

Multiplexing and Sampling Theory

THE ECONOMY OF MULTIPLEXING

Sampled-Data Systems

An ideal data acquisition system uses a single ADC for each measurement channel. In this way, all data are captured in parallel and events in each channel can be compared in real time. But using a multiplexer, Figure 3.01, that switches among the inputs of multiple channels and drives a single ADC can substantially reduce the cost of a system. This approach is used in so-called sampled-data systems. The higher the sample rate, the closer the system mimics the ideal data acquisition system. But only a few specialized data acquisition systems require sample rates of extraordinary speed. Most applications can cope with the more modest sample rates typically offered by mainstream data acquisition systems.

Solid State Switches vs. Relays

A multiplexer is an array of solid-state switches or electromechanical relays connected to several input channels. Although both approaches are used in a wide variety of applications, neither one is perfect; each type comes with various advantages and disadvantages. Electromechanical relays, for example, are relatively slow, about 1,000 samples/sec or less for the fastest reed relays, but they can handle large input voltages, and some can isolate voltages of several kV. A relay's size and

contact type determine its current carrying capacity. For instance, laboratory instrument relays typically switch up to 3A, while industrial applications use larger relays to switch higher currents, often 5 to 10A.

Solid-state switches, on the other hand, are much faster than relays and can reach sampling rates of several MHz. However, these devices can't handle inputs higher than 25V, and they are not well suited for isolated applications. Moreover, solid-state devices are typically limited to handling currents of only one mA or less.

Another characteristic that varies between mechanical relays and solid-state switches is called ON resistance. An ideal mechanical switch or relay contact pair has zero ON resistance. But real devices such as common reed-relay contacts are 0.010 Ω or less, a quality analog switch can be 10 to 100 Ω , and an analog multiplexer can be 100 to 2,500 Ω per channel. The ON resistance adds directly to the signal source impedance and can affect the system's measurement accuracy if not compensated.

Analog switching devices have another undesirable characteristic called charge injection. This means that a small portion of input-gate drive voltage is coupled to the analog input signal and manifests as a spike in the output signal. This glitch produces measurement errors and can be seen riding on the input signal when the source impedance is too high. A compensating circuit can minimize the effects of charge injection, but the most effective method is to keep source impedance as low as possible to prevent it from developing in the first place.

Channel-to-channel cross talk is another non-ideal characteristic of analog switching networks, especially integrated circuit multiplexers. Cross talk develops when the voltage applied to any one channel affects the accuracy of the reading in another channel. Conditions are optimum for cross talk when signals of relatively high frequency and high magnitude such as 4 to 5V signals are connected to one channel while 100 mV signals are connected to an adjacent channel. High frequency multiplexing also exacerbates cross talk because the signals couple through a small capacitance between switch channels. Low source impedance minimizes the cross talk and eliminates the charge injection.

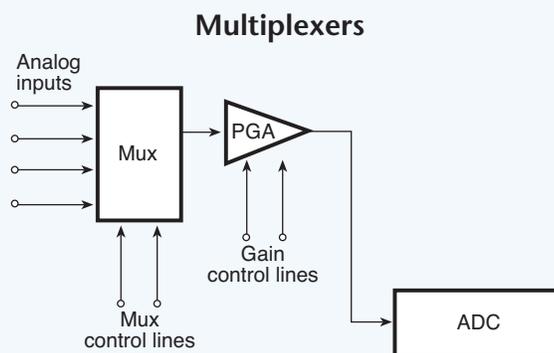


Fig. 3.01. The multiplexer (MUX) is a fast switch that sequentially scans numerous input-signal channels and directs them in a preprogrammed manner to a single ADC for digitizing. This approach saves the cost of using an ADC for each channel.

