electricity you see demonstrations of glass rods being rubbed with cat's fur.[136] With some charge stuck on the glass, the electric field that it produced caused the electrodes of the reed switch to acquire a small amount of opposing charge when separated.

Here's what we did to confirm this conjecture: we took an inch of woven metal cloth,[137] grounded one end with a clip lead, wrapped it around the reed switch, and slid it along the length several times, to provide an opportunity for any static charge on the glass surface to take a hike. And it did: after one treatment the output-voltage step was reduced to a negligible ~2 mV.

The moral: sensitive measurements (here we're talking femtoamps and picofarads) can reveal effects that are so small that you've never thought about them. They can seriously disrupt your work... but there's a redeeming delight in discovering them for yourself. And then eliminating them.

8.13.8 Noise potpourri

Herewith a collection of interesting, and possibly useful, facts.

1. The averaging time required in an indicating device to reduce the fluctuations of a rectified noise signal to a desired level for a given noise bandwidth is

$$\tau \approx \frac{1600}{B\sigma^2} \text{ seconds,} \qquad (8.61)$$

where τ is the required time constant of the indicating device to produce fluctuations of standard deviation σ percent at the output of a linear detector whose input is noise of bandwidth B.

2. For band-limited white noise, the expected number of maxima per second is

$$N = \sqrt{\frac{3(f_2^5 - f_1^5)}{5(f_2^3 - f_1^3)}}, \qquad (8.62)$$

where f_1 and f_2 are the lower and upper band limits. For $f_1 = 0$, $N = 0.77f_2$; for narrowband noise ($f_1 \approx f_2$), $N \approx (f_1 + f_2)/2$.

3. rms-to-average (i.e., average magnitude) ratios:

Gaussian noise: $\text{rms/avg} = \sqrt{\pi/2} = 1.25 = 1.96\,\text{dB}$,

sinewave: $\text{rms/avg} = \pi/2^{\frac{3}{2}} = 1.11 = 0.91\,\text{dB}$,

square wave: $\text{rms/avg} = 1 = 0\,\text{dB}$.

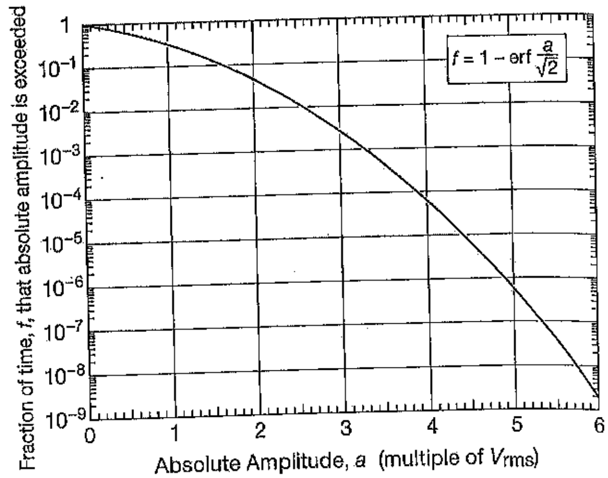

Figure 8.115. Relative occurrence of amplitudes in Gaussian noise. Potentially useful for estimating false trigger rates, required "crest factor" in rms measurements, and the like.

4. Relative occurrence of amplitudes in Gaussian noise. Figure 8.115 plots the fractional time that a given amplitude level is exceeded by a Gaussian noise waveform of unit rms amplitude.

5. Positive-threshold crossing rate of lowpass-filtered Gaussian white noise of unit rms amplitude is

$$\text{TCR} = \frac{\text{BW}}{\sqrt{3}} \exp\left(-V_{\text{th}}^2/2\right) \text{ crossings/second,} \qquad (8.63)$$

where V_{th} is the positive threshold voltage and BW is the brick-wall lowpass bandwidth.[138]

6. Standard deviation of noise resulting from quantization error is

$$\sigma_n = \frac{\text{LSB}}{\sqrt{12}} \approx 0.3\,\text{LSB}. \qquad (8.64)$$

8.14 Signal-to-noise improvement by bandwidth narrowing

As luck would have it, the signals you often want to measure are buried in noise (where the "noise" may include other signals nearby in frequency, i.e., interference), frequently to the extent that you can't even see them on an

136 Or is it a cat being rubbed with glass wool?

137 Wrapped around some foam, this handy stuff is used to make flexible conductive shielding gaskets. Check out the self-stick "fabric-over-foam" gasket materials from Laird Technologies, e.g., their rectangular 4046 or D-shaped 4283.

