Chapter 14

Digital filtering

TEC/NOT/019

This paper introduces analog filters and their application to flight test instrumentation (FTI).

The following topics are discussed:

14.1 Overview

This paper compares analog filters with the type of digital filter used in Acra KAM-500 modules. It is not intended to be an exhaustive explanation of filtering and sampling theory. The most commonly used FTI techniques and their advantages and disadvantages are qualitatively discussed. This paper shows why digital filtering can provide superior performance to both traditional analog and switched-capacitance filters.

Most transducers used in FTI produce an analog signal that is related to a measurement of interest in some predefined way. For example, Wheatstone bridges formed from strain gages produce signals (on the order of tens of millivolts), which are proportional to the mechanical strains, and therefore the forces some structural members of an aircraft undergo.

Analysis of such signals using modern techniques requires that the data be acquired and stored in digital form. To accomplish this the signal is typically converted to a voltage (if not already so), amplified, offset added and then sampled at regular intervals.

This process necessarily requires the signal to be filtered. The filtering removes the unwanted portion of the signal while preserving the portion of interest. Without this filtering there is the chance that the sampled signal will be aliased and falsely show characteristic not really present in the original signal.

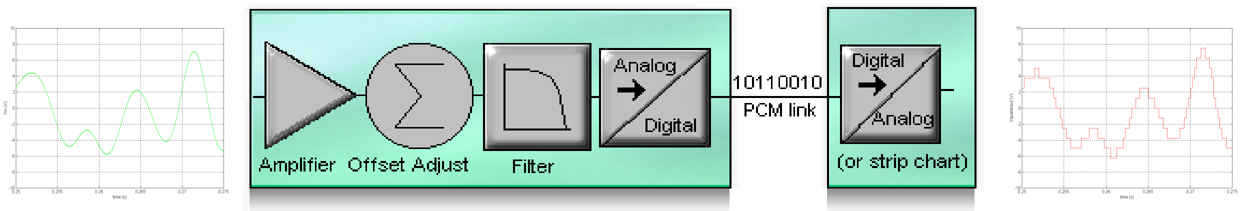

Figure 14-1: An FTI channel

The primary design goal of this filter is to remove the unwanted components of the signal as completely as possible without distorting the portion of the signal of interest. There are three major areas where this design goal may be compromised:

1. The attenuation (reduction of amplitude) characteristic of the filter that reduces the unwanted portion may also affect the signal of interest.
2. The delay through the filter circuit of each component frequency that makes up the desired signal may be different. This shifting distorts the captured signal's time characteristics. For example, a signal used to modulate a carrier may be phase aligned to the carrier before passing through the filter but have an unwanted time offset afterward.

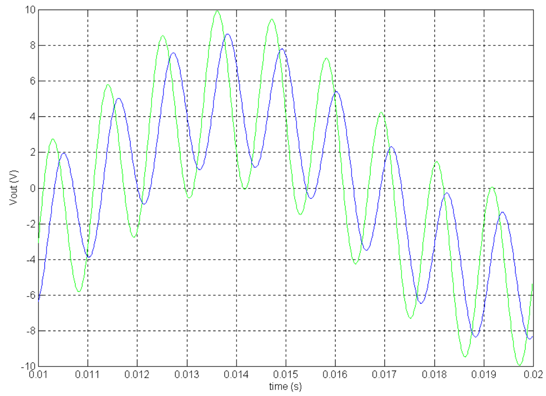

Figure 14-2: Time offset before and after filtering

3. Converting a signal from analog to digital introduces uncertainty (quantization noise) because of finite resolution of the device used to convert the signal. This noise is always present in analog design, but it can be greatly reduced if a digital filter with oversampling and downsampling is used.

This paper will show that using digital filters minimizes or completely eliminates all three of these errors and is, therefore, the best choice.

The analog and digital approaches are very similar. Both have amplifiers and filters before the analog-to-digital conversion. They differ only in that the digital system's additional filtering and fine-tuning of gain and offset is carried out after the analog-to digital-conversion. This significantly simplifies the amplification, offset and in particular the filtering before the A/D. This paper describes the merits of this method.

14.2 Reproducing the signal

The best way to analyze both digital and analog filtering methods is to compare the captured signal to the original. Any difference can be considered an error.