Scan Sequencer

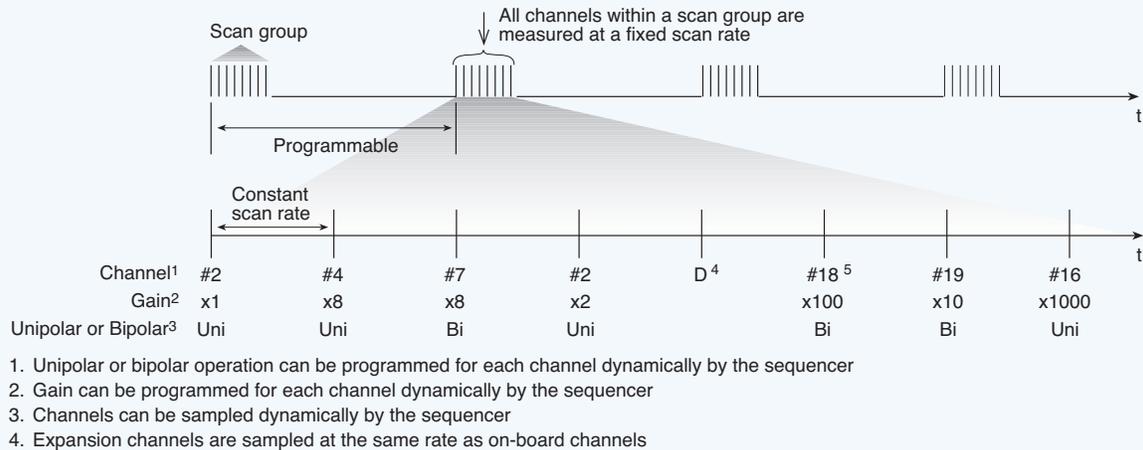


Fig. 3.02. This is an example of a 512 location scan sequencer operating in a 100 kHz data acquisition system. The sequencer can program each channel dynamically for unipolar or bipolar operation. The sequencer also can program the gain for each channel on the fly and change the order in which the channels are scanned.

Speed

Multiplexing reduces the rate at which data can be acquired from an individual channel because of the time-sharing strategy between channels. For example, an ADC that can sample a single channel at 100 kHz is limited to a 12.5 kHz/channel sampling rate when measuring eight channels.

Unfortunately, multiplexing can introduce yet other problems. For instance, the multiplexer's high source impedance can combine with stray capacitance to increase settling time and generate cross talk between channels. Multiplexer impedance itself also can degrade signals. A solid-state multiplexer with an impedance of tens or hundreds of ohms is worse than a relay with a typical resistance of 0.010 Ω or less.

In spite of these negative issues, the advantages of multiplexing outweigh its disadvantages, and it has become a widely used technique to minimize cost without compromising performance. Because the measurement errors are known and specified, they can be compensated at each stage of the data acquisition system to ensure high accuracy at the output.

Sequence vs. Software-Selectable Ranges

Most data acquisition systems accommodate a variety of input ranges, although the manner in which they do so varies considerably. Some data acquisition systems allow the input range to be switched or jumper selected on the circuit board. Others provide software-selectable gain. This is more convenient, but a distinction should be made between data acquisition systems whose channels must all have the same gain and other systems that can sequentially select the input range for each channel. The more useful system accepts different input ranges on different channels, especially when measuring signals from different transducers. For example, thermocouples and strain gages require input ranges in tens of mV and use special signal conditioners, while other sensors might output several volts.

A data acquisition system with a software-selectable range can measure different ranges on different channels (but at a relatively slow rate) with a command to change the gain between samples. But the technique has two problems. First, it is relatively slow. That is, issuing a software command to change the gain of a programmable-gain amplifier (PGA) can take tens or hundreds of ms, lowering the system's sample rate to several Hz. Second, the speed of this sequence is often indeterminate due to variations in PC instruction cycle times. Cycling through the sequence continuously generates samples with an uneven (and unknown) spacing in time. This complicates time-series analysis and makes FFT analysis impossible because its algorithm requires evenly spaced samples.

