Introduction to Digital Signal Processing

1.0 Introduction

When a signal varies with time, we are usually concerned not only with its magnitude but also with how it changes. Oscilloscopes, strip chart recorders and other analog recording devices enable us to make observations of the signal by continuously recording and displaying the measurement data in the time domain. When digital computers are utilized for this purpose, however, the magnitude of the signal is sampled only at fixed intervals of time with a complete loss of continuity between. For data acquired in this form, the mathematics of digital signal processing can be used to analyze the signal in both the time and frequency domains. That is, we can know not only how the magnitude of the signal varied with time, but also what the amplitudes of any oscillations were over a spectrum of frequencies.

This Tech Note is intended as a guide to understanding the necessary considerations for converting an analog signal into a useful series of discrete digital data, particularly as related to StrainSmart Data Systems.
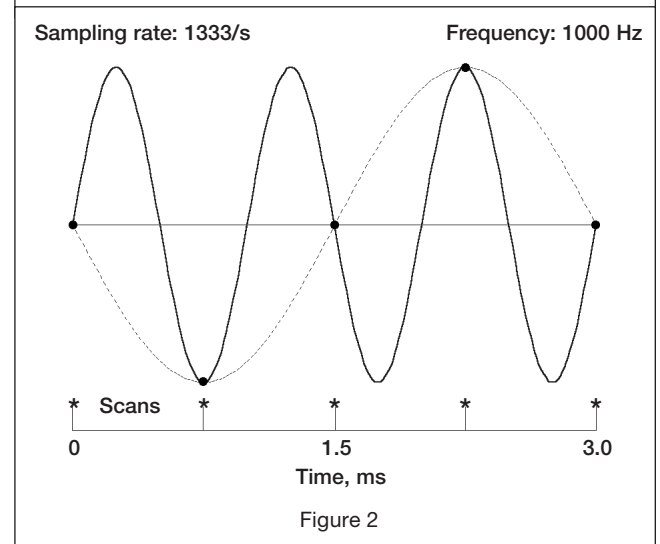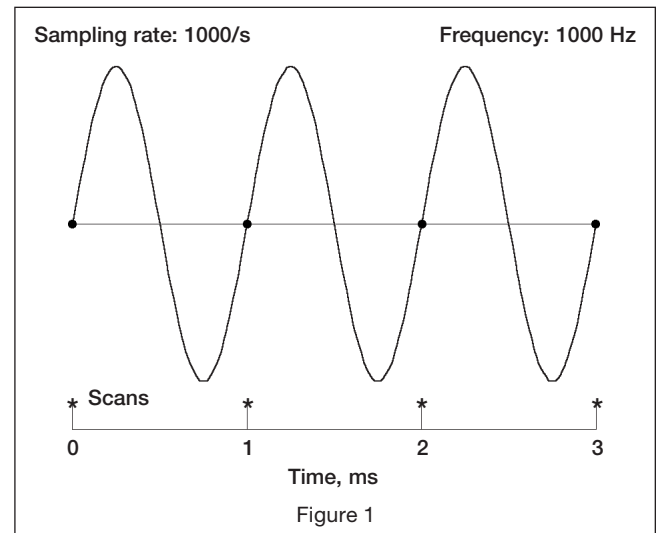
2.0 Signal Aliasing

2.1 Fundamentals

Care must be taken when dealing with digital data to avoid the creation of false, lower-frequency signals by a phenomenon called aliasing. Do you remember seeing the spoke wheels of a wagon that appeared to turn backwards as the wagon rolled across a television or movie screen? That's a false visual impression caused by aliasing. When the wheel is rotating at a slightly slower rate than that at which the frames are projected, the wheels appear to run in reverse. Further, if the wheels were rotated at the same rate as the frames are projected, the wheel would appear to be static, or not turning at all!

Similar things can happen with an analog signal that is sampled periodically. Consider a sinusoidal signal oscillating at, say, 1000 Hz, or 1000 times per second. If we sample this signal at the same rate as the oscillations (Figure 1), we might think the signal was static, not varying with time.

All sampling between 1000 and 2000 times per second would produce a lower-frequency alias. Shown here is the 333-Hz alias that we would "see" at a sample rate of 1333 per second (Figure 2).