138 The crossing rate through a pair of symmetrical thresholds (i.e., the rate of crossings of *magnitude* greater than V_{th}) is double that obtained from the formula. We thank Phil Hobbs for this fact, found along with many others in his fine book referenced on page 551.

oscilloscope. Even when external noise isn't a problem, the statistics of the signal itself may make detection difficult, as, for example, when counting nuclear disintegrations from a weak source, with only a few counts detected per minute. Finally, even when the signal is detectable, you may wish to improve the detected signal strength in order to make a more accurate measurement. In all these cases some tricks are needed to improve the signal-to-noise ratio. They all amount to a narrowing of the detection bandwidth in order to preserve the desired signal while reducing the total amount of (broadband) noise accepted.

The first thing you might be tempted to try when thinking of reducing the bandwidth of a measurement is to hang a simple lowpass filter on the output, in order to average out the noise. There are cases where that therapy will work, but most of the time it will do very little good, for a couple of reasons. First, the signal itself may have some high frequencies in it, or it may be centered at some high frequency. Second, even if the signal is in fact slowly varying or static, you invariably have to contend with the reality that the density of noise power usually has a $1/f$ character, so as you squeeze the bandwidth down toward dc you gain very little. Electronic and physical systems are twitchy, so to speak.

In practice, there are a few basic techniques of bandwidth narrowing that are in widespread use. They go under names like signal averaging, transient averaging, boxcar integration, multichannel scaling, pulse-height analysis, lock-in detection, and phase-sensitive detection. All of these methods assume that you have a repetitive[139] signal; that's no real problem, since there is almost always a way to force the signal to be periodic, assuming it isn't already. Here we discuss an important one of these techniques, known as "lock-in" or "synchronous" detection.

8.14.1 Lock-in detection

This is a method of considerable subtlety. It consists of two steps. (1) Some parameter of the source signal is *modulated*; for example, an LED might be driven with a square wave at a fixed frequency. (2) The detected (and noisy) signal is *demodulated*, for example by multiplying it by a fixed-amplitude reference signal at the same modulating frequency. The modulation moves the source signal spectrum up to the modulating frequency, above noisy $1/f$ low-frequency backgrounds, and away from other noise sources (such as ambient-light fluctuations in the case of the LED example). The demodulation step creates a dc output proportional to the signal, which can be lowpass filtered (a simple RC filter may be adequate) to narrow the detected bandwidth.

To understand the method, it is necessary to take a short detour into the phase detector, a subject we first take up in §13.13.2.

A. Phase detectors

In §13.13.2 we describe phase detectors that produce an output voltage proportional to the phase difference between two digital (logic-level) signals. For purposes of lock-in detection, you need to know about *linear* phase detectors, because you are nearly always dealing with analog voltage levels.

The simplest circuit[140] is shown in Figure 8.116. An analog signal passes through a linear amplifier whose gain is reversed by a square-wave "reference" signal controlling a FET switch (see Table 3.3 or 13.7 for candidates). The output signal passes through a lowpass filter, RC. That's all there is to it. Let's see what you can do with it.

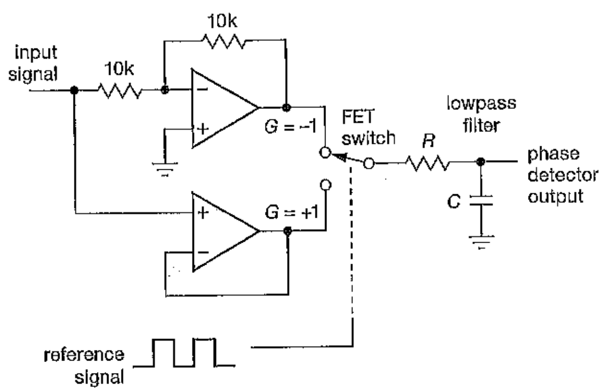

Figure 8.116. Phase detector for linear input signals. Most simply, you can implement this with a dual op-amp and a CMOS switch IC. This scheme is used in the monolithic AD630.

Phase-detector output

To analyze the phase-detector operation, let's assume we apply a signal

$$E_s \cos(\omega t + \phi)$$

to such a phase detector, whose reference signal is a square wave with transitions at the zeros of $\sin \omega t$, i.e., at

[139] Or, more generally, a *known* signal variation to which the measured signal can be correlated.

[140] But less than ideal: the square-wave modulation causes response at odd harmonics. The use of an analog multiplier like the AD633 or AD734, driven by a sinewave reference, eliminates this deficiency.