If a slow-moving sine wave (here, slow means the frequency is less than half the sampling rate) is input to the system discussed in the previous section, the reproduced signal is seen to be a series of steps as the output attempts to follow the input. The size of these steps depends on the resolution of the A/D (the errors caused by this are discussed in “14.4 Quantization noise” on page 4).

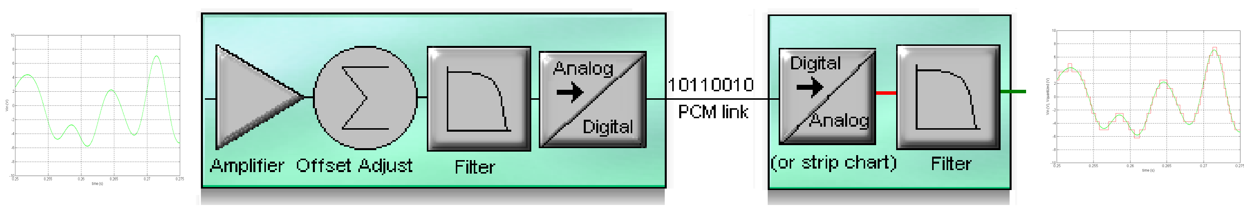

Figure 14-3: Slow signal input into digital filter

For a slow moving sine wave the output appears to follow the input. The following figure shows what happens if a fast sine wave is connected (fast means the frequency is more than twice the sampling rate).

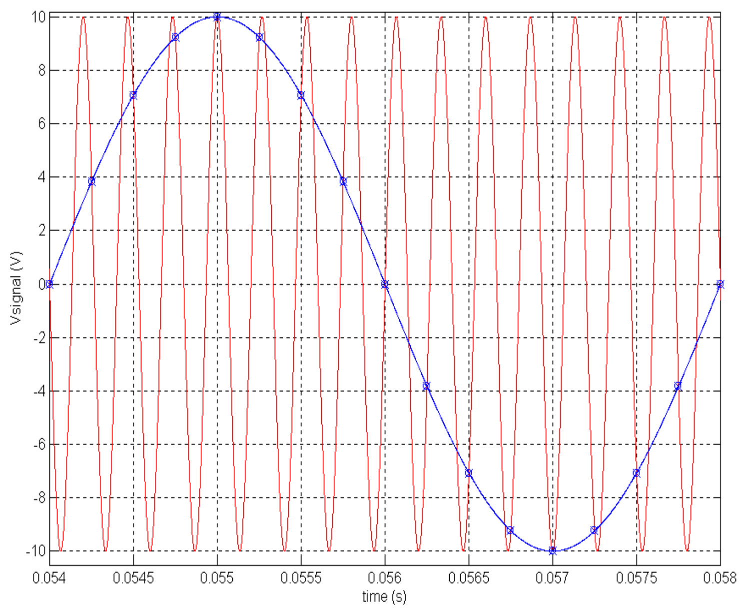

Figure 14-4: Aliasing due to undersampling

The fast sine wave will be reproduced as a slower sine wave. In 1928 H. Nyquist working in Bell telephone laboratories in New York studied this phenomenon and reached the following conclusions for all signals (not just sine waves).

If a signal is sampled at twice the bandwidth of the signal, then the signal can be reproduced exactly.

For low-pass signals this can be paraphrased as:

Assume a perfect filter that does not attenuate sine waves below half the sampling rate, but removes sine waves above half the sampling rate. With one of the filters before the A/D and one after the D/A then:

The filtered signal can be sampled and reproduced exactly* even at each point in time between the samples.

*There will be an error due to the step size of the A/D and D/A which is discussed in “14.4 Quantization noise” on page 4.

These filters are called anti-aliasing (or pre-sample) filters and in FTI are usually low-pass (in that they pass lower frequencies only). The crux of the problem is that the ideal filters discussed above do not exist. Real-life filters are a series of compromises with respect to attenuation, time distortion, power consumption and PCB (Printed Circuit Board) space. The next section discusses some of these trade-offs with respect to analog filters.

14.3 Analog filters

One method of designing analog filters is to use the Sallen-Key low-pass circuit in the following figure *(a)*. At very low frequencies the capacitors act as open circuits and the filter behaves like the voltage follower in the following figure (b). At very high frequencies the capacitors behave as short circuits and the inputs (and therefore output) becomes 0V.