Sampling rate: 1000/s   Frequency: 1000 Hz

Scans

Time, ms

Figure 1

Sampling rate: 1333/s   Frequency: 1000 Hz
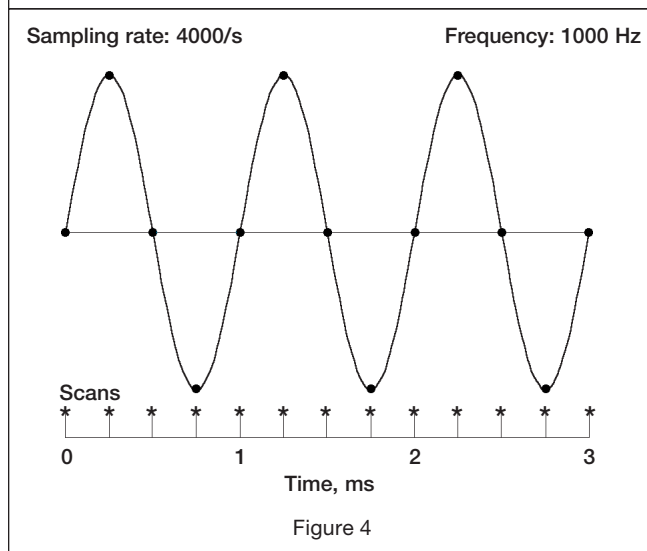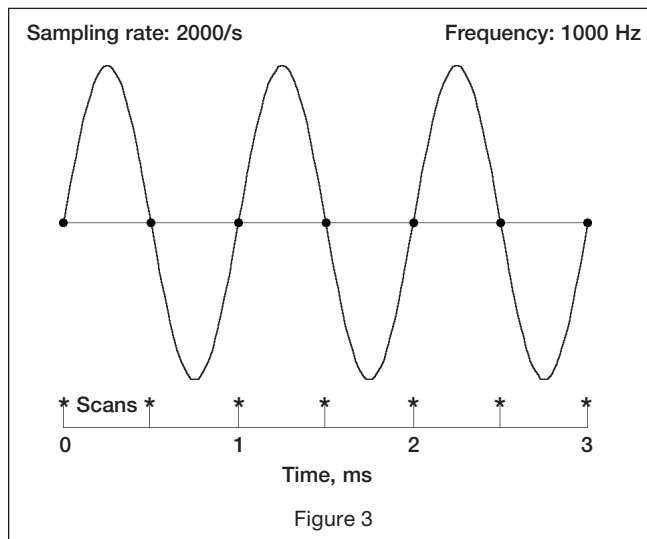
Scans

Time, ms

Figure 2

Sampling at 2000 times per second, the signal would again appear to be static (Figure 3).

So what is it that we must do to ensure that our instrument always sees the wagon wheel turning in the right direction? The answer is to always sample at more than twice the highest rate of oscillation. When sampling at 4000 times a second (Figure 4), we can see that it is impossible to produce either a static or lower-frequency alias from the measurements.

TECH NOTE

Introduction to Digital Signal Processing

Sampling rate: 2000/s   Frequency: 1000 Hz

Figure 3

Sampling rate: 4000/s   Frequency: 1000 Hz

Figure 4

Just how fast should we sample? The mathematics used in digital signal processing can construct a model of the analog waveform that produced a particular set of data by looking at how the magnitudes of the measurements data vary over time. For frequency domain analysis, the sampling rate can be quite close to twice the maximum rate of oscillation. Reconstruction of the actual signal itself would require a sampling rate of ten or more times the highest frequency.

2.2 Anti-aliasing Filters

In the general case, we do not know the frequency of any of the oscillations that might be present in the signal being measured. But, as was just shown, we do know those of half or more the sampling rate will produce aliases during acquisition. Therefore, to ensure that aliases are

prevented, it is necessary to remove all components of the signal and noise with frequencies of half or more the sampling rate with a low-pass anti-aliasing filter. This is an analog circuit through which the input signal must pass on its way to the analog-to-digital converter. Of course, the filter not only eliminates any aliasing in the digital data, but also attenuates any true signals — wanted or unwanted — above the stopband of the filter. Thus, when the data is subsequently analyzed with digital signal processing, we need not worry about any false lower-frequency signals being left in the data by the analog-to-digital conversion process.