A better implementation hosts a sequencer that sets the maximum acquisition rate and controls both channel selection and associated amplifier gain at random. For example, one widely used data acquisition system running at 100 kHz and 1 MHz uses software selectable channel gain and sequencing. (See Figure 3.02.) The 100 kHz system provides a 512 location scan sequencer that lets operators use software to select each channel and its input amplifier gain for both the built-in and expansion channels. Each scan group can be repeated immediately or at programmable intervals. The sequencer circuitry overcomes a drastic reduction in the scan rate for expansion channels, a major limitation encountered with many plug-in data acquisition boards.

All channels are scanned, including expansion channels, at 100 kHz, (10 μ s per channel). (See Figure 3.03.) Digital inputs also can be scanned using the same scan sequence employed for analog inputs, enabling the time correlation of acquired digital data to acquired analog data. Such systems permit each scan group (containing up to 512 channel/gain combinations) to be repeated immediately or programmed to intervals of up to 12 hours. Within each scan group, consecutive channels are measured at a fixed 10 μ s per channel rate.

Multiplexer With Sequence-Selectable Gain

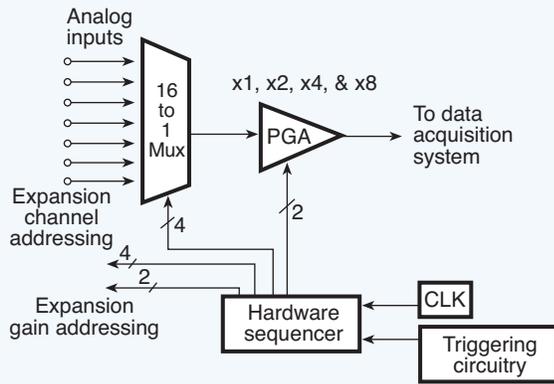


Fig. 3.03. A multiplexer with sequence-selectable gain lets the data system operate at its maximum acquisition rate.

Buffers

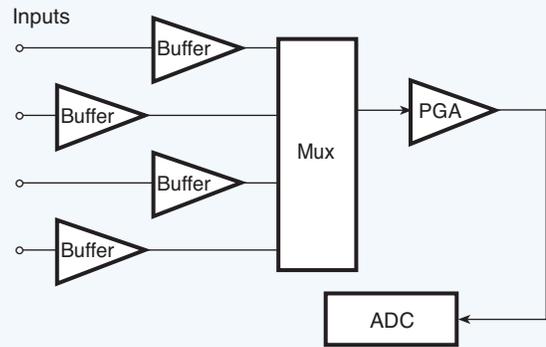


Fig. 3.04. Buffering signals ahead of the multiplexer increases accuracy, especially with high-impedance sources or fast multiplexers

FUNDAMENTAL CONCEPTS

Sampling Rates

When ADCs convert an analog voltage to a digital representation, they sample the value of the measured variable several times per second. Steady or slowly changing dc voltages might require sampling rates of only several Hz, but measuring various ac and sine waves is different. Enough samples/sec must be made to ensure that the wave being measured in the continuous-time domain can be faithfully reproduced in both the continuous-time domain and the discrete-time domain.

Source Impedance

Most signal sources have impedances of less than 1.5 kΩ, so such a maximum source impedance is usually not a problem. However, faster multiplexer rates require lower source impedances. For example, a 1 MHz multiplexer in a 12-bit system requires a source impedance less than 1.0 kΩ. When the source impedance exceeds this value, buffering is necessary to improve accuracy. A buffer is an amplifier with a high input impedance and extremely low output impedance. (See Figure 3.04.) A buffer on each channel located between the transducer and the multiplexer ensures higher accuracies by preventing the multiplexer's stray capacitance from discharging through the impedance of the transducer.

