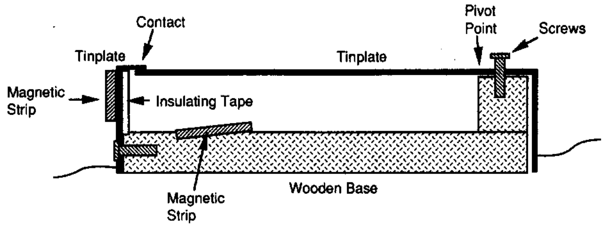

Fig. 2. Constructional details for the arrival sensor.

minate final operation. The design shown in Fig. 2 was found to be satisfactory in this regard. The switch contacts are normally held closed by the weak magnetic forces induced by the short strip of magnetic material; this is of the type used to clamp notices to steel surfaces. The impact of the ball opens the contacts as the large plate pivots about the edge of the wooden support. The securing screws are a loose fit and act only to retain the plate after impact. The second magnetic strip on the wooden base traps the pivoting plate after the switch opens and eliminates any possibility of switch bounce. The plastic insulating tape prevents the ball from subsequently shorting the contacts if it finally settles within the sensor unit after being dropped.

The timer was compared with a precision counter timer, and it was found that time intervals were systematically overestimated by 7.5 ms when an average over many readings was taken, probably due to the interface pulsewidth and the lack of synchronization between the external events and the timer's internal clock. The combined timer/sensor system was tested in a free-fall experiment for heights x between 5 and 50 cm, taking the average of five measurements for each height. The variation in measured time was rarely more than 0.01 s. The results, after the systematic error had been accounted for, are plotted in Fig. 3, together with a one-parameter least-squares fit of the relation $x = 0.5gt^2$. The results confirm the quadratic variation and give $g = 10.0 \pm 0.4$ m/s^2. If required, a more precise determination of g may be made using a greater range of heights; for example, over 3 m, a value of 9.77 m/s^2 was obtained.

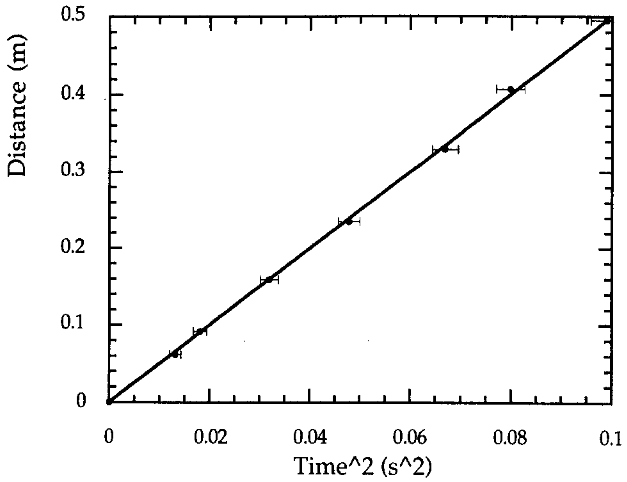

Fig. 3. Distance fallen versus time of flight (squared) for a ball-bearing projectile released from rest.

In conclusion, the free-fall timing apparatus described here is very low in cost, may be constructed with the resources of most teaching laboratories, and has adequate accuracy for the purpose of a teaching experiment. The unit requires no external power, and the timer/stopwatch may be detached and used for other timing purposes.

The lock-in amplifier: A student experiment

Richard Wolfson
Department of Physics, Middlebury College, Middlebury, Vermont 05753

(Received 18 December 1990; accepted for publication 15 January 1991)

The lock-in amplifier is an exceptionally useful instrument for extracting signals from noise, even when noise and signal frequencies are close. Lock-in amplifiers are widely used in physics research, and for this reason it is important that undergraduate students be exposed to the device. The principles of lock-in amplification are covered in common electronics textbooks,[1–4] and several of these offer fairly detailed circuit diagrams. Some years ago this Journal published a description of an inexpensive but rather complex (nine operational amplifiers) lock-in circuit that could be built for student use.[5] But neither the textbook descriptions of lock-in amplifiers nor the full-blown lock-in circuit published earlier provide students with simple experimental insight into the basic workings of this important instrument. This report describes a simple laboratory experiment, easily completed in 2–3 h by students in a sophomore-level electronic course, that illustrates the basic principle of lock-in amplification.