3.0 Digital Filters

The sampling rate of the ADC is typically much higher than that required to extract the necessary information from the signal within the frequency range of interest. In addition to allowing unwanted higher frequency components (noise) to remain in the data, these higher sampling rates will also increase data storage requirements and analysis time.
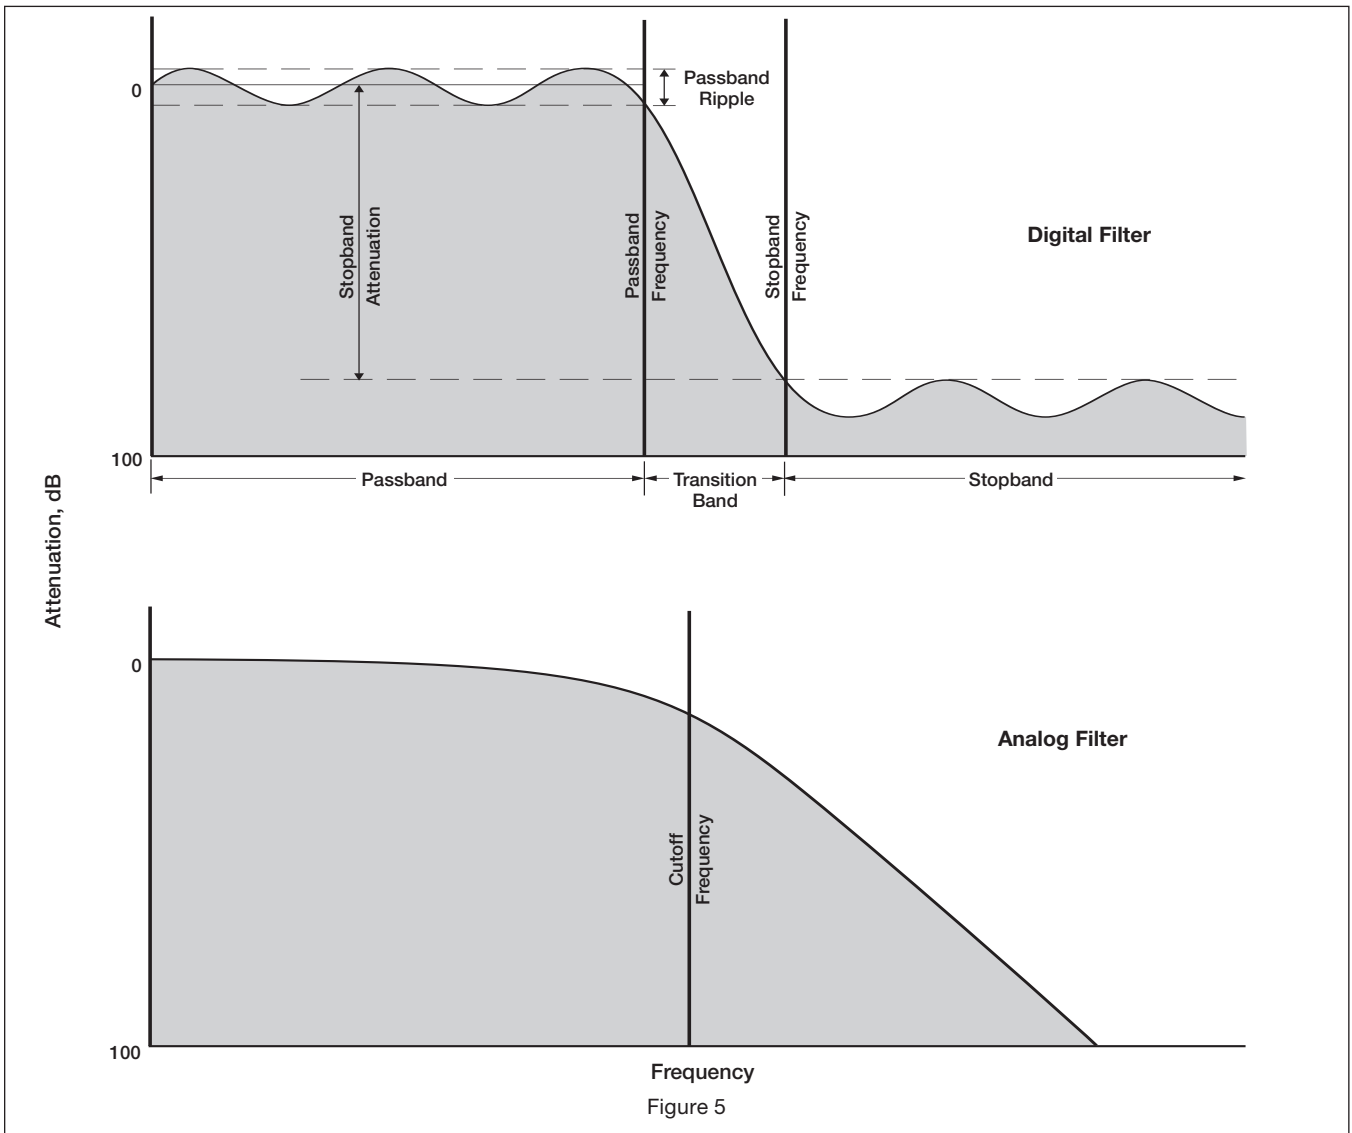
In preparation for acquiring data in digital form, the analog signal being measured is typically passed through an analog filter to ensure that all components of the signal, and noise with frequencies corresponding to a half, or more, of the digital sampling rate of the analog-to-digital converter (ADC), are removed. As described in *Section 2.0*, this helps ensure that false lower-frequency signals — called aliases — are not introduced into the digital data by the sampling process itself.

In order to acquire data at a lower rate while avoiding aliasing errors, it would be necessary to make physical changes to the analog anti-aliasing filter and slow down the sampling rate of the ADC. The disadvantage of this approach is that different analog filter components are required for each sampling rate. A more practical solution is to leave the analog filter and ADC sampling rate unchanged and to mathematically eliminate any unwanted components from the measured signal by passing the digital data through a digital filter.

Specifically, low-pass digital filters enable the digital data coming from the ADC to be "decimated." Instead of the software processing and storing data from every analog-to-digital conversion, the digital filter allows data to be sampled at intervals corresponding to every n^{th} conversion, effectively reducing the sampling rate without introducing aliases. And, at the same time, any unwanted higher frequency components of the measured signal are eliminated from the digital data as well.

