Part II: Technical Series on Data Acquisition – Data Sampling

Introduction

Most data acquisition systems involve the sampling of one or more continuous analog signals. This time-series of discrete samples is typically used to represent the original continuous time-varying signal. To accurately represent a continuous time-varying signal, it is necessary to sample the signal at a sufficiently high rate to correctly represent the time dependence of the signal.

Also, when this discrete time-series of samples is used to extract time or frequency dependent information present in the original signal (typically the case), a phenomenon known as *aliasing* can result. This can cause a significant loss of accuracy and lead to misconceptions about the process being measured. Note that aliasing *does not* infer that the "sampled value" is in error, but rather means a distortion of the inferred time dependence of a series of samples.

Aliasing occurs whenever there are frequency components in the original signal above the Nyquist frequency—half the sampling rate—at the point where the signal is sampled (typically at the ADC). A common oversight is to inject a nice "clean" sine wave from a generator into a data acquisition system, then under-sample the signal. When the resulting digitized waveform is displayed, it appears to have little resemblance to the *expected* sine wave, causing confusion about whether the signal generator or the data acquisition system has possibly malfunctioned.

Sampling Theorem

The ultimate goal of any data acquisition system is to capture critical *information* about the process being monitored. Since digital data acquisition by nature must sample the various signals present, the question is, "How fast should the sampling be done to acquire the desired information?” Obviously, the process and sensor play crucial roles. Generally, the dynamics of the process and the sensor limit the generated signal bandwidth. For example, a thermocouple buried in a copper block will only respond so fast when thrust into boiling water.

Furthermore, the information of interest may lie at a much lower frequency compared to the process dynamics and possibly the sensor. For example, a simple undamped float for measuring the gasoline level in a car would be almost useless, since the information needed is the average level of fuel, not how much it "sloshes" around when the car is moving.

The answer to the sampling question is that a data acquisition system must sample at a rate of at least twice the highest frequency of interest. It must also limit the bandwidth of the signal prior to sampling to avoid aliasing.

In the thermocouple case, the thermal time constant typically limits the maximum rate of change of any signal observed and thus the maximum frequency of interest. The minimum sampling rate is twice the maximum frequency of interest.

The undamped gasoline gauge is more complex. Here, the *maximum frequency* of interest is well below the sensor signal bandwidth—probably an effective time constant of several minutes works. It is necessary to average the input signal in someway to eliminate the *high frequency* noise due to the fuel sloshing in the tank. This can be accomplished in several ways. One technique is to use a low-pass filter before the signal is sampled. The filtered signal is then sampled at a minimum rate of twice the highest frequency in the filtered signal. The alternative is to sample at a much higher rate determined by the raw sensor signal bandwidth and *average* or *filter* the signal using digital techniques.

Even in the thermocouple case, filtering is often desirable to attenuate noise introduced on the sensor cable.

Aliasing

Signal aliasing occurs whenever a signal is sampled that has frequency components above the Nyquist frequency (half the sampling frequency). A simple illustration of this phenomenon is shown in Figure 1. In this figure, a 900 Hz sine wave is sampled at 1000 Hz. The 900 Hz signal is illustrated in light gray and the sampled points by circles. If one looks only at the sampled points, one infers that the original signal was a 100 Hz sine wave, i.e., the 900 Hz signal was aliased to 100 Hz. Obviously, had one sampled the 900 Hz signal at a rate above 1800 Hz (twice the highest input frequency), the inferred signal would be 900 Hz.

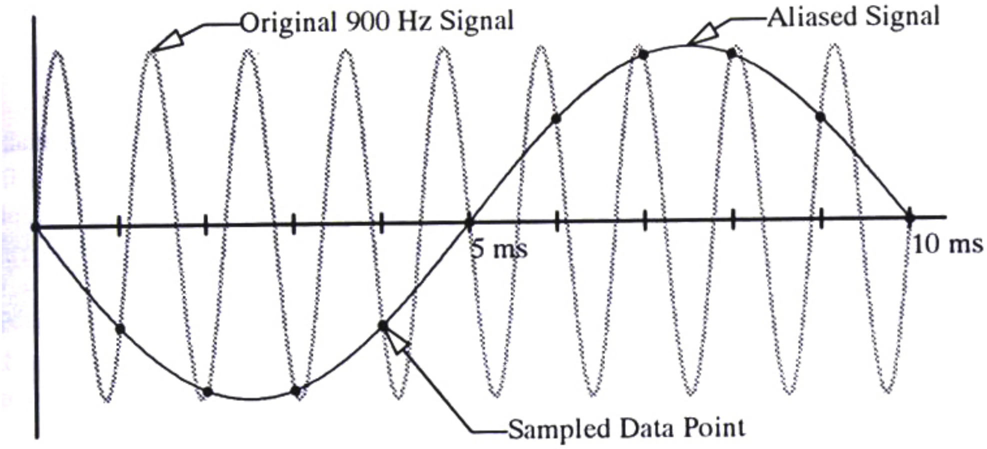

Figure 1: Signal aliasing illustration

Note that the aliasing effect has to do with the act of "sampling" and applies whether the "signal" is a continuous analog signal or a series of discrete digital values. For example, "sampling" every 10th point from a series of digitized points representing a time dependent signal can result in just as much of an aliasing effect as having digitized the original signal at 10th the sampling rate. The graph in Figure 1 is, in fact, composed of discrete samples at 300 dots per inch. Also note that once a signal is aliased, there is no magic to undo the aliasing.

To take this one step further, if one could imagine a frequency domain plot of a signal, the effect of aliasing is to fold the plot accordion-style at multiples of the Nyquist frequency. Thus, a frequency component just above the Nyquist frequency aliases an equal distance below it. A frequency just below twice the Nyquist aliases just above zero frequency, etc.

Anti-aliasing Filtering

The question frequently is "Does one need to worry about aliasing?" The answer is probably *yes* unless the transducers are low bandwidth and one can afford to over-sample, or one is lucky enough to have one of those few applications where all the information is in the sampled signal value, like the charge deposited in a detector by a nuclear particle, and *not* in the time history of the signal.

The other solution to the aliasing question is to limit the bandwidth of the signal *before* it is sampled. This is accomplished with a low pass filter prior to any multiplexing or sampling (*always remember that aliasing is a sampling phenomenon*). The low-pass filter limits the bandwidth of the signal prior to sampling. Such a filter is required for the hypothetical gasoline gauge discussed above.

Establishing a Sampling Rate

One very serious decision which the designer of a data acquisition system must make is that of sampling rate. No one wants to miss information or get erroneous results from aliasing. Therefore, one might conclude that the simplest answer is a "high" sampling rate. The down side of this solution is typically cost, particularly where there are many I/O points and/or a high bandwidth requirement. Most modest systems will fairly easily accommodate an aggregate throughput (sampling rate of each channel times number of channels) of several thousands of samples/second. Aggregate throughputs to 100,000 samples/second are not too difficult to achieve with today's technology. Generally system costs rise quickly for aggregate sampling rates above 500,000 samples/second.

Visit www.kscorp.com/Support/whitepapers for upcoming data acquisition exclusives from KineticSystems. Next month’s featured whitepaper: ***Analog Front-End***.

As a general rule, a little effort to limit the amount of data is well placed since it reduces system cost, storage requirements, and analysis time. This is especially true when dealing with higher throughput applications.