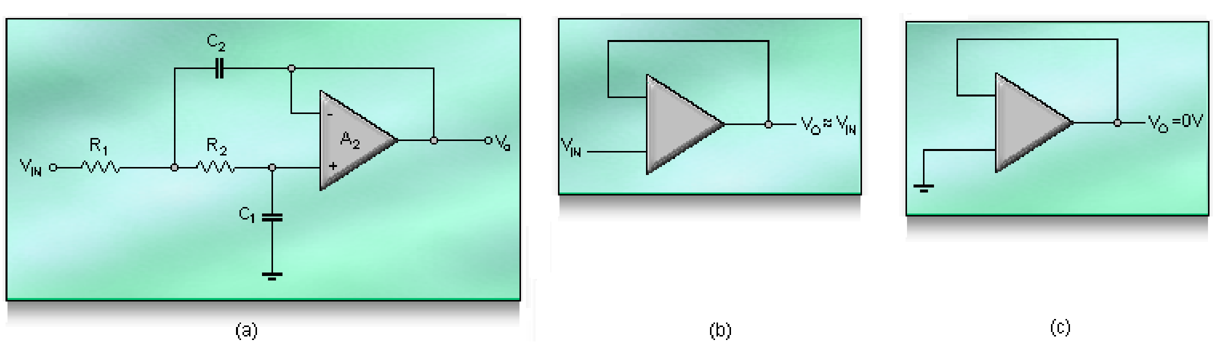

Figure 14-5: (a) Unity-gain Sallen Key 2-pole filter (b) Behavior of filter for very low frequencies (c) Behavior of filter for very high frequencies

This circuit is called a 2-pole low-pass filter. A 2-pole filter must have at least two capacitors. Other types such as 4, 6 or 8-pole filters can be made by cascading two, three or four of these blocks. In FTI, Butterworth and Bessel (Thomson) filters are often used. The attenuation as a function of frequency and the delay as a function of frequency for both Butterworth and Bessel filters are shown in the following figure.

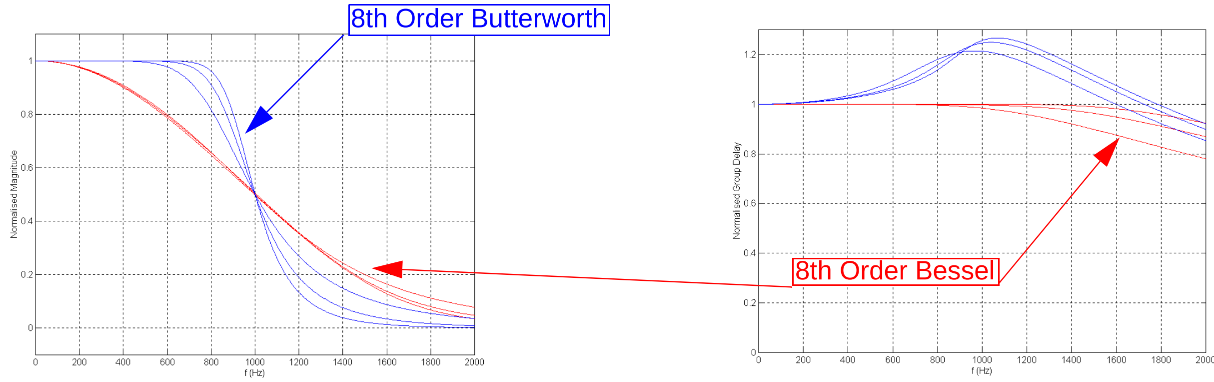

Figure 14-6: Normalized amplitude response and group delay for 4th, 6th and 8th order Bessel and Butterworth low-pass filters

The Butterworth filter has the flattest amplitude response for lower frequencies of any analog filter, however different frequencies are delayed by different amounts causing shape distortion. The Bessel filter has a more consistent delay than any other analog filter type (less shape distortion) but has poor attenuation (it rolls off too slowly thus attenuating signals of interest) in the passband.

A real-life illustration of both of these types of distortion is graphed in “14.7 Conclusions” on page 7.