Like an analog filter, the digital filter is selected on the basis of which frequencies in the signal are to be retained and which are to be rejected. Low-pass filters, the most common type, are designed to allow signal components

TECH NOTE

www.vishaymg.com   For technical support, contact   Document Number: 11067
170   micro-measurements@vishay.com   Revision 15-Aug-07

Introduction to Digital Signal Processing

Figure 5

from 0 Hz (dc) to some nonzero passband frequency, f_o, to pass essentially unaltered (Figure 5). The filter does introduce a series of small positive and negative deviations from the actual signal in the passband. When this "ripple" exceeds a certain amount, typically 0.01 dB, it defines the passband frequency. For frequencies in the transition band between the passband frequency and higher stopband frequency, the signal is increasingly attenuated. When the attenuation reaches a certain level, typically in the vicinity of 95 dB, it defines the stopband frequency of the digital filter.

When using digital filters, the user should pay attention to both the stopband and the transition band. In some cases, particularly those with lower passband frequencies, the

transition band may be as great as, or even greater than, the passband itself.

Further, as shown in Figure 5, it should be noted that the passband and stopband frequencies of digital filters differ from the cutoff frequency of the commonly used Bessel and Butterworth analog filters. The cutoff frequency of an analog filter, typically specified at an attenuation of 3 dB, usually lies in the transition band between the passband and stopband frequencies of comparable digital filters.

Digital filters are a combination of mathematical algorithms and fast digital circuits that operate on a series of digital data acquired over a period of time. The necessity of using a series of data leads to a delay as the data passes through

Document Number: 11067
Revision 15-Aug-07

For technical support, contact
micro-measurements@vishay.com

www.vishaymg.com
171

TECH NOTE

Introduction to Digital Signal Processing

the filter. After each new sample is taken, the oldest data drops off the front of the series, the remaining data is moved forward in the series, and the data just acquired is added to the end of the series. Then the algorithm is applied to the series of data to obtain a calculated value for the filtered data. The delay, calculated as the time a particular sample takes to get midway through the series, is a function of the ADC sampling rate, the number of terms used in the series, and the passband frequency. Accordingly, the same digital filter should be selected for all measurement channels to ensure that all data acquired at the same time emerges from the digital filters at the same, but delayed, time.

4.0 Throughput Rates of Digital Systems

The electrical resistance strain gage is an inherently analog device that utilizes changes in the relative resistance of the gage to quantify mechanical strains in the surface to which it is attached. Of course, as readers probably already know, the strain gage is typically connected to some form of instrumentation that incorporates a Wheatstone bridge circuit to provide an analog electrical signal that varies as the strain changes. Indeed, most sensors — whether they be strain-gage-based transducers, LVDTs, thermocouples, piezoelectric devices, or a wide variety of others — ultimately produce such a signal.

Unfortunately, the digital computers increasingly incorporated into measurement systems are inherently incompatible with these analog signals. To store measurement data in digital form, the analog signal must be sampled at various points in time and converted to numbers, i.e., the signal must be digitized. Ideally, the time between samples should be vanishingly small (approaching zero). But we know from arithmetic that anything divided by zero is infinitely large. And, of course, the computer can handle only a finite number of data points. The question then becomes how infrequently to sample. If the signal is oscillating on a regular basis, then a minimum of ten data points per period, for the highest frequency component to reasonably reconstruct the signal in the time domain, are typically required. In the frequency domain, any rate of more than two samples per period will suffice. In order for these conditions to be met, a digital measurement system must have a sufficient throughput rate.

In the simplest of terms, the throughput rate is little more than an indication of how much digital data a specific combination of hardware and software can acquire per unit of time. At the instrumentation level, it is primarily controlled by the number of analog-to-digital converters (ADCs) being used in the system, and the rate at which the analog signals being measured can be sampled and digitized. The useful throughput rate of the overall data acquisition system, however, is typically much slower because of such things as (1) the need for oversampling to eliminate aliasing in dynamic signals, (2) the presence of bottlenecks

in the communications link between instrumentation and computer hardware, and (3) limitations in the rate at which software can acquire, reduce, store, and/or present the digital data.

A calculation of throughput also requires knowledge of how the instrumentation hardware acquires the data. The simplest approach is to sequentially sample each data channel in the system at fixed intervals (Figure 6). System 4000, the original Vishay Micro-Measurements data system, acquires data in this fashion, using a single ADC at a throughput rate of 25 or 30 samples per second (depending upon the frequency of the mains power).
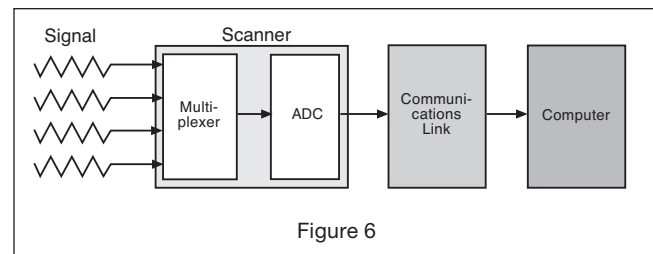
Figure 6