Sample and Hold ADCs

Time Skew

A multiplexed ADC measurement introduces a time skew among channels, because each channel is sampled at a different time. Some applications can't tolerate this effect. But a sample and hold circuit placed on each input ahead of the multiplexer remedies time-skew problems. In a simultaneous sample and hold circuit (SS&H) each channel is equipped with a buffer that samples the signal at the beginning of the scan sequence. The buffer output holds the sampled value while the multiplexer switches through all channels, and the ADC digitizes the frozen signals. In a good

Sample and Hold

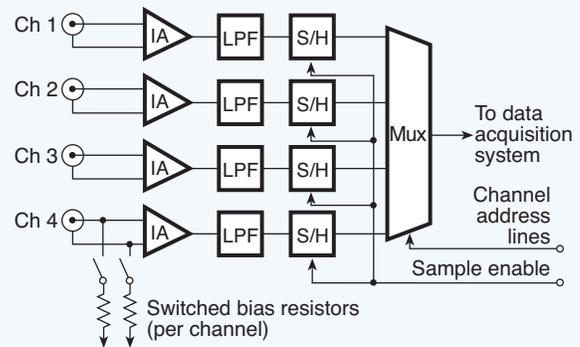


Fig. 3.05. A four-channel multiplexer with low-pass filters and simultaneous sample-and-hold circuits ensures that all the samples are taken within a few ns of each other.

simultaneous sample and hold implementation, all channels are sampled within 100 ns of each other.

Figure 3.05 shows a common scheme for SS&H. Each input signal passes through an instrumentation amplifier (IA), a low-pass filter, and into a sample-and-hold buffer (S/H). When the sample enable line goes high, each S/H samples its input signal and holds it while the multiplexer switches through the readings. This scheme ensures that all the samples are taken within 50 ns of each other, even with up to 256 simultaneous channels connected to a single instrument.

Nyquist Theorem

Transforming a signal from the time domain to the frequency domain requires the application of the Nyquist theorem. The Nyquist sampling theorem states that if a signal only contains

Inadequate Nyquist Sampling Rates

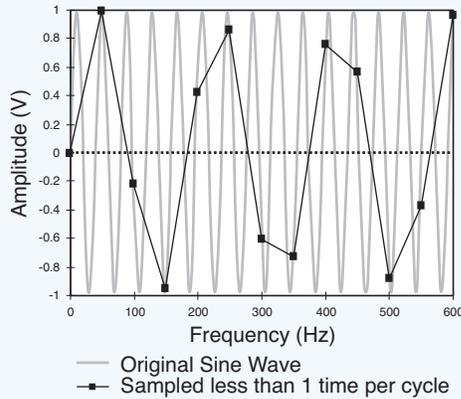


Fig. 3.06. When sampling an ac signal at less than two times the Nyquist frequency, the original waveform cannot be reproduced faithfully.

Acceptable Nyquist Sampling Rates

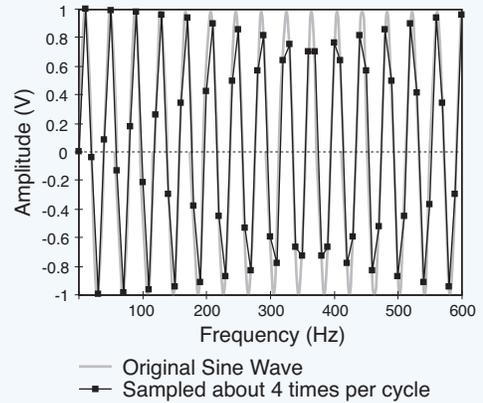


Fig. 3.07. When ac inputs are sampled more than twice the Nyquist frequency of the sine wave, the frequency content of the signal is preserved, and all the Fourier components of the periodic waveform are recovered.

frequencies less than cutoff frequency, f_c , all the information in the signal can be captured by sampling it at a minimum frequency of $2f_c$. This means that capturing a signal with a maximum frequency component of f_{max} requires that it must be sampled at $2f_{max}$ or higher. However, common practice dictates that while working in the frequency domain, the sampling rate must be set more than twice and preferably between five and ten times the signal's highest frequency component. Waveforms viewed in the time domain are usually sampled 10 times the frequency being measured to faithfully reproduce the original signal and retain accuracy of the signal's highest frequency components.

