

Basics of data acquisition

This general application note is an introduction to data acquisition for those new to it, to help to ensure the best results when recording with a PowerLab system. (Its text is included in recent Chart and Scope user's guides.)

Written by staff of ADInstruments.



Introduction

The purpose of the PowerLab system is to acquire, store, and analyse data. Figure 1 summarises the acquisition. Usually, the raw input signal is in the form of an analog voltage whose amplitude varies continuously over time. This voltage is monitored by the hardware, which can modify it by amplification and filtering, processes called 'signal conditioning'. Signal conditioning may also include zeroing, for example the removal of an unwanted steady offset voltage from a transducer's output.

After signal conditioning, the analog voltage is sampled at regular intervals. The signal is then converted from analog to digital form before transmission to the attached computer (computers need digital data). The computer software usually displays the data directly; it plots the sampled and digitised data points and reconstructs the original waveform by drawing lines between the points. Digital data can be stored on disk for later retrieval. Software can also easily manipulate and analyse the data in a wide variety of ways.

Most of the parameters that affect acquisition can be set by the user through the software. To make a good recording, the parameters must be appropriate for the signals being recorded. In some disciplines you may be able to find tables of suggested sampling rates, ranges,

and filter settings, but these should not be applied blindly. You still need to know the science (what you are recording, why you are recording it, and what relation it bears to real phenomena) and the technique (how best to record, and what limitations or compromises are inherent in the process).

Sampling rate

Sampling replaces the original continuous analog signal by a series of discrete values (samples) taken at regular time intervals. The appropriate sampling rate depends on the signal to be measured. If the sampling rate is too low, information is irreversibly lost and the original signal will not be represented correctly (Figure 2). If it is too high, there is no loss of information, but the excess data increases processing time and results in unnecessarily large disk files.

Recordings of periodic waveforms that have been sampled too slowly may be misleading as well as inaccurate because of aliasing (Figure 3). An analogy to aliasing can be seen in old films: spoked wagon wheels may appear to stop or even go backwards when their rate of rotation matches the film frame speed — obviously not an accurate record.

Figure 1.
A graphical summary of data acquisition using a PowerLab system




Figure 2. Undersampling: recording rat blood pressure at 20/s and 200/s clearly shows that the former sampling rate is too low to depict the signal accurately





Figure 3. Aliasing: sampling a 3.5 Hz signal at 4 samples/second gives a misleading waveform, an apparent signal of 0.5 Hz (real-life effects tend to be more subtle, though)



To prevent aliasing, the sampling rate must be at least twice the rate of the highest expected frequency of the incoming waveform. This sampling rate is known as the Nyquist frequency, the minimum rate at which digital sampling can accurately record an analog signal. For example, if a signal has maximum frequency components of 100 Hz, the sampling rate needs to be at least 200 Hz to record it accurately. To provide a safety factor to guard against information loss, it is usual to sample at five to ten times the highest expected frequency rather than the minimum two times.


In most cases, the highest expected frequency will be known. It may well be limited by the transducer used: a bridge transducer to measure mechanical force will not produce high frequencies, for instance. If you are unsure of the frequency range (bandwidth) of your signal, a useful rule of thumb is to choose a sampling rate high enough to allow at least 5 to 20 samples for any transient peaks or recurring waves in the signal.

The highest frequencies in a signal can be formally determined by sampling the signal at the maximum rate, and looking at the spectrum of the signal (using the Spectrum Chart extension, or Scope's FFT display setting). High-frequency components in the spectrum with less than 2% or so of maximum amplitude usually contribute little to recording accuracy.

Filtering

Any analog waveform can be described mathematically as the sum of a number of pure sine waves at various frequencies and amplitudes. Low frequencies characterise the slowly changing parts of a waveform; high frequencies, the quickly changing parts. A filter removes selected frequencies from a signal: for instance, a low-pass filter lets low frequencies pass and stops high frequencies. Low-pass filters are commonly used to help reduce noise and give a smoother signal. A high-pass filter removes any steady component of a signal; it also removes slow fluctuations.