A more complicated approach, at the opposite end of the spectrum, is to simultaneously sample each data channel in the system (Figure 7). System 6000 does this at rates of up to 10 000 samples per second per channel. Because a separate ADC is used for each instrumentation card, the theoretical throughput rate of the system is 10 000 times the number of channels used in the system. For a 100 channel system, that would be a million samples per second. However, because of the limitations in the digital communications link between instrumentation hardware (where the data is acquired) and computer (where the data is stored), the practical maximum throughput of System 6000 utilizing Model 6100 Scanners is about 200 000 samples per second per system for data acquisition and storage. That total can be from a combination of 20 channels acquiring data at 10 000 samples a second, or of 1000 channels acquiring 200 samples a second.
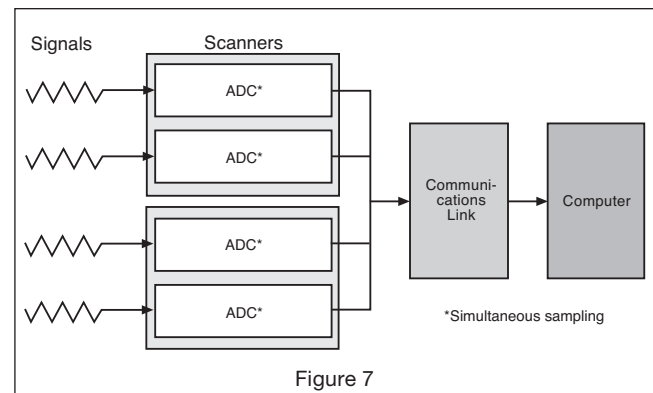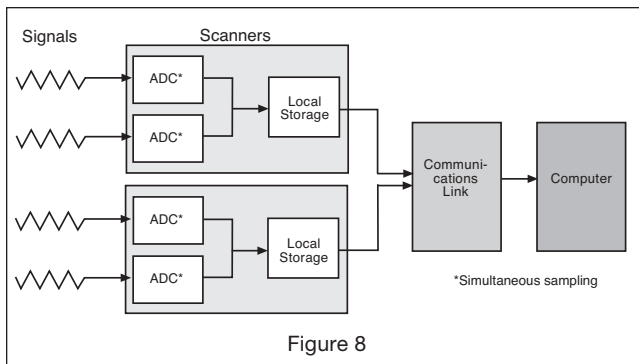
Figure 7

www.vishaymg.com

172

For technical support, contact
micro-measurements@vishay.com

Document Number: 11067
Revision 15-Aug-07

Introduction to Digital Signal Processing

A substantially higher throughput can be obtained with System 7000 (or System 6000 using the Model 6200 Scanners), which simultaneously sample and store data locally on each scanner (Figure 8). With a full complement of sixteen cards, each System 7000 scanner has a practical throughput of 256 000 samples per second (128 channels at 2000 samples/second/channel) or 160 000 samples/second (16 channels at 10 000 samples/second/channel) for the Model 6200. With this combination, the communications link bottleneck is virtually eliminated, and the total system throughput is limited only by the number of scanners used. A System 7000 with 10 scanners, each filled with sixteen cards, would have a maximum throughput rate for data acquisition and storage of 2 560 000 samples per second, for example.

Figure 8

In addition to sequential and simultaneous sampling methods, it is also possible for systems to utilize a hybrid of the two (Figure 9). System 5000 is an example. In this case, all scanners begin sampling simultaneously, but the data from each channel in each scanner is sampled and converted sequentially. Consequently, while the maximum scanning rate for this hardware is 50 scans per second per channel and the maximum number of channels is 1200 per system, the maximum throughput rate of the system is limited by the communications link between scanners and computer to about 12 500 samples per second. That could be 250 channels running at 50 samples a second, or 1000 channels collecting and storing data at 10 samples a second.

Figure 9
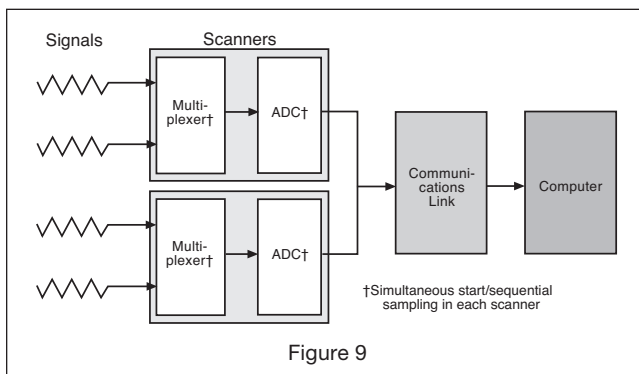
As shown here, the useful throughput of a digital system is a function of many parameters. While it is clear that the throughput rate can be no greater than the sum of the analog-to-digital conversions taking place, it is sometimes less obvious that other processes (digital filtering; transfer of the data to the computer; data storage, reduction and display) are often the limiting factors. Indeed, the real question of throughput is not always how much analog data can be digitized, but rather how much of the digitized data can actually be utilized.