With analog filters there is a trade-off between amplitude attenuation in the pass-band and shape distortion. Analog filters use capacitors, which means that the characteristics vary from channel to channel and with temperature. Higher order analog filters require even more power and PCB space than lower order filters and introduce more components in the signal path, thus increasing offset, gain, linearity and reliability concerns.

14.4 Quantization noise

Before looking at digital filtering it is worthwhile looking at the effects of the finite step-size of the A/D in more detail. The following figure shows a random signal with no frequency components above 1000 Hz. The signal is sampled at 4000 Hz using a 4-bit A/D.

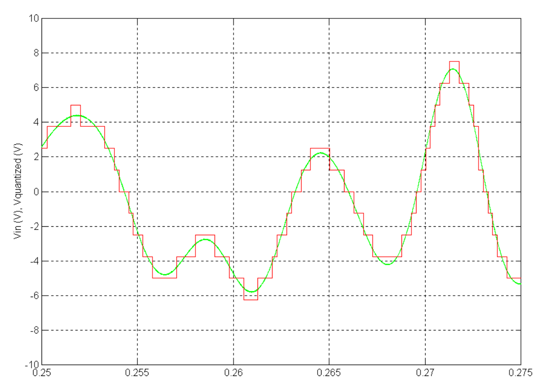

Figure 14-7: Band-limited random signal sampled with a 4-bit A/D at 4000Hz

Allowing for delays in the system the difference between the original signal and the captured signal is as shown in the following figure. This error is due to the finite step size of the 4-bit A/D. For a 4-bit A/D the step size is 1.25V for a range of ±10V (20/16), this means a worst-case error of ±0.625V.

This error has frequency components of up to half the sampling rate. However, if the same signal is sampled with the same A/D operating at, for example, eight times faster than before and then digitally filtered and downsampled [2], then the error will be reduced by the same factor (that is, 1/8). It is important to emphasize here that the traditional analog approach yields the large error (that is, it does nothing to reduce this noise) and the digital approach yields the small error (that is, it significantly reduces this noise).

Another advantage to over-sampling the signal is that real-life noise components above the signal bandwidth (Dither) cause the A/D to toggle between levels rather than return a constant value for slow signals. These components are removed by the filter after the D/A, but in the meantime, have helped provide a more accurate capture of the signal.

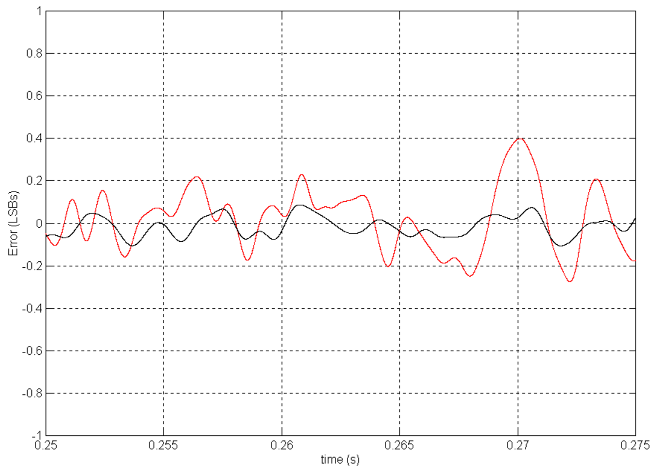

Figure 14-8: Quantization errors with a 4-bit A/D at 4000 Hz and the error from the same A/D at 32 kHz with dither added and filtering after the A/D

Another advantage is that non-linearities (variations in step-sizes) are reduced by filtering after the A/D.

14.5 Digital filtering

Digital filters are implemented after a signal has been sampled. They are mathematical calculations performed on a data series with known and controllable properties. They are not subject to component tolerances. Digital filters can be constructed that are equivalent to analog filters, but the converse is not always true. That is, some digital filters cannot be reproduced with analog circuits.

The following figure shows a generic FIR (Finite Impulse Response) filter. Outputs are calculated by multiplying past inputs by certain coefficients.

