate. Since the filtering is now digital, channel-to-channel gain and phase matching are nearly identical to within the precision of the DSP, and implementing a wide variety of filters with custom characteristics is feasible. Further, it is possible to implement digital filters that are beyond today's analog technology.

By decimating the resulting filtered signal at the DSP, it is possible to reduce the effective sampling rate down to *a* value consistent with the passband characteristics of the digital filter and the sampling theorem. This reduces the bandwidth and storage requirements on the rest of the system.

One of the techniques that critically depend on this approach is the Sigma-Delta ADC. With this device, a very fast one-bit ADC is used to highly oversample the signal, and digital filtering techniques are used to "average" the sampled data to gain precision at a lower effective data rate.

Visit www.kscorp.com/Support/whitepapers for other data acquisition exclusives from KineticSystems.