5.0 Scanning Rates

As mentioned in *Section 4.0*, the throughput rate of a digital data system was defined as the total number of useful data points that can be acquired, stored, reduced, or displayed by the system, per unit of time. While many hardware and software factors contribute to the throughput rate, one of the most important is the sampling rate, which can be defined as the total number of useful datum per unit of time *for each signal channel*. While the throughput and scanning rates may be the same for a system containing a single signal channel, the scanning rate is more commonly the throughput rate divided by the number of channels.

For static signals, where the measurand does not vary during the measurement process, the sampling rate is of little consequence. For dynamic signals, unless the user has an understanding of how the system acquires data, how many measurement channels are being sampled, and how the data are to be used, the term can be particularly misleading, resulting in data that is somewhat inaccurate, or completely wrong.

For example, consider a digital measurement system that acquires data sequentially, with a single analog-to-digital converter (ADC) running at 100 samples/second, and a single strain gage on a test part vibrating at 10 Hz. This sampling rate of 100 samples/second/channel provides for 10 datum/cycle/channel, and should be adequate to reconstruct the signal in the time domain. But if the frequency of the signal increases to 100 Hz, the system can provide for only 1 datum/cycle/channel—clearly insufficient to reconstruct the signal. *While both the throughput and sampling rates are unchanged, the data becomes meaningless when the frequency of the signal changes.*

Changes affecting the sampling rate can cause similar cases of undersampling. Consider again the same digital measurement system acquiring data at 100 samples/second and a single strain gage on a test part vibrating at 10 Hz. As before, this sampling rate of 100 samples/second/ channel provides for 10 datum/cycle/channel, and should be adequate to reconstruct the signal in the time domain. But if the number of channels is increased to 10 and a multiplexer is used to sample each channel sequentially, the scanning rate is reduced to only 10 scans/second/channel

Document Number: 11067
Revision 15-Aug-07

For technical support, contact
micro-measurements@vishay.com

www.vishaymg.com
173

Introduction to Digital Signal Processing

and the system can now provide for only 1 datum/cycle/channel. *While the throughput rate and frequency of the signal were unchanged, the data became meaningless when the number of channels (and thus the sampling rate) was changed.*

The only cure for undersampling, of course, is to increase the sampling rate. For systems operating at a fixed throughput rate (like System 4000, the original Vishay Micro-Measurements data system), that usually means decreasing the number of channels being sampled. More sophisticated systems — like the Vishay Micro-Measurements System 5000, 6000 and 7000 data systems — allow the sampling (scanning) rates to be adjusted until the maximum scan rate, maximum throughput rate, or both are reached. System 5000, for example, can scan at 1, 10, 50, and 100 samples/second/channel with a maximum throughput of 12 500 samples/second/system. System 6000 will scan at 10, 100, 200, 500, 1000, 5000, and 10 000 samples/second/channel with a maximum throughput of about 200 000 samples/second/system when using Model 6100 Scanners, and virtually unlimited throughput when using Model 6200 Scanners. System 7000 supports scan rates of 2048, 2000, 1024, 1000, 512, 500, 256, 200, 128, 100, and 10 samples/second/channel with virtually unlimited throughput.

Frequency Response and Sampling Rate

For an analog instrument designed to measure changing input signals, a key parameter is frequency response, or bandpass, stated in hertz (Hz). That is a direct indication of how rapidly the input signal can change, while still being properly amplified and conditioned by the instrument. For a digital instrument used in the same application, the key parameter is the sampling rate, which is the number of samples that can be acquired in a specified period of time, typically one second. Sampling rate is sometimes specified in terms of frequency response (Hz), but while the two are *related*, they are not the same, and should not be used interchangeably.

5.1 Time Skewing

Data acquired from various channels are often functionally related. In stress analysis work, for example, three separate channels might provide strain data from the three grids of a strain gage rosette, which are used together to calculate principal stresses and strains. In this case, *when* the measurements were made is important if the three signals vary with time. When such signals are sampled sequentially, the resulting data are all taken at different times, and are said to be skewed. The errors produced by this skewing depends upon the nature of the signal, the scanning interval (inverse of the scanning rate) and the number of intervals between the sampling of any two data points.
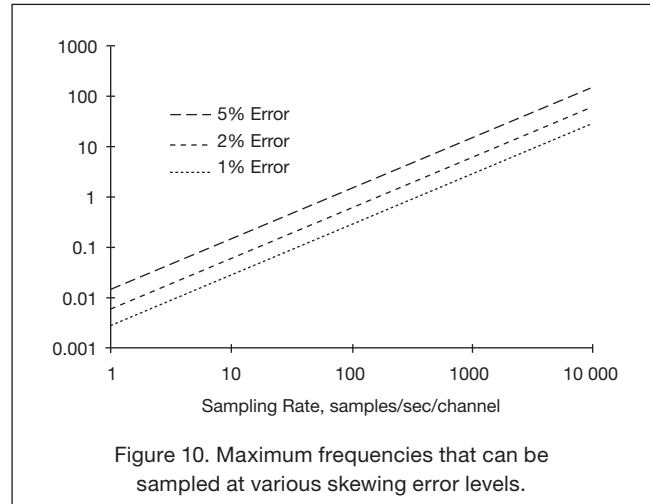

Figure 10. Maximum frequencies that can be sampled at various skewing error levels.