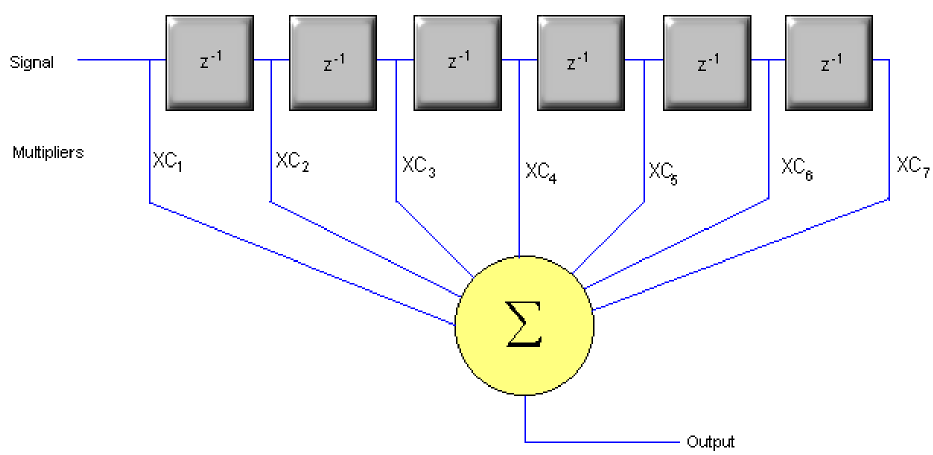

Figure 14-9: Generic FIR filter

Digital filters may be implemented with digital signal processors, lookup tables, or other digital circuitry. A couple of practical concerns are the speed at which the filter calculation can be performed and the precision of coefficients and results.

In Acra KAM-500 modules, signals are sampled many times faster than specified by the user. This over-sampling has the advantage of removing some of the quantization noise as discussed in “14.4 Quantization noise” on page 4. More importantly, the constraints on the analog anti-aliasing filter are much reduced and so it will not typically suffer from the problems of a higher order multi-stage design.

A simple, fixed, second-order anti-aliasing filter is usually sufficient so the problem with noise of cascaded stages is eliminated. When using an oversample-and-decimate approach the cutoff point of this filter may be fixed thus eliminating the noise and inconvenience associated with altering the passive components of the filter.

The digital filter used in Acra KAM-500 modules is a 31-tap half-band filter (sometimes called a Kaiser filter). The following graphs show that it has flatter pass-band response than the flattest (Butterworth) analog and a more consistent delay then the most consistent (Bessel) analog filter.

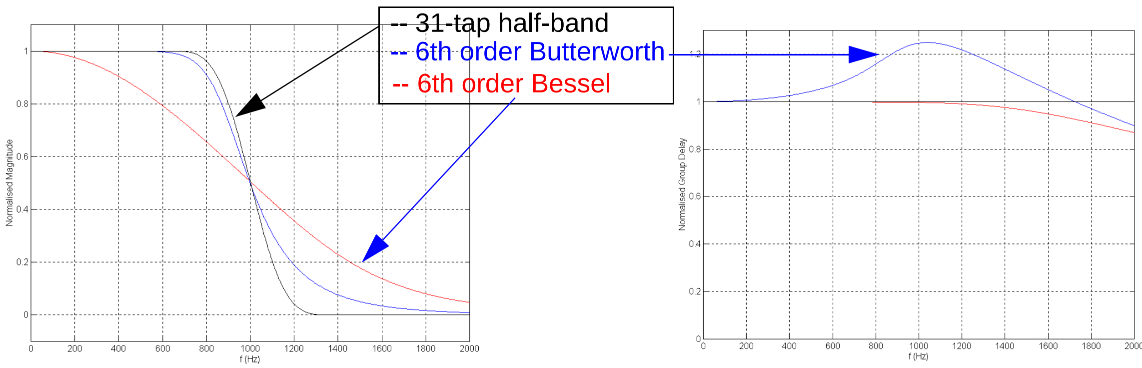

Figure 14-10: Normalized amplitude response and group delay for a 6th order Butterworth, bessel and a 31-tap half-band filter

14.6 Switched Capacitor Filters (SCFs)

In “14.3 Analog filters” on page 3, analog filters built around a Sallen-Key building block were discussed. The statevariable method uses the integrator building block shown in the following figure (a). This integrator can be approximated by the switched capacitor circuit in the following figure (b) where capacitors are charged and discharged through MOSFET switches. Adding this scheme to an active filter allows the adjustment of the resistance and hence the cutoff of the filter by only changing the frequency of a reference clock signal used to switch the MOSFET.