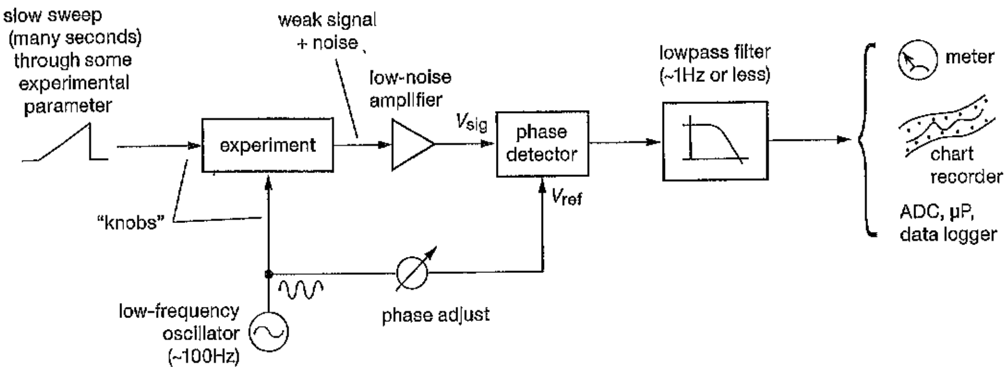

Figure 8.117. Lock-in amplifier detection.

$t = 0, \pi/\omega, 2\pi/\omega$, etc. Let us further assume that we average the output, V_{out}, by passing it through a lowpass filter whose time constant is much longer than one period:

$$\tau = RC \gg T = 2\pi/\omega.$$

Then the lowpass filter output is

$$\frac{1}{2}\langle E_s \cos(\omega t + \phi)\rangle|_0^{\pi/\omega} - \langle E_s \cos(\omega t + \phi)\rangle|_{\pi/\omega}^{2\pi/\omega},$$

where the angle brackets represent averages and the minus sign comes from the gain reversal over alternate half-cycles of V_{ref}. As an exercise, you can show that

$$\langle V_{out}\rangle = -(2E_s/\pi)\sin\phi.$$

Exercise 8.9. Perform the indicated averages by explicit integration to obtain the preceding result for unity gain.

Our result shows that the averaged output, *for an input signal of the same frequency as the reference signal*, is proportional to the amplitude of E_s and sinusoidal in the relative phase.

We need one more result before going on: what is the output voltage for an input signal whose frequency is close to (but not equal to) the reference signal? This is easy, because in the preceding equations the quantity ϕ now varies slowly, at the difference frequency:

$$\cos(\omega + \Delta\omega)t = \cos(\omega t + \phi) \quad \text{with } \phi = t\Delta\omega,$$

giving an output signal that is a slow sinusoid:

$$V_{out} = (2E_s/\pi)\sin(t\Delta\omega),$$

which will pass through the lowpass filter relatively unscathed if $\Delta\omega < 1/\tau = 1/RC$ and will be heavily attenuated if $\Delta\omega > 1/\tau$.

B. The lock-in method

Now the so-called lock-in (or phase-sensitive) amplifier should make sense. First you make a weak signal periodic, as we've discussed, say at a frequency in the neighborhood of 100 Hz. The weak signal, contaminated by noise, is amplified and phase detected relative to the modulating signal. Look at Figure 8.117. In many cases you'll want to measure the weak signal as some experimental condition is varied – you'll have an experiment with two "knobs" on it, one for fast modulation to do phase detection, and one for a slow sweep through the interesting features of the signal (in NMR, for example, the fast modulation might be a small 100 Hz modulation of the magnetic field, and the slow modulation might be a frequency sweep 10 minutes in duration through the resonance). The phase shifter is adjusted to give maximum output signal, and the lowpass filter is set for a time constant long enough to give a good signal-to-noise ratio. The lowpass-filter rolloff sets the bandwidth, so a 1 Hz rolloff, for example, gives you sensitivity to spurious signals and noise only within 1 Hz of the desired signal. The bandwidth also determines how fast you can adjust the "slow modulation," because now you must not sweep through any features of the signal faster than the filter can respond; people use time constants of fractions of a second up to tens of seconds or more.[141]

Note that lock-in detection amounts to *bandwidth narrowing*, with the bandwidth set by the post-detection lowpass filter. Another way to reduce the detection bandwidth is with the technique of *signal averaging*, in which the results of repetitive measurements (e.g., frequency sweeps) are accumulated; this is a common option on instruments such as spectrum analyzers. In either case the effect of the modulation is to center the signal at the fast modulation

141 In the old days the slow modulation was done with a geared-down clock motor turning an actual knob on something.

frequency, rather than at dc, in order to get away from $1/f$ noise (flicker noise, drifts, and the like).