At the heart of a lock-in amplifier is a phase-sensitive detector, a circuit that gives preferential treatment to the desired signal based on information about the signal's phase. In typical applications, a signal—usually optical or electrical—is "tagged" by chopping or otherwise modulating it at a frequency much higher than that of any changes the signal itself might undergo. The signal is then put to use in some experiment, as, for example, when a light beam is used to probe the optical properties of a material. The signal may emerge from the experiment literally buried in

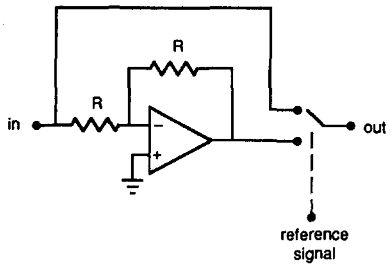

Fig. 1. A simple phase-sensitive detector consists of a unity-gain inverting amplifier and a switch that selects alternately between noninverted and inverted signals. Switching is controlled by the same reference signal that was used to chop the signal before it became mixed with noise.

noise. But a phase-sensitive detector locked to the modulation that was applied *before* the signal was overwhelmed by noise can extract the signal in what would otherwise appear a hopeless situation.

A variety of schemes are available for modulation and subsequent phase-sensitive detection. In optical systems, a light beam is often chopped by a rotating shutter; alternately, an electrically pulsed source or shutter can be used for modulation. The subsequent phase-sensitive detection also admits a variety of implementations, including analog multiplication and heterodyning. A particularly simple phase detector, shown in Fig. 1, alternately samples the noisy signal in its original and in its inverted form using an analog switch at the output of a unity-gain inverting amplifier. Any input signal components that are not phase locked to the driver of these alternations will, over time, be equally likely to be sampled in the inverted as in the noninverted state. Over many cycles of the alternation, these components will average to zero. But an input component that is phase locked with the alternation will be sampled in only one of the two states, so its time average will not vanish. A low-pass filter provides the requisite time averaging. The phase-sensitive detector and filter will then extract the signal from larger amplitude noise—even if the noise frequency is close to but not phase locked to the signal frequency.

Our simple student experiment consists of a phase-sensitive detector of the type shown in Fig. 1, preceded by a summing amplifier that mixes a modulated "signal" with "noise." For pedagogical reasons, we use a pure sine wave for the noise; that way the students can see more clearly exactly what happens to both signal and noise as they progress through the circuit. For further clarity, we give the signal a different shape from the noise; a triangle waveform proves appropriate.

Lock-in amplifiers are most commonly used with slowly varying signals; time scales for signal variation are commonly on the order of seconds or longer. An appropriate output device for such a signal is a strip chart recorder, while modulation frequencies in the range from tens to hundreds of hertz are often used. Because students in introductory electronics courses are familiar with and generally have available oscilloscopes rather than strip chart recorders, our lock-in demonstration runs at higher frequency. The signal frequency is taken about as low as can comfortably be viewed on an oscilloscope; 50 Hz is a good choice. The circuit can successfully pull that 50-Hz signal from a much stronger noise signal at 51 Hz or even closer, and toward the end of the experiment students should so demonstrate. But to start, it is clearer if the noise frequency is sufficiently different that a glance at the oscilloscope screen can distinguish the signals by virtue of their different periods. A noise frequency around 70 Hz is therefore reasonable.