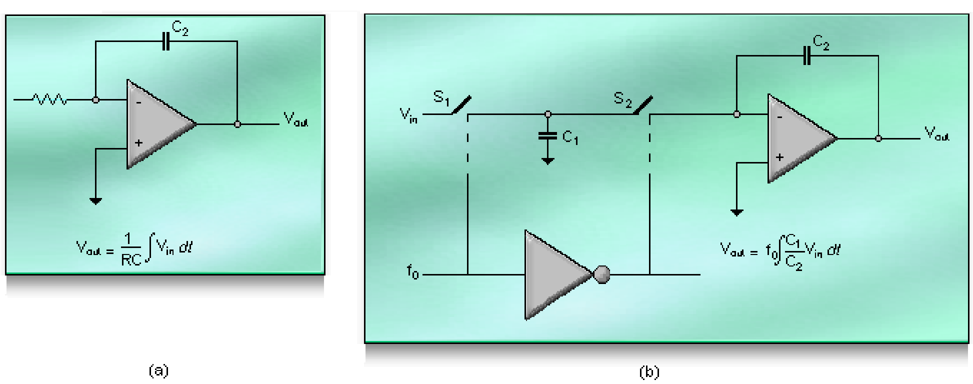

Figure 14-11: (a) Conventional integrator (b) Switched capacitor integrator

This approach uses a minimum number of components, and is very flexible and accurate in terms of adjusting the cutoff frequency. ICs typically have 5th order Bessel or Butterworth responses.

However, after much experimentation with these devices, Curtiss-Wright decided to stop using them in 12-bit systems for the following reasons:

- Harmonics of the noise caused by high-speed switching was very difficult to remove. Ideally a filter would be added before and after the switched capacitor filter (SCF).
- Power consumption per channel due to high-speed switching is almost 125 mW per channel.
- DC offsets, non-linearities and drifts were in the order of 0.1% which is not consistent with a 12-bit system.
- Significant PCB space is required as multiple channels cannot share the resource.

- The SCFs are sampling devices and need anti-aliasing filters. To allow cutoff frequencies more than a decade (10x) apart, multiple SCFs need to be cascaded (just like the FIR filter in “14.5 Digital filtering” on page 5). However, in the case of SCFs that means doubling the errors, PCB space, and power consumption.

14.7 Conclusions

In this paper the following points were demonstrated:

- Sampling at twice the highest frequency component of a signal means that it can be reproduced at the sampling points to within ±1/2 LSB.
- If the reproduced signal is passed through a filter then it can be reproduced even between the sample points to within ±1/2 LSB.
- Signals above half the sampling rate appear as slower frequencies (much like stage coach wheels in old western movies) and as such must be removed before sampling.
- These components are traditionally removed using either a 6th-order Butterworth or 6th-order Bessel analog filter.
- To maximize the amount of signal pass-band, these filters should have cutoff points set as a function of the sampling rate (not easy for active filters).
- Butterworth filters have strong attenuation in the stop-band and little attenuation in the pass-band. However, signal delay varies with frequency thus causing shape distortion. See Figure 14-12 on page 7.
- Bessel filters have less delay distortion then any analog filter. However, they have weak attenuation in the stop-band and significant attenuation in the pass-band. See Figure 14-12 on page 7.
- Switched capacitor filters cannot be used in 12-bit systems where power and space are at a premium.

Digital filters offer the following advantages:

- The 31-tap half-band filter has a flatter amplitude response than the flattest analog filter (Butterworth) (see the 900 Hz component in the following figure).

