The Phase Sensitive (Lock-in) Detector

The "lock-in amplifier" is an instrument used in many physics experiments because of its special effectiveness in reducing noise in electrical measurements. But unlike instruments such as oscilloscopes and various types of meters, its operating principle is somewhat subtle. This laboratory will lead the student through the principle of phase sensitive detection—the heart of the lock-in amplifier—and will explore the basic function, features and limitations of the lock-in amplifier—the commercial box built around the phase sensitive detector.

1 Introduction

Many measurements in experimental physics involve the detection of an electrical quantity, either a voltage or a current. Some physical quantities are intrinsically electrical in nature, for example the voltage drop across a diode, or the emission current in a vacuum tube. Other quantities such as temperature, pressure, displacement, or light level can be converted to electrical quantities by means of transducers (sometimes called sensors). Perversely, the electrical quantity (signal) of interest is accompanied by noise, the latter sometimes orders of magnitude greater than the former. Various techniques exist to recover the signal of interest from the composite of signal + noise, and one technique, phase sensitive (lock-in) detection will be explored in this experiment.

Consider the following question: Can one, by being sufficiently careful and clever, reduce the noise content of an electrical signal coming from a transducer to arbitrarily low levels? The answer, in a word, is no. The signal from any transducer with resistive or diode-like characteristics will, from physical principles, have an irreducible amount of (Johnson or shot) noise on its output signal. In addition, some flavor of reducible noise ($1/f$, electromagnetic interference, microphonic, to name a few) is almost always present on transducer output, buttressing the claim that noise will be present at some level on any transducer signal. Finally, just amplifying the composite signal won't help make the signal of interest more distinguishable from noise, since amplifiers boost the level of everything present at the input (noise included), and contribute noise of their own to boot! (For a discussion of the various kinds of electrical noise see Horowitz and Hill, pp. 430-436, as noted in the references at the end of this section).

However, if the experiment in question is one in which measurements are made of the response to a controlled excitation, there may be a way out of the noise quagmire. An example of such an experiment is the measurement of the resistance of a circuit element: apply a known current and measure the corresponding voltage drop. Another example is an optics experiment involving the measurement of the ratio of scattered to incident light intensity. In such situations the experimenter may be able to impart some unique characteristic to the excitation and then measure the response by a method that strongly selects for that characteristic. For instance, in the resistance example, one could use an AC current of a particular frequency f_m and then measure the AC voltage with an instrument tuned to f_m.

The obvious way to accomplish this tuned measurement would be to first modulate the excitation with an AC signal at f_m, place a narrow bandpass filter also tuned to f_m in the response-signal path to filter out all signals but the ones at f_m, and then rectify the result. In many cases the improvement in noise rejection will be sufficient. But what if you want to change the frequency of the excitation? And what if you would like to vary the bandpass width? More sophisticated filter designs might be needed. Moreover, in practice it is difficult to make an extremely narrow bandwidth filter. The bandwidth of a filter may be specified by its Q value, which is defined as the ratio of the center frequency f_0 to the range of frequencies Δf between the points where the filter response falls by one-half: $Q = f_0/\Delta f$. Practically, it is hard to build and use filters with $Q >\approx 50$.

The phase sensitive detector makes an elegant end run around these problems by reversing the order: it first rectifies the signal and then filters it. We'll look at how this is done, and then at why this works so well. Figure 1 shows the basic phase sensitive detector or PSD, greatly simplified. The PSD is composed of two basic circuits: a synchronous switch and a amplifier/filter.