Aliasing and Fourier Transforms

When input signals are sampled at less than the Nyquist rate, ambiguous signals that are much lower in frequency than the signal being sampled can appear in the time domain. This phenomenon is called aliasing. For example, Figure 3.06 shows a 1 kHz sine wave sampled at 800 Hz. The reconstructed or transformed frequency of the sampled wave is much too low and not a true representation of the original. If the 1 kHz signal were sampled at 1,333 Hz, a 333 Hz alias signal would appear. Figure 3.07, on the other hand, shows the signals when sampling the same 1 kHz wave at more than twice the input frequency or 5 kHz. The sampled wave now appears closer to the correct frequency.

Conversely, input frequencies of half or more of the sampling rate will also generate aliases. To prevent these aliases, a low-pass, anti-aliasing filter is used to remove all components of these input signals. The filter is usually an analog circuit placed between the signal input terminals and the ADC. Although the filter eliminates the aliases, it also prevents any other signals from passing through that are above the stop band of the filter, whether they were wanted or not. In other words, when selecting a data acquisition system, make certain that the per channel sampling frequency is more than twice the highest frequency intended to be measured.

Another example of aliasing is shown in Figure 3.08 for a square wave after passing through a Fourier transform. A Fourier transform

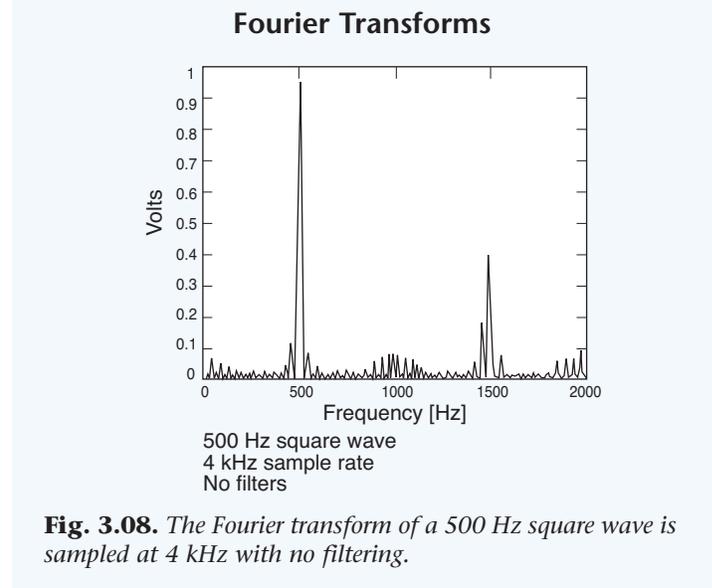


Fig. 3.08. The Fourier transform of a 500 Hz square wave is sampled at 4 kHz with no filtering.

is a frequency spectrum display of the sampled data. It shows how much energy at a given frequency is in a particular signal. For the purpose of illustration, assume that the example processes only frequencies under 2 kHz. Ideally, a Fourier transform of a 500 Hz square wave contains one peak at 500 Hz, the fundamental frequency, and another at 1,500 Hz, the third harmonic, which is 1/3 the height of the fundamental. Figure 3.08, however, shows how higher frequency peaks are aliased into the Fourier transform's low-frequency range. The low-pass filter with cutoff at 2 kHz shown in Figure 3.09 removes most of the aliased peaks.

When the sampling rate increases to four times the highest frequency wanted, the Fourier transform in the range of interest looks even better. Although a small peak remains at 1 kHz, it is

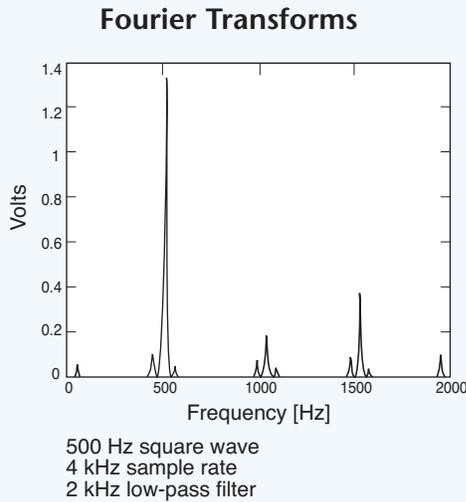


Fig. 3.09. The Fourier Transform of a 500 Hz square wave is sampled at 4 kHz with a low-pass filter cutoff at 2 kHz.