Figure 4. The effects of filtering on a mixed-frequency signal: a high-pass filter removes low frequencies; a low-pass filter removes high frequencies




Filters are imperfect. A 200 Hz low-pass analog filter, say, might leave frequencies up to 150 Hz untouched, reduce a 200 Hz signal to 0.7 of its original amplitude (this is its cut-off frequency), and reduce higher frequencies more and more. If you were expecting useful frequencies of up to 100 Hz, you could sample at 400 samples per second and filter out any higher frequencies using a 200 Hz filter. If higher frequencies are left unfiltered, then they could possibly be aliased, resulting in the appearance of spurious lower frequencies.

Filtering can change the signal to some extent: its use must be balanced against the distortions it can remove, such as noise, baseline drift, and aliasing. If the filter settings overlap the bandwidth of a signal, then the signal will have components removed. If you were interested in a waveform with components down to 5 Hz and used a 20 Hz high-pass filter (filtering out 0 to 20 Hz signals), then vital information would be lost from the signal.

Filtering can also be applied to the recorded digital data after acquisition. Chart and Scope have smoothing functions available to help remove noise, clutter, and unwanted high frequencies from signals. These act as simple low-pass filters by averaging adjacent data points, but shouldn't be used as substitutes for the correct low-pass filter during recording. They are most useful in helping clean up signals recorded at high sampling rates.

Figure 5. An electrocardiogram (ECG) with various low-pass filter settings: some filtering cleans up noise, but too much distorts the high-frequency spikes



Digitisation

When analog data such as voltage amplitude is converted to digital form for use by a computer, it must conform to a fixed number of values (for instance, a digital thermometer might only measure temperature to the nearest degree). Any analog values between these values are rounded up or down. The approximation is usually very good, since the number of values is very large. The analog-to-digital converters (ADCs) that perform the digitisation usually generate a number of values equal to a power of two: an ADC with 12-bit resolution can resolve a signal into 2^{12} or 4096 possible amplitude values, which is adequate for most biological signals; an ADC with 16-bit resolution can resolve a signal into 2^{16} or 65,536 possible amplitude values.

Most PowerLab recording units use 16-bit ADCs (some early models had 12-bit ADCs, and some early versions of software used only 12 bits of the data anyway). Chart v4.0 and later, for both Macintosh and Windows, fit 64,000 of the 65,536 values to the range, allowing some leeway above and below it for signal extremes. The values in a 10 V range, say, would be divided into 64,000 fixed values from -10 V to $+10$ V; the minimum change in voltage that could be discerned at that range would be 0.3125 mV. At 10 mV range, the minimum discernible voltage change would be 0.3125 μ V. ADC resolution is part of the hardware, and cannot be set by the user.

Range

Range is inversely proportional to gain, the amount of amplification, and is a more useful concept than gain since it relates directly to the signal being measured. The range can be set independently for each channel on a PowerLab recording unit.

If the signal amplitude exceeds the range, there will be severe loss of information. (This is the same as 'clipping' in a stereo system where music is severely distorted when the amplitude exceeds the capabilities of the amplifier.) Any signal exceeding the range is 'out of range', a condition indicated where no amplitude can be given. If there is any possibility of this condition occurring, you should set the range to a larger value.

For the best resolution, the maximum amplitude of the signal you are interested in should be reasonably close to the chosen range without exceeding it. That way, the


minimum change in voltage discernible in digitisation remains small in relation to the signal being measured (the signal is digitised after it is amplified). If a signal is very small in relation to the range, then its resolution will be degraded. In extreme cases, the recorded waveform may appear stepped rather than smooth. Even though you could see a ± 380 mV signal easily enough at the default 10 V range, you would use the 500 mV range to measure it at maximum resolution. It would be safer in practice to use a 1 V or 2 V range, though, since unexpectedly large peaks could exceed the 500 mV range if the signal was not well-behaved.

Note that changing the display of the waveform on screen (by enlarging it in the Zoom window or by stretching or shrinking its Amplitude axis, for instance) does not affect its resolution, just its appearance.

Noise

Noise is defined as 'unwanted signal'. It is likely to be a problem at lower range settings, when you are trying to measure very small signals. Random noise, such as thermal noise, is inherent in all electronic circuits, including those of the PowerLab recording unit, and generally must be filtered to minimise it. The trick is to choose a low-pass filter setting that removes most of the background noise without unduly changing the signal of interest.