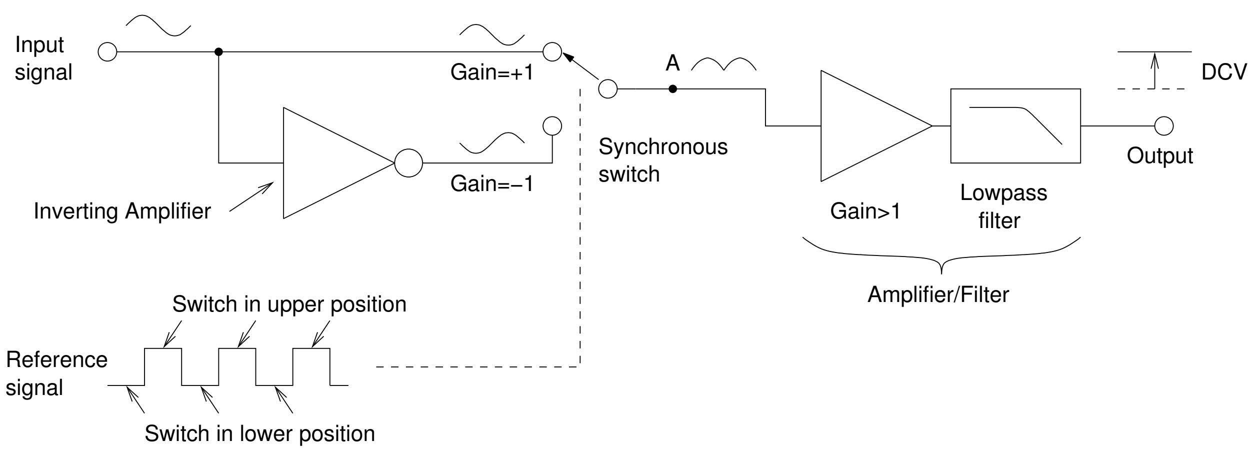

Figure 1: Simplified schematic of the phase sensitive detector.

Consider, first, the action of the switch. It is designed so that it spends an equal amount of time in each position corresponding to the period of a reference signal $T_m = 1/f_m$. In the upper position, it passes the signal unchanged; in the lower position it passes the signal inverted. If we apply a sine wave to the input which has the same frequency as the reference, $V(t) = V_0 \sin(2\pi f_m t + \phi)$, the signal at the output of the switch (at point A in the figure) will depend upon the phase angle ϕ between the reference signal and the input signal.

If the phase angle $\phi = 0$, the two signals are in phase, and the result at point A looks like the lower trace of Fig. 2a: the sine wave has been rectified. If the result is then passed to the amplifier/filter stage, with a time constant τ of the filter sufficiently long, the output signal will be a constant (DC) signal whose value will be proportional to the amplitude V_0. However, if $\phi = 90°$, as in Fig. 2b, the signal at point A will be symmetric about the zero-volt axis, and the output from the amplifier/filter would be zero. Likewise, it is easy to see that if $\phi = 180°$, the output would be proportional to $-V_0$. This dependence on the phase between the input and the reference signals is why this method is called "phase sensitive detection".

Qualitatively, it is now easy to see why this method helps to significantly reduce noise. If one were to apply a signal to the input at a different frequency f than the reference frequency f_m, the phase ϕ would be constantly changing, and the net effect would be to produce a signal that would average to zero.

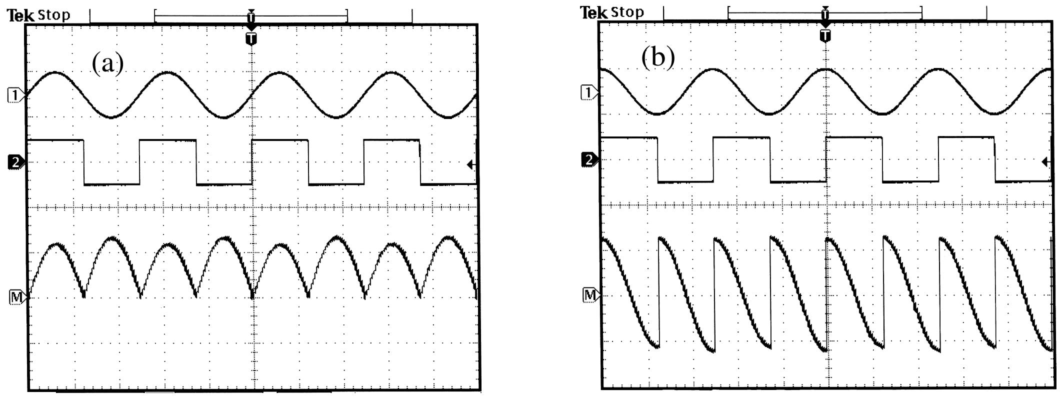

Figure 2: Oscilloscope traces simulating the effect of the synchronous switch section of the phase sensitive detector. (a) Reference (square wave) and signal (sine wave) in phase produce a net positive waveform; (b) Reference and signal 90° out of phase produce a waveform with a zero time average.