For sinusoidal signals with sequential sampling, the worst-case errors will occur as the signal is crossing through the inflection points. At these points, the maximum frequency that can be sampled without any *detectable* skewing (signal change of 1 LSB, or less, of the full scale signal over a single scanning interval) is a function of the sampling rate and the number of "bits" into which the ADC digitizes the full-scale signal. These frequencies are a few hertz at best, even for relatively high sampling rates. And, of course, the situation worsens as the number of intervals between data points increases.

Skewing errors greater than 1 LSB of full scale *are* detectable and, as shown in Figure 10, the measurable frequency range is increased moderately if the accompanying errors are acceptable.

Of course, all skewing errors can be virtually eliminated by using a digital data system, like System 6000 and System 7000, that features *simultaneous sampling of all signal channels.* In these systems, the maximum time difference between samples is typically in the nanosecond range, such that skewing errors for all measurable frequency ranges are undetectable.

6.0 Uncertainties in Digital Measurement of Peak Signals

For most engineering parameters that require measurement, the test signal varies continuously over time. A plot of these signals, made on any analog data recorder as they vary from one instant of time to the next, produces a line consisting of an infinite number of data points. From a visual inspection of this line, we can immediately see the nature of the variable. Does it increase or decrease? Is it cyclical? What are the maximum and minimum values, and when did they occur? When measurements are made digitally, the time between each conversion of an analog

TECH NOTE

www.vishaymg.com
174

For technical support, contact
micro-measurements@vishay.com

Document Number: 11067
Revision 15-Aug-07

Introduction to Digital Signal Processing

measurement to a digital datum, and the finite data storage capacity of a computer, limits us to the measurement of only a few of the data points on the line of interest. The question then is how many digital data are needed to make a "good enough" reconstruction of the analog signal for us to obtain meaningful measurement results. That depends upon the nature of the analog signal, the digital sampling rate, when the samples are taken, and the accuracy required.

Suppose, as shown in Figure 11, that a signal varies in a sinusoidal way with time. Further suppose that digital samples are taken at ten even intervals throughout each cycle, beginning at the point where the amplitude passes through zero. As we can further see when these measurement values are superimposed on the plot of the analog signal, it is possible to get a vague notion that the signal is sinusoidal. But, the largest values actually measured are only about 95% of the peak value.
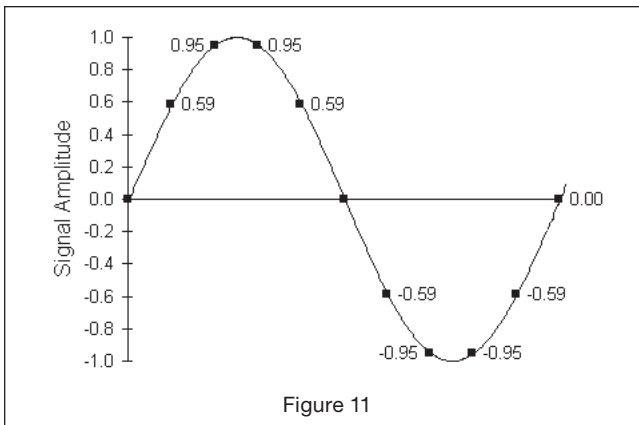
Figure 11

If, however, the start of sampling is delayed by a twentieth of a cycle, as shown in Figure 12, we can still get the same notion about the nature of the signal. But, in this case, the largest measurement values will correspond to the peak values of the signal.
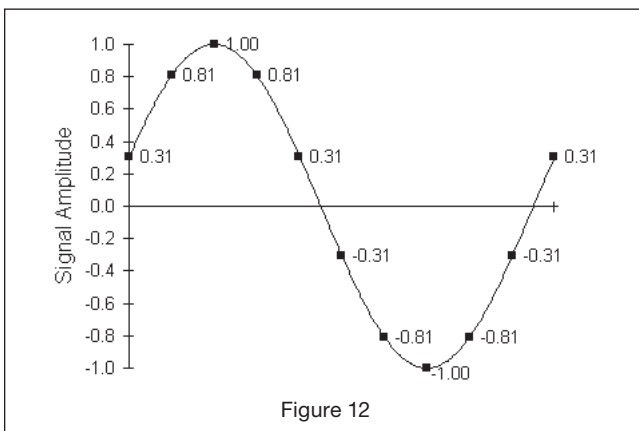
Figure 12

