

So you want to measure something?
– noise, averaging, and bandwidth
 Henk W.Ch. Postma – November 11, 2015

Manual readings

Suppose there is a signal that we want to measure, and it is a little bit noisy. It could be a small voltage V across a resistor that has a current running through it. We are not interested in the fluctuations, but only the average value. Let us assume that we have taken great care to remove line interference (at multiples of the line frequency of 50 or 60 Hz).

We take one reading, call it V_1 , and while we're at it, we take a few more. We get

$$V_1 = 1.3423 \text{ V}; \quad V_2 = 0.40399 \text{ V}; \quad V_3 = 0.51914 \text{ V}$$

and the average is $\langle V_{1..3} \rangle = 0.75513 \text{ V}$. Just to be sure, we repeat the three readings again, and we get a new average $\langle V_{4..6} \rangle = 1.1404 \text{ V}$ (figure 1). We see that there is a significant difference between the readings, and we are not happy with that. So we take the average of all 6 numbers, and get $\langle V_{1..6} \rangle = 0.94779 \text{ V}$.

Computer readings

By this time, we are ready to hook up a computer to the signal and take some of this work off our hands, and we go all out: we take 1000 readings and get $\langle V_{1..1000} \rangle = 0.97096 \text{ V}$. Our sample rate is 100 Hz, so this reading takes 10 s. To verify that our readings are converging to the actual average value, we run a short script in matlab/gnu octave, and arrive at the following plot of the average voltage V_n as a function of how many averages we were taking n (figure 2).

We see that indeed the average values fluctuate quite a lot in the beginning when we're averaging over a low number of readings, but converge to an average value of $\sim 1 \text{ V}$.

Bandwidth and aliasing

When we were doing readings by hand, we took perhaps a few seconds to read the value, record it in our notebook, and then we calculated the average. While we were doing our recordings, however, the signal kept fluctuating. Look at the blue signal in figure 1, and then look at the red marked readings that we took by hand. The fluctuations happened on a time scale that was shorter than we were reading at.




Figure 1: Repeated measurements of a time-varying signal.




Figure 2: Convergence of average readings to the mean.




Figure 3: Fourier transform of signal.

To illustrate, look at the fourier transform of the final readings we were doing by computer in the figure. Indeed we have high-frequency fluctuations that got folded into our measurements even though we were not reading at that high frequency. That is known as ‘aliasing’.

Anti aliasing

In order to get a better reading, we should discard these higher frequencies. Because even if we’re reading at a low frequency, our readings still contain the effect of higher-frequency fluctuations. Removing the high frequencies that occur faster than our sampling rate is known as applying an ‘anti aliasing’ filter. There are many ways you can accomplish this.

We could put a low-pass filter with a cut-off frequency f_0 slightly above what we will be reading at. For instance, we could use a simple first-order RC low-pass filter with a transfer function

$$|H(f)| = \frac{1}{\sqrt{1 + f^2/f_0^2}}, \quad f_0 = \frac{1}{2\pi RC}.$$

This approach requires a-priori knowledge of our measurement frequency and maybe some soldering.

Alternatively, we could measure the signal as quickly as possible with the computer, and average it in software. Let’s say, as in our experiment, we sample the signal at 100 Hz, then load all that data into the computer for averaging.

How averaging works

The signal has a specific power spectral density S_V , in V^2/Hz , and when we read the signal, we get a root-mean-square (RMS) level of fluctuations into the reading equal to

$$V_{RMS}^2 = \int_0^\infty S_V(f) df$$

If we limit the signal to a specific bandwidth B , we basically terminate the integral before f reaches infinity

$$V_{RMS}^2 = \int_0^B S_V(f) df$$

and the RMS signal is smaller, i.e. we have less fluctuations.

More accurately, when we filter, we modify S_V itself because we send it through a filter. If it is a simple RC low-pass filter as suggested above, it has a transfer function with absolute value

$$|H(f)| = |V_{out}/V_{in}| = 1/\sqrt{1 + f^2/f_0^2},$$

where f_0 is the cutoff frequency of the filter. If the noise is white, i.e. $S_V(f) = S_V$, a frequency-independent level of fluctuations, the RMS signal becomes

$$V_{RMS}^2 = \int_0^\infty S_V \frac{1}{1 + f^2/f_0^2} df = S_V f_0 \pi / 2$$

The lower we make the cutoff frequency f_0 , the smaller the level of fluctuations is.

If we average the signal, we can describe that as integrating the signal. That, in turn, can be described as convolving the signal with a scaled rectangular function in the time domain,

$$V_{av}(t) = \int V(t') h(t - t') dt' \quad .$$

If the duration of the average is τ , we are convolving the signal with a modified rectangular function $h(t)$

$$h(t) = \text{rect}'(t) = \begin{cases} 0 & \text{if } |t| > \tau/2 \\ \frac{1}{2\tau} & \text{if } |t| = \tau/2 \\ 1/\tau & \text{if } |t| < \tau/2 \end{cases}$$

and its fourier transform is

$$H(f) = \int_{-\infty}^\infty \text{rect}'(t) e^{-2\pi i f t} dt = \int_{-\tau/2}^{\tau/2} \frac{e^{-2\pi i f t}}{\tau} dt = \frac{\sin(\pi f \tau)}{\pi f \tau} = \text{sinc}(\pi f \tau)$$

Therefore, if the noise is white, the RMS level is (using $\int \text{sinc}^2(x) dx = \frac{1}{2}$)

$$V_{RMS}^2 = S_V \int_0^\infty \left(\frac{\sin(\pi f \tau)}{\pi f \tau} \right)^2 df = \frac{S_V}{2\pi\tau}$$

The longer we average, the smaller the fluctuations become, and the RMS level scales like

$$V_{RMS} \propto \tau^{-1/2} \quad .$$

This behavior is very similar to the discrete case, where the *standard deviation of the mean* is $\sigma \propto n^{-1/2}$.

Let's look again at how the average reading gets better and better the longer we average. We expect the deviation of an average over n numbers ($V_{1..n}$) from the final average value to get smaller and converge as $\propto \tau^{-1/2}$ towards V_∞ . But the deviation could be positive as well as negative. Therefore, if we square the deviation, we get

$$(V_n - V_\infty)^2 \propto \frac{1}{\tau}$$

If we look at the bottom of figure 2, we see that indeed it appears to follow that behavior.

Summary

Measure as fast as you can so you can get many readings in. If you cannot measure as quickly as your signal is varying, filter out the fluctuations above your sample frequency. If you do not have access to a computer that can read the signal quickly, you can filter the signal yourself. Or, you can use a digital multimeter that allows you to increase the integration time.