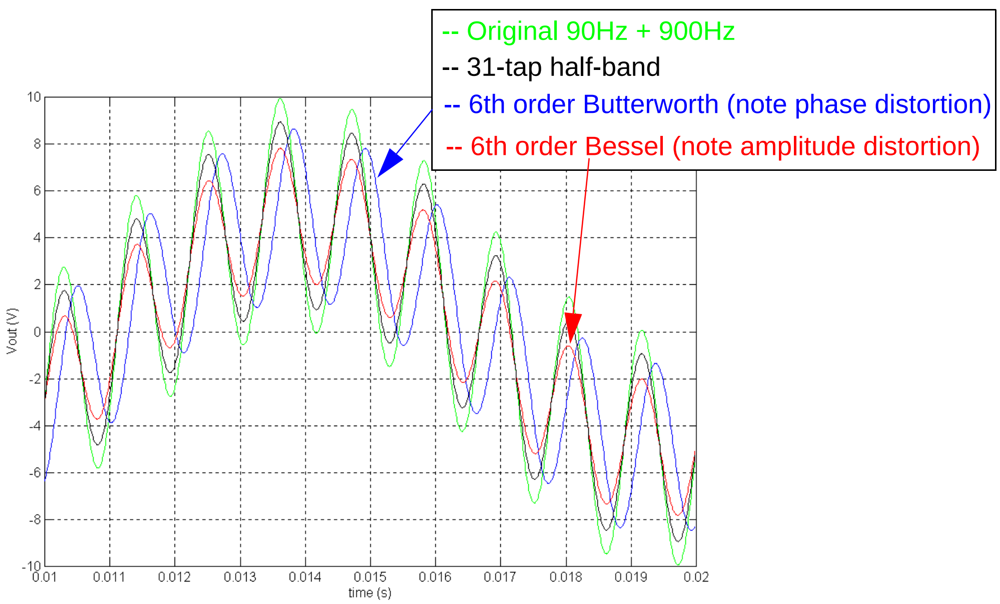

Figure 14-12: A sum of 90 Hz and 900 Hz components passed through a half-band, Butterworth and Bessel filter

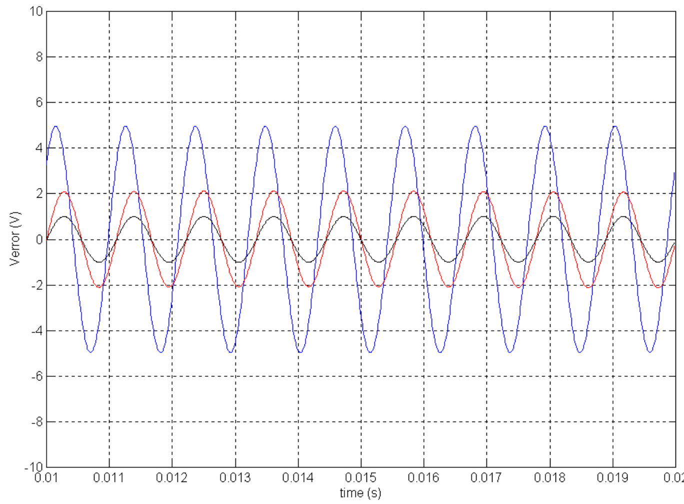

Figure 14-13: Time and amplitude distortions through a half-band, Butterworth and Bessel filter

It can be clearly seen in the previous figure that the digital filter has no time distortion and the least amplitude distortion of all three types of filters. As stated above, this is the design goal of the filtering, and so digital filters are the best choice to accomplish this goal.

- The 31-tap half-band filter has less delay variation than the most consistent analog filter (Bessel).
- Over-sampling and decimating digital filters significantly improves accuracy by reducing quantization noise and non-linearity errors.
- Channel-to-channel response matching of digital filters is exact compared to 5% variations with active analog filters and 0.5% variations with switched capacitor filters.
- Digital cutoff frequencies are an exact function of the final sampling rate. Therefore they are optimum and programmable.
- Doing the filtering and fine-tuning of gain and offset after the A/D means there are considerably fewer components (sources of error) in the signal path.
- Savings in PCB space gained with digital filters and smaller package sizes of today's A/Ds allow good channel density to be achieved with a dedicated A/D for each channel. This eliminates the errors associated with analog multiplexers.

14.8 References

For a detailed mathematical treatment of the issues involved with sampling and digital signal processing see:

[1] Digital Signal Processing - A Practical Approach

Emmanual C. Ifeachor, Barrie W. Jervis

Addison-Wesley Publishing Company

[2] Discrete-Time Signal Processing

Alan V. Oppenheim, Ronald W. Schafer, John R. Buck

Prentice Hall Publishing

[3] A Basic Introduction to Filters - Active, Passive, and Switched-Capacitor

National Semiconductor

Application Note 779

Kerry Lacanette

April 1991

To experiment with A/Ds and filters including the three types discussed in this paper or to reproduce the graphs used:

The MatLAB software package from The MathWORKS® Inc.