C. Two methods of "fast modulation"

There are several ways to do the fast modulation: the modulation waveform can be either a very small sinewave or a very large square wave compared with the features of the sought-after signal (line shape versus magnetic field, for example, in NMR), as sketched in Figure 8.118. In the first case the output signal from the phase-sensitive detector is proportional to the *slope* of the line shape (i.e., its derivative), whereas in the second case it is proportional to the line shape itself (providing there aren't any other lines out at the other endpoint of the modulation waveform). This is the reason all those simple NMR resonance lines come out looking like dispersion curves (Figure 8.119).

For large-shift square-wave modulation there's a clever method for suppressing modulation feedthrough, in cases where that is a problem. Figure 8.120 shows the modulation waveform. The offsets above and below the central value kill the signal, causing an on–off modulation of the signal at *twice* the fundamental of the modulating waveform. This is a method for use in special cases only; don't get carried away by the beauty of it all!

Large-amplitude square-wave modulation is a favorite with those dealing in infrared astronomy, where the telescope's secondary mirrors are rocked to switch the image back and forth on an infrared source. It is also popular in radio astronomy, where it's called a Dicke switch.

Commercial lock-in amplifiers have a variable-frequency modulating source and tracking filter, a

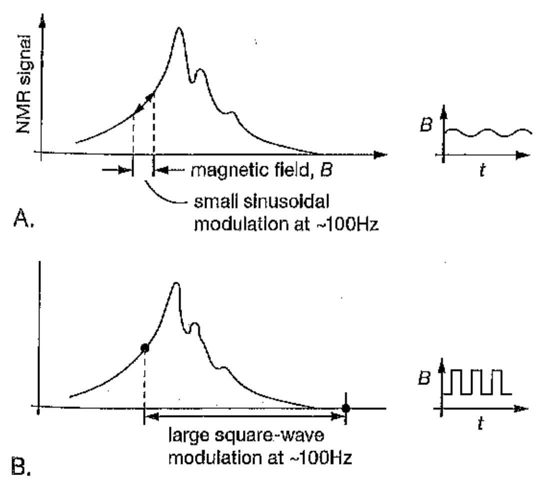

Figure 8.118. Lock-in modulation methods. A. Small sinusoid. B. Large square wave.

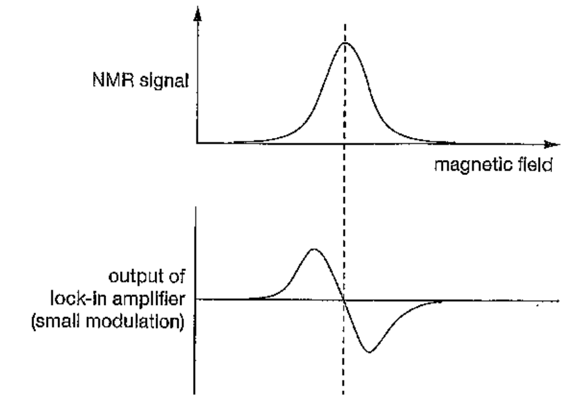

Figure 8.119. Line-shape differentiation resulting from lock-in detection.

switchable-time-constant post-detection filter, a good low-noise wide-dynamic-range amplifier (you wouldn't be using lock-in detection if you weren't having noise problems), and a nice linear phase detector. They also let you use an external source of modulation. The phase shift is adjustable, so you can maximize the detected signal. The whole item comes packed in a handsome cabinet, with a meter or digital display showing the output signal. Typically these things cost a few thousand dollars. Stanford Research Systems has a nice selection of lock-in amplifiers, including several that use digital signal-processing methods for enhanced linearity and dynamic range. In these the amplified input signal is accurately digitized (to 20 bits), the "oscillator" is a computed lookup table of (quadrature) sines and cosines, and the "mixer" is a numeric multiplier. Ordinarily lock-in amplifiers are of rather limited signal bandwidth, typically 100 kHz. But by using the radiofrequency "heterodyne" technique (input frequency band translation via linear mixing with a "local oscillator"), the lock-in method can be extended into high radio-frequencies. For example, the SR844 goes to 200 MHz; it uses a hybrid of analog techniques (input filtering and downconversion) and digital signal processing (baseband digitizing and synchronous detection).