With noise and signal frequencies on the order of 100 Hz, an appropriate modulation frequency is in the 10-kHz range. For pedagogical purposes, it is valuable to have students study the early stages of the circuit with a modulation frequency around 700 Hz; that way they can see clearly what the modulation is doing. Switching to the higher frequency is necessary for optimum performance when examining the final output of the device. The modulation is achieved with an analog switch that swings between the signal and ground; the output of this modulator is then mixed with the noise to provide a noisy signal for input to the phase-sensitive detector. The circuit can easily extract the signal when the noise is 5 times stronger; we therefore use an input signal of 0.5-V peak and a noise of 2.5-V peak, for a signal-to-noise ratio of 1:5.

Figure 2 shows the full circuit for the lock-in experi-

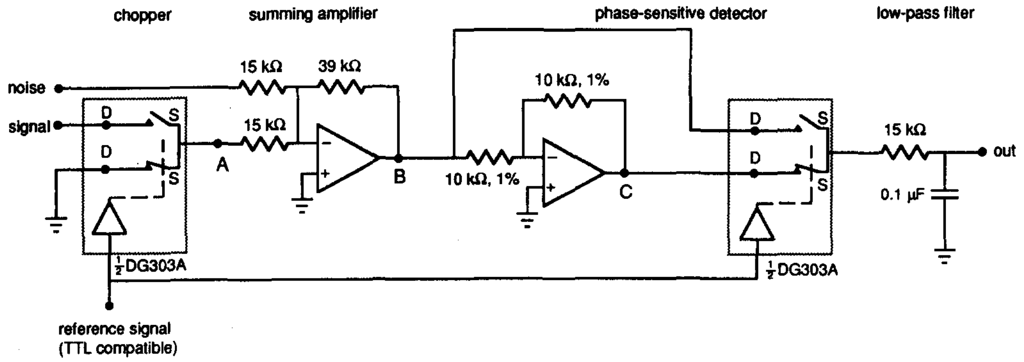

Fig. 2. Circuit for the lock-in experiment. Input signal is a 0.5-V peak, 50-Hz triangle wave; noise is a 2.5-V peak, 70-Hz sine wave. The reference signal is a TTL-compatible square wave at about 10 kHz, although initial study of the early stages of the circuit uses a 700-Hz reference signal for clarity. Gray boxes each contain half of a DG303A analog switch, whose pin connections are shown in Fig. 3. Gain of the summing amplifier is not critical. The operational amplifiers can be 741, 411, or other general-pupose op–amps.

ment. The signal and noise are supplied from standard laboratory function generators, while the TTL-compatible output of a third function generator drives the analog switches that provide modulation and phase-sensitive detection. The DG303A analog switched contain two independently driven pairs of CMOS switches in a single package. One switch in each pair is normally open, the other normally closed, so together they function like a single SPDT switch. Each pair is operated by a TTL-compatible driver. (See Fig. 3 for DG303A details and pin connections.) Component values in Fig. 2, especially the resistors that determine the gain of the summing amplifier, are not crucial. It is important, however, that the resistors in the phase-sensitive detector be closely matched so that the inverting amplifier has a gain of very nearly -1. The components in the low-pass filter are chosen to give as clean an output signal as is possible without significantly degrading the triangular signal waveform. Even more impressive signal-to-noise enhancement could be achieved with heavier filtering, at the expense of fidelity to the input waveform. It is important for students to recognize that the filter would be useless in directly separating the original 70-Hz noise from the 50-Hz signal; what it does, instead, is attenuate the noise signal after the phase-sensitive detector has set it swinging rapidly—at the high modulation frequency—between inverted and noninverted values.