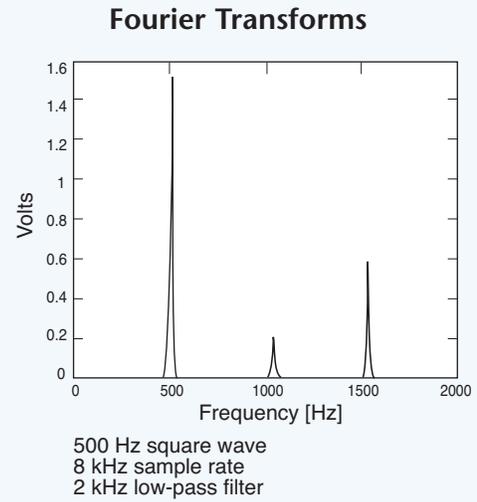


Fig. 3.10. The Fourier Transform of a 500 Hz square wave is sampled at 8 kHz with a low-pass filter cutoff at 2 kHz.

most likely the result of an imperfect square wave rather than an effect of aliasing. (See Figure 3.10).

Discrete Fourier Transform

When ac signals pass through a time-invariant, linear system, their amplitude and phase components can change but their frequencies remain intact. This is the process that occurs when the continuous time domain ac signal passes through the ADC to the discrete time domain. Sometimes, more useful information can be obtained from the sampled data by analyzing them in the discrete time domain with a Fourier Series rather than reconstructing the original signal in the time domain.

The sampled data pass through a Fourier transform function to cull out the fundamental and harmonic frequency information. The amplitude of the signal is displayed in the vertical axis, and the frequencies measured are plotted on the horizontal axis.

Windowing

Real-time measurements are taken over finite time intervals. In contrast, Fourier transforms are defined over infinite time intervals, so limiting the transform to a discrete time interval produces sampled data that are only approximations. Consequently, the resolution of the Fourier transform is limited to approximately $1/T$ where T is the finite time interval over which the measurement was made. Fourier transform resolution can only be improved by sampling for a longer interval.

A finite time interval used for a Fourier Transform also generates spurious oscillations in the transform's display as shown in Figure 3.11. From a mathematical viewpoint, the signal that's instantaneously turned on at the beginning of the measurement and then suddenly turned off at the end of the measurement produces the spurious oscillations. These spurious oscillations are usually eliminated with a function called windowing, that is, multiplying the sampled data by a weighting function. A window function that rises gradually from zero decreases the

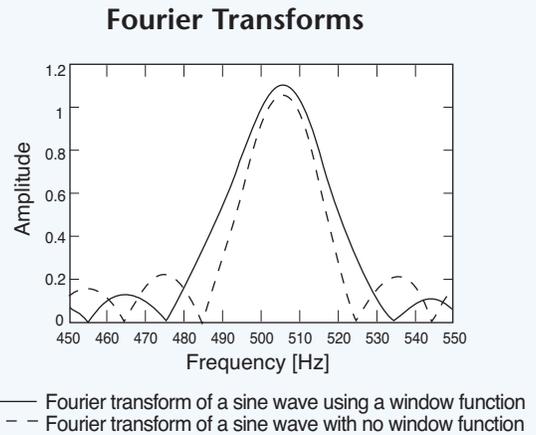


Fig. 3.11. A Fourier Transform is shown with and without a window function.

spurious oscillations at the expense of a slight loss in triggering resolution. Many possible window functions may be used, including Hanning, Hamming, Blackman, Rectangular, and Bartlett. (See Figures 3.12, 3.13, and 3.14.)