This problem of timing is present in nearly all digital measurements because we can seldom ensure that samples are taken at the peak values. Accordingly, all measurements of analog signals with digital systems will contain some amount of uncertainty with regard to capture of the peak values. In the case of a sinusoidal signal, this uncertainty is a function of the ratio of the digital sampling rate and
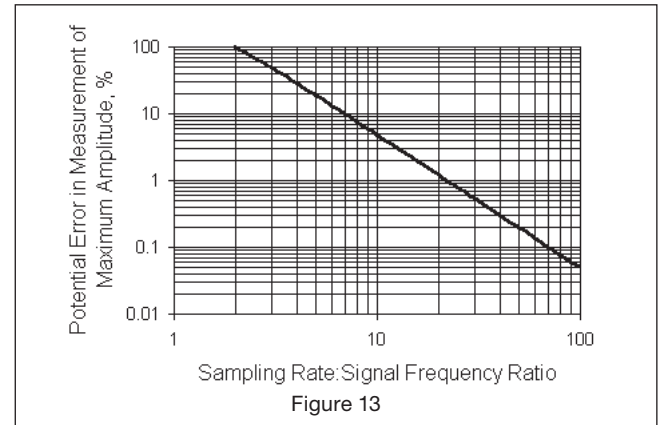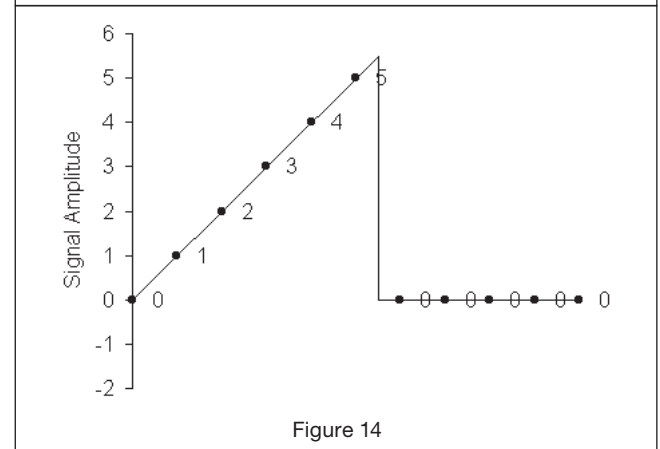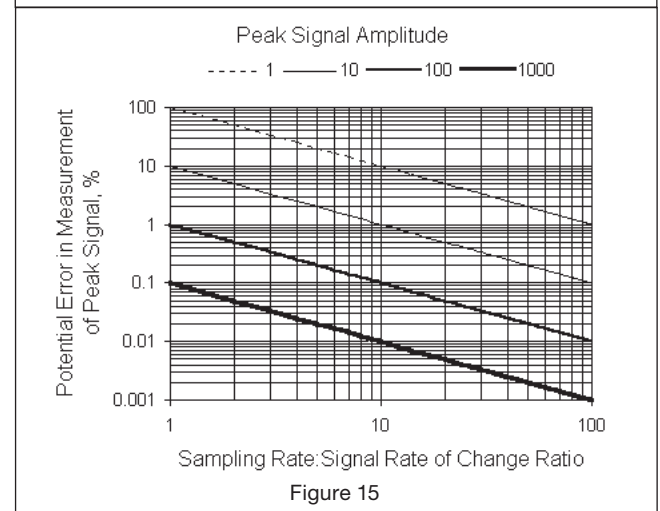
Figure 13

Figure 14

Figure 15

Document Number: 11067
Revision 15-Aug-07

For technical support, contact
micro-measurements@vishay.com

www.vishaymg.com
175

Introduction to Digital Signal Processing

the signal frequency, as shown in Figure 13 for a sinusoidal signal oscillating about zero (with no zero offset).

Of course, many dynamic signals do not vary in a sinusoidal fashion. Consider the further case of a discontinuous signal (Figure 14) that varies linearly with time to some maximum or minimum value before instantly returning to zero (such as would be experienced by a load-bearing structure undergoing a uniformly increasing load until it breaks).

Here the worst case scenario is for the signal discontinuity to occur one sampling interval after the start of the last measurement. (And, because the event can occur at any time during a sampling interval, there is always one sampling interval of uncertainty.) The extent of the error caused by

failure to read the peak value depends not only upon the rate of sampling, but also upon the rate of change in the signal and the peak value of the signal. The uncertainty associated with this error is shown in Figure 15 for various peak values as a function of the ratio of sampling rate and signal rate of change.

Uncertainties, unlike errors, cannot be eliminated from measurements. At best, they can be minimized. And, in the case of digital measurements of peak values, the only recourse for minimizing them is to increase the sampling rate. Particular care should be taken if the per-channel sampling rate of the measurement system decreases with an increasing number of measurement channels.

TECH NOTE

www.vishaymg.com

176

For technical support, contact
micro-measurements@vishay.com

Document Number: 11067
Revision 15-Aug-07