In laboratory the students build the circuit in sections, examining the operation of each section before going on to another. They first set the input and noise signal to their requisite values, and build the summing amplifier to see that the signal does indeed get lost in the noise. (The summing amplifier simulates the effect of a noise-producing experiment environment.) Figure 4 shows the signal and noise, displayed on an oscilloscope at the same scale so the 1:5 signal-to-noise ratio is evident. Next students construct the chopper and insert it between the signal source and the input to the summing amplifier circuit. With the chopping signal set to 700 Hz, the effect of the chopping is clearly evident on an oscilloscope monitoring point A—the chopper output—in the circuit of Fig. 2. Figure 5 shows the voltage at point A as well as at the output of the summing amplifier. Not only is the chopping obvious, but it is also clear that the sinusoidal, 70-Hz noise signal dominates the output of the summing amplifier.

Finally students add the phase-sensitive detector and low-pass filter. Figure 6 shows the detector and filter outputs. The noise signal still seems dominant even after phase-sensitive detection—but it also shows clearly a rapid alternation between positive and negative values. The low-pass filter eliminates that rapid alternation, leaving only the true signal whose sign was not commutated by the phase-sensitive detector. The final output from the filter—the upper trace in Fig. 6—is a resonably faithful reproduction of the original input signal, showing clearly both its 50-Hz frequency and its triangular shape. (The quality of the output signal is fairly sensitive to the modulating frequency; experience suggests that modulating frequencies within a factor of 2 of 10 kHz work best with the circuit and input signal parameters of Fig. 2.)

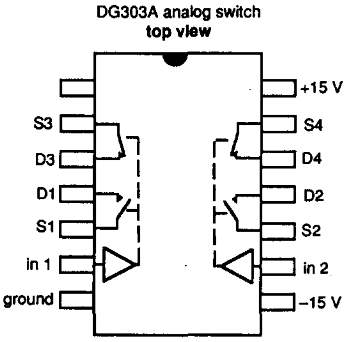

Fig. 3. Pin connection for the DG303A analog switch. This CMOS device has TTL-compatible switch drivers. Each side of the device contains a pair of switches, one normally open and the other normally closed. Connecting their S contacts together gives an SPDT switch.

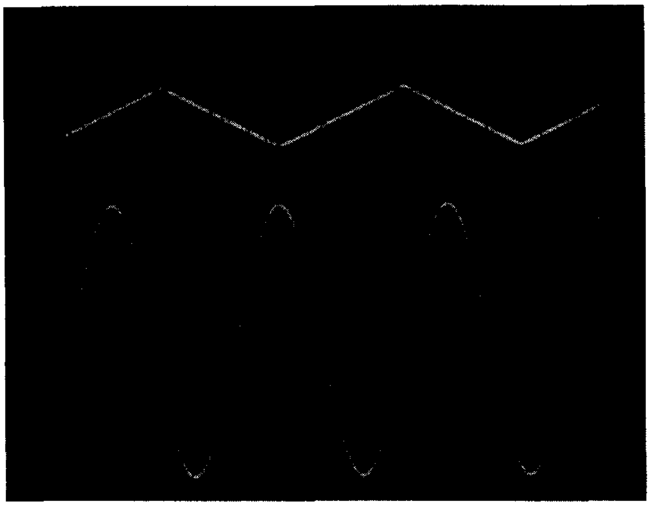

Fig. 4. Signal (upper trace) and "noise" (lower trace), shown at the inputs to the circuit of Fig. 2. Both traces are displayed at 1 V/division, with the sweep rate at 5 ms/division. The signal-to-noise ratio is 1:5.