Fast Fourier Transforms

The Fast Fourier Transform (FFT) is so common today that FFT has become an imprecise synonym for Fourier transforms in general. The FFT is a digital algorithm for computing Fourier transforms of data discretely sampled at a constant interval. The FFT's simplest implementation requires 2^n samples. Other implementations accept other special numbers of samples. If the data set to be transformed has a different number of samples than required by the FFT algorithm, it is often padded with zeros to meet the required number. Sometimes the results are inaccurate, but most often they are tolerable.

Measurement Without Window Function

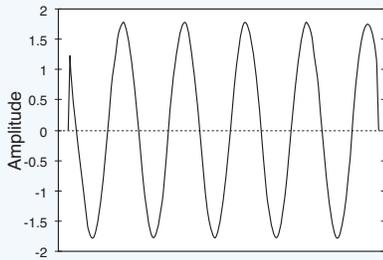


Fig. 3.12. The frequency information obtained from a measurement can be incorrect when the waveform contains abrupt beginning or ending points.

Measurement With Window Function

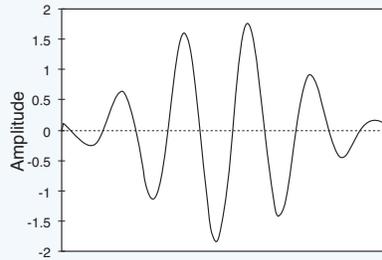


Fig. 3.13. In this case, the frequency information is preserved when multiplied by a windowing function. It minimizes the effect of irregularities at the beginning and end of the sample segment.

Common Window Function

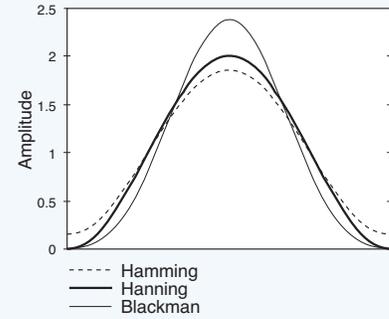


Fig. 3.14. The three most common types of window functions include Hamming, Hanning, and Blackman.

Digital Filtering

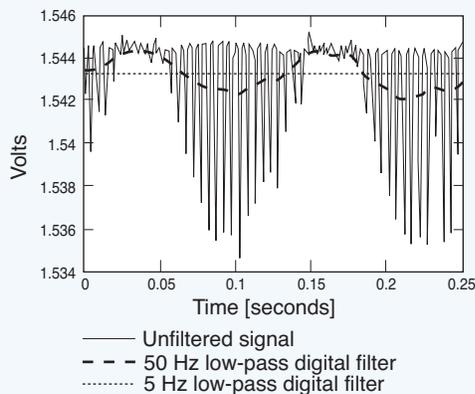


Fig. 3.15. The solid line is the noisy, unfiltered signal, and the two dashed lines show the same signal at the output of the 5 and 50-Hz low-pass digital filters.

Standard Fourier Transforms

A Standard Fourier Transform (SFT) can be used in applications where the number of samples cannot be arranged to fall on one of the special numbers required by an FFT, or where it cannot tolerate the inaccuracies introduced by padding with zeros. The SFT is also suitable where the data are not sampled at evenly spaced intervals or where sample points are missing. Finally, the SFT can be used to provide more closely spaced points in the frequency domain than can be obtained with an FFT. In an FFT, adjacent points are separated by $1/T$, the inverse of the time interval over which the measurement was made.

Many standard numerical integration techniques exist for computing SFTs from sampled data. Any other technique selected for the problem at hand probably will be much slower than an FFT

of a similar number of points. This is becoming less of an issue, however, as the speed of modern computers increases.

Digital vs. Analog Filtering

Digital filtering requires three steps. First, the digital signal is subjected to a Fourier transform. Then the signal amplitude in the frequency domain is multiplied by the desired frequency response. Finally the transform signal is inverse Fourier transformed back into the time domain. Figure 3.15 shows the effect of digital filtering on a noisy signal. The solid line represents the unfiltered signal while the two dashed lines show the affect of different digital filters. Digital filters have the advantage of being easily tailored to any frequency response without introducing phase error. However, one disadvantage is that a digital filter cannot be used for anti-aliasing.