To illustrate the power of lock-in detection, we set up a small demonstration for our students. We use a lock-in to modulate a small LED of the kind used for panel

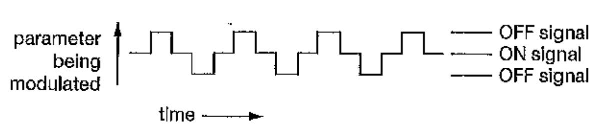

Figure 8.120. Modulation scheme for suppressing modulation feedthrough.

indicators, with a modulation rate of a kilohertz or so. The current is very low, and you can hardly see the LED glowing in normal room light. Six feet away a phototransistor looks in the general direction of the LED, with its output fed to the lock-in. With the room lights out, there's a tiny signal from the phototransistor at the modulating frequency (mixed with plenty of noise), and the lock-in easily detects it, using a time constant of a few seconds. Then we turn the room lights on, at which point the signal from the phototransistor becomes just a huge messy waveform, jumping in amplitude by 50 dB or more. The situation looks hopeless on the oscilloscope, but the lock-in just sits there, unperturbed, calmly detecting the same LED signal at the same level. You can check that it's really working by sticking your hand in between the LED and the detector. It's darned impressive.

At the other end of the cost spectrum, synchronous detection is used to accomplish the same rejection of ambient light in some inexpensive light-beam-detection components, for example the S6809/46 and S6986 from Hamamatsu. These ICs come in a clear plastic case (several available package styles) containing an integrated photodiode, preamp, and synchronous detector with logic-level output; included also is the internal oscillator and output driver for the external LED light source. They cost about $6 in small quantities.

8.15 Power-supply noise

Amplifier circuits that do not have a high degree of power-supply rejection are susceptible to noise (and signals) on the dc power supplies. If the dc rails are noisy, the output will be too, so you've got to keep them quiet. The problem is not as bad as it could be, though, because supply noise appears unamplified at the output, to be compared with the amplifier's signal gain of, say, $\times 100$. Still, dc power supplies are rarely quiet even to the level of $100\,\text{nV}/\sqrt{\text{Hz}}$ (100 times a reasonable input noise target of, say, $1\,\text{nV}/\sqrt{\text{Hz}}$). That is why the dc rails in many of the circuits in this chapter are marked "quiet."

How noisy are typical dc bench power supplies? You'll see specs like "0.2 mVrms, 2 mVpp," which sounds respectable enough until you realize that signal levels in a sensitive circuit may be far smaller. For example, the output noise level over the 20 kHz audio band of the preamp of Figure 8.42 is just $1\,\mu$Vrms ($v_{n(out)}=G_V e_{n(in)}\sqrt{\Delta f}$). So such a supply's specified noise is 46 dB above the output noise floor.

Specs are one thing, actual performance is another. To get a measure of the power supply noise scene we measured two dozen dc supplies from our lab's collection, with the resulting spectra of Figure 8.123.[142]

The spread in noise performance is stunning – the outstanding performer (#4) turned out to be a "precision power source" half a century old (we bought it in 1967), with an ovenized compartment for the zener reference and (discrete) error amplifier. By comparison, a contemporary bench supply with elegant digital readout like #12 is nearly 100 times (40 dB) as noisy. Not to be outdone, the real screamers turn out to be a simple switchmode cellphone charger (#22), and an unregulated wall-wart (#23) whose powerline-tracking output wanderings dominate the low-frequency spectrum. At the other end, nothing beats a lead-acid battery (#3, right down at our measurement noise floor) for the utmost in inherently quiet dc sources.[143]

What accounts for these large differences? Switching supplies are inherently noisy, of course. But even among the linear supplies there is a $100\times$ spread (40 dB) in noise voltage. A quiet dc regulated supply must have plenty of loop gain, implemented with low-noise amplifiers. And it's critically important to select a low-noise voltage reference (particularly down at low frequencies, where it can't be quieted with filters); see the discussion and plots in §9.10.

8.15.1 Capacitance multiplier

A nice trick for cleaning up a noisy supply is a "capacitance multiplier" circuit (Figure 8.121). We introduced this back in §8.5.9, where we explored the properties of a BJT input stage with $e_n{\sim}0.07\,\text{nV}/\sqrt{\text{Hz}}$. Here we choose R small enough so that there is at most a volt or so drop at maximum load current, then choose C for a long-enough time constant RC to attenuate adequately the portion of the noise spectrum you care about. For our ribbon-mic amplifier we elaborated on this a bit, with a 2-stage RC filter of ~2 s time constant; it's shown in the top portion of the BJT noise-measuring circuit of Figure 8.92. There we increased RC until the output noise spectrum was down to the analyzer's noise floor.

Note that the capacitance multiplier compromises the dc output regulation – there's no feedback from its output pin. However, this technique is especially effective when

[142] These supplies were tested "as is," with no effort to confirm their operation within original specifications. The reader is cautioned not to rely on these data when making purchasing decisions.

[143] The battery was being trickle charged at the same current as its load for curve #3. If the battery is being discharged (without replenishment) by the load, the slight downward voltage "tilt" appears as a low-frequency excess, seen here in curve #2.