Having convinced themselves that phase-sensitive detection works, students can next explore the limits within which the circuit is effective. One of the most impressive demonstrations is to move the noise frequency right through the signal frequency, either manually or with a

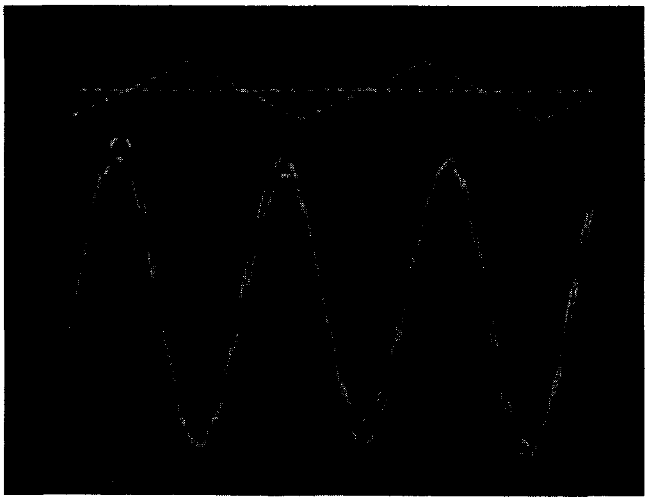

Fig. 5. Upper trace: Output of the chopper, point A in Fig. 2, shows that the input to the summing amplifier alternates ground and the input signal of Fig. 4. Chopping frequency has been set to about 700 Hz to make the alternation evident. Lower trace: Output of the summing amplifier (point B in Fig. 2) is dominated by the 70-Hz sinusoidal "noise," although the chopped signal is also evident. Upper trace is at 1 V/division, lower trace at 2 V/division. Sweep rate is 5 ms/division.

Fig. 6. Lower trace: Output of the phase-sensitive detector (point C in Fig. 2) shows rapid (~10 kHz) alternation between inverted and noninverted waveforms. Not obvious is the fact that only the "noise" is being inverted. Upper trace: Output of the entire circuit, at the junction of the 15-kΩ resistor and the 0.1-μF capacitor that form the low-pass filter. The 50-Hz triangle-wave signal is obviously the dominant feature. Upper trace is at 0.5 V/division; lower trace at 5 V/division. Sweep rate is 5 ms/division.

sweep-frequency generator. Based on differing waveform shapes, the circuit can still extract the signal from noise with five times the amplitude, even when signal and noise frequencies differ by less than 1 Hz. This demonstration should convince students of the importance of the modulation: It's not the noise and signal frequencies that the circuit must separate, but rather the noise and modulation frequencies. Indeed, supplying square-wave noise slightly decreases the effectiveness of the circuit, presumably because of the presence of higher frequency harmonics.

Our circuit demonstrates the principles of phase-sensitive detection, but it hardly constitutes a complete lock-in amplifier. A full lock-in circuit would include a tuned amplifier with adjustable gain, a phase shifter to bring the reference signal (i.e., the modulation) into the optimum phase relation with the signal, a low-pass filter with adjustable time constant, metering, and other options. And our summing amplifier would not be included—mixing of signal and noise really occurs in the experimental apparatus, external to the lock-in amplifier. To help students understand the relation of their laboratory experiment to a "real" lock-in, we have a commercial lock-in amplifier set up in the laboratory, with a noisy input signal and a chart recorder at the output. After they have completed their simple phase-sensitive detection experiment, students learn how to make the connections and control adjustments that allow this "real" lock-in to pull the signal from noise. They leave the laboratory with an appreciation of this instrument, and an eagerness to use it when the need arises in subsequent research.

[1] Paul Horowitz and Winfield Hill, *The Art of Electronics* (Cambridge U. P., New York, 1989), 2nd ed., pp. 1031–1034.

[2] Richard J. Higgins, *Electronics with Digital and Analog Integrated Circuits* (Prentice-Hall, Englewood Cliffs, NJ, 1983), pp. 509–515.

[3] A. James Diefenderfer, *Principles of Electronic Instrumentation* (Saunders, Philadelphia, 1979), pp. 384–387.

[4] Robert E. Simpson, *Introductory Electronics for Scientists and Engineers* (Allyn and Bacon, Boston, 1987), 2nd ed., pp. 344–354.

[5] Paul A. Temple, "An introduction to phase-sensitive amplifiers: An inexpensive student instrument," Am. J. Phys. **43**, 801–807 (1975).