Figure 6. Noise in a signal (top) can mar the original signal (bottom), unless filtered out



For biological signals, differential inputs are preferred, and can reduce common-mode noise due to ground loops. (Ground loops occur where multiple connected pieces of recording equipment are connected to mains power grounds.) PowerLab single-sided inputs are 'quasi-differential', and automatically neutralise up to a fraction of a volt of ground loop noise, the most that normally appears.

Other important causes of noise are stray electromagnetic and electrostatic fields, and include interference (often at the mains frequency of 50 Hz or 60 Hz) from unshielded power lines, switching equipment, fluorescent tubes, transformers, computers, network cables, VDUs, and so on. The interference can interact significantly with a recorded signal. Reasonable care in the arrangement and shielding of equipment and cables will reduce such interference. Particularly delicate measurements, however, may require special apparatus and a controlled environment.

Display and expectations

Interpretation of many measurements, particularly biological ones, is based on empirical evidence: thousands of measurements have been made over the years and a profile built up of normal and abnormal readings. The human brain is particularly good at pattern recognition, and a trained user can quickly assess such information when it is presented graphically. Because the expected shape of a waveform is based on what has been seen previously, it may be difficult at first to interpret waveforms presented in a new manner.


For example, electrocardiographs have been around for at least 75 years, and have a standard output on chart paper with 1 mm and 5 mm divisions, where the recording speed is 25 mm/s, and a 1 mV signal would typically measure 10 mm vertically. Chart and Scope allow a flexible data display, so that a waveform can be reshaped and resized at will. A perfectly valid waveform may look abnormal if one expects it to have the same size and aspect ratio as a standard ECG and it is displayed differently. If the underlying data are valid and you want to make the waveform's appearance more standard, you should change the display settings, not the recording settings.

The screen display on a monitor is generally about 72 dots per inch, so the apparent resolution may be poor if

the display is kept small (as it would be, say, if you were trying to keep an ECG the same size as standard chart-paper traces). The signal might appear jagged and unresolved on screen. Fortunately the resolution of the underlying recorded data is independent of the resolution of the display: even if the channel is very thin or not visible on the screen, sampled data are recorded at full fidelity, as can be seen by expanding the channel display or examining the waveform in the Zoom window.

High-resolution printing will show the selected waveform accurately. Precisely because it is at a higher resolution, a print-out might not appear as smooth as the pen output of a standard electrocardiograph, say, but could easily be smoothed using the smoothing functions of Chart and Scope to give a more standard-looking result.

Figure 7.
An ECG looked at with various display settings changed — this is the same waveform in each case



It is important to check display settings and axis labels carefully when examining a waveform, to be sure of what you are looking at, especially if the settings on your machine might have been altered. A waveform may look very strange if it has been stretched vertically, compressed horizontally, and smoothed as well.

References

This introduction has provided a basic look at some aspects of data acquisition technique. More detailed information on experimental technique and analysis should be available for your particular field. The recommended electrophysiology reference is:

J. Dempster, *Computer Analysis of Electrophysiological Signals* (Academic Press, London, 1993).

Document Number: AGB01c

Copyright © 2002 ADInstruments. All rights reserved.

MacLab, PowerLab, and PowerChrom are registered trademarks, and Chart and Scope are trademarks, of ADInstruments. Other trademarks are the properties of their respective owners.

Contacts

International (Australia)

Tel: +61 (2) 9899 5455
Fax: +61 (2) 9899 5847
E-mail: info@adi.com.au
Web: www.adinstruments.com

North America

Tel: +1 (888) 965 6040
Fax: +1 (866) 965 9293
E-mail: info@adinstruments.com

Europe

Tel: +44 (1865) 891 623
Fax: +44 (1865) 890 800
E-mail: info@adi-europe.com

Japan

Tel: +81 (3) 5820 7556
Fax: +81 (3) 3861 7022
E-mail: info@adi-japan.co.jp

Asia

Tel: +86 (21) 5830 5639
Fax: +86 (21) 5830 5640
E-mail: info@adinstruments.com.cn



www.ADInstruments.com