In contrast to digital filters, analog filters can be used for anti-aliasing. But changing the frequency response curves is more difficult because all analog filters introduce some phase error.

Settling Time

Source impedance and stray capacitance influence the multiplexer input settling time. Their effect can be predicted with a simple equation:

EQN. 3.01 Time Constant

$$T = RC,$$

Where:

$$T = \text{time constant, s}$$

$$R = \text{source impedance, } \Omega$$

$$C = \text{stray capacitance, F}$$

For example, determine the maximum tolerable source impedance for a 100 kHz multiplexer. The time between measurements on adjacent channels in the scan sequence is $10 \mu\text{s}$. During time $T = RC$, the voltage error decays by a factor of 2.718, or one time constant, T . But a time interval 10 times longer ($10 T$) is required to reduce the error to 0.005%. Consequently, a fixed time of $10 \mu\text{s}$ between

scans (T_{scan}) and an error of 0.005% requires that $T = 1 \mu s$. But in a typical multiplexed data acquisition system, this is insufficient settling time, so the data would continue to be in error. The relationship can be explained as follows. Most 100 kHz converter's sample and hold circuits are set to acquire the signal at about 80% of the sample window ($10 \mu s$) to allow $8 \mu s$ of settling time. Subtracting this from the scan time results in a sampling time of:

EQN. 3.02 Settling Time

$$\begin{aligned} T_{sample} &= T_{scan} - T_{settle} \\ T_{sample} &= 10 \mu s - 8 \mu s \\ T_{sample} &= 2 \mu s \end{aligned}$$

In a typical 16-bit data acquisition system, the internal settling time (T_{int}) may be $6 \mu s$. The external settling time may then be computed as follows:

EQN. 3.03 External Settling Time

$$\begin{aligned} T_{ext} &= \sqrt{T_{settle}^2 - T_{int}^2} \\ T_{ext} &= 5.29 \mu s \end{aligned}$$

For a 16-bit data acquisition system with 100 pF input capacitance, (C_{in}) and a multiplexer resistance (R_{mux}) of 100Ω , the maximum external resistance is:

EQN. 3.04 External Resistance

$$\begin{aligned} R_{ext} &= \frac{T_{ext}}{(C_{in})(\ln 2^{16}) - R_{mux}} \\ R_{ext} &= 4,670 \Omega \end{aligned}$$

The simplified examples above do not include effects due to multiplexer charge injection or inductive reactance in the measurement wiring. In actual practice, the practical upper limit on source resistance is between 1.5k and 2 k Ω .

Data Acquisition Solutions

LOW-COST MULTIFUNCTION MODULES [USB-1208, 1408, 1608 Series](#)



- Up to 8 analog inputs
- 12-, 14-, or 16-bit resolution
- Up to 200 kS/s sampling
- Digital I/O, counters/timers
- Up to 2 analog outputs
- Included software and drivers

LOW-COST, HIGH-SPEED, 12-BIT MODULES [USB-1208HS Series](#)



- 1 MS/s sampling
- 8 analog inputs
- 16 digital I/O, counters/timers
- Up to 4 analog outputs
- Included software and drivers

HIGH-ACCURACY MULTIFUNCTION [USB-2416 Series](#)



- 24-bit resolution
- Measure thermocouples or voltage
- 32 single ended/16 differential inputs
- Up to 4 analog outputs, digital I/O, counters
- Expandable to 64 channels
- Included software and drivers

HIGH-SPEED, MULTIFUNCTION MODULES [USB-1616HS Series](#)



- 1 MS/s sampling
- Measure voltage or thermocouples
- 16-bit resolution
- Up to 64 analog inputs
- 24 digital I/O, 4 counters/timers
- Up to 4 analog outputs
- Included software and drivers