

PC-Based Digital Lock-In Detection of Small Signals in the Presence of Noise

Philip Kromer, Ralph Robinett, Roger Bengtson, Charles Hays²

Department of Physics, University of Texas at Austin

Abstract

Several interesting experiments in the advanced laboratory require an accurate measurement of a slowly varying, extremely small voltage. Lock-in detection is a powerful technique to recover such a signal, even in the presence of broadband noise whose magnitude is several times greater than the signal itself. We have implemented a versatile, low-cost digital lock-in analyzer completely in software. No specialized hardware is required beyond a general-purpose data acquisition board and a low-noise amplifier, yet the detector has a sensitivity of 20 nV and negligible offset drift. Since all signal processing takes place on the computer, students can display the waveform – as a time series or power spectrum – as it progresses through the instrument, which makes it an extremely useful teaching tool. We describe the implementation of the lock-in detector and describe its use in a measurement of the resistance of a superconductor as it undergoes its superconducting transition. Its versatility, resolution, and teaching utility make this tool excellently suited to the advanced laboratory.

Introduction

We have implemented a PC-based digital lock-in analyzer (LIA) completely in software¹. It is both a powerful measurement device and a useful tool for teaching the basics of lock-in detection. We use the LIA to measure the resistance of a superconductor as it undergoes its superconducting transition, but it is a versatile device. It may be applied to any experiments which requires detection of a small (less than a microvolt), slowly varying (stable over several seconds) voltage that is in response to an externally applied voltage:

- The Hall effect in semiconductors¹¹,
- The use of a thermistor bridge to detect small temperature changes¹²,
- Many other examples from physics and engineering^{11,12}.

Measuring small signals

It is difficult to measure such a small signal for several reasons^{5,6,7}. Measurement systems suffer from offset and drift – several effects (common-mode error, thermoelectric offsets, internal offsets, and rectification of noise) will add or subtract to the signal level you measure, and these offsets will vary with temperature and time. These errors may be mitigated but not eliminated.

Small signals will also be contaminated by noise. The experiment should be set up to minimize the effects of noise as far as possible, and to eliminate those noise sources (such as ground loops and electromagnetic interference) which may be eliminated. However, there are other sources of noise (such as Johnson noise, $1/f$ noise, and shot noise) which are due to fundamental processes and may not be eliminated. Since some of these noise sources have a $1/f$ spectrum – the noise amplitude increases as the frequency goes to zero – their effect will be worst at DC.

If, however, we were able to move the signal of interest away from DC – if the quantity appeared as a periodic signal at some other frequency – several benefits would result. Offset and drift would no longer be a concern, as long as they did not introduce a significant common-mode error. We could choose a frequency high enough to avoid $1/f$ noise, yet low enough to be “quasi-static” (to not materially affect the phenomena under study), and which is far from any significant noise source. Since our signal would be stable in frequency and phase, it will be easy to distinguish from noise.

Application: Resistance of a Superconductor

For example, these concerns have consistently plagued our experiment to measure resistance versus temperature for a $Y_1Ba_2Cu_3O_{7-x}$ sample¹⁵ as it undergoes its superconducting transition^{13,14}. We use the four-point configuration shown in Figure 1. A constant current is applied at points **A** and **D**, and the resulting voltage drop appears across **B** and **C**. A “T”-type thermocouple measures the temperature. (It is important to use a “four-point” technique – different contacts to apply the current than to measure the voltage – since the contact resistance is typically several ohms, much more than even the room temperature resistance of the sample).




Figure 1: Setup to Measure Superconductivity

Our sample shows a room temperature resistance near 20 mΩ. An applied current of 10 mA gives a voltage drop, at room temperature, of 200 μV. In order to be convincingly “super-” conducting, we would like to observe the resistance drop by several orders of magnitude. This means we must measure a voltage less than 200 nV in the superconducting state. Unfortunately, the large temperature difference (300 K) between the sample and the amplifier cause thermoelectric offsets, and the student must carefully null the amplifier before beginning the lab. (It is also unconvincing to depend on carefully zeroing the amplifier – after all, this only means that you are as “superconducting” as your shorting terminator.) Furthermore, the amplifier we use gives 1 μV of input noise at DC.

Our solution is to modulate the applied current and use lock-in detection. We apply a 37 Hz, 10 mA current across points **A** and **D** (Figure 1). The resistive voltage drop will show up across **B** and **C**, at 37 Hz and in phase with the applied current. The lock-in analyzer extracts only that component of the signal. It discards offset errors, 1/f noise, and thermoelectric effects; only that portion of the broadband noise near 37 Hz survives, yielding a remarkable improvement in sensitivity.

Overview of Lock-In Detection

A lock-in detector takes a periodic reference signal and a noisy input signal and uses a *phase-sensitive detector* (PSD) to extract only that part of the output signal whose frequency and phase match the reference¹⁰. To see how the phase sensitive detector works, consider a reference signal, S_{ref} , which is a pure sine wave with frequency ω_{ref} :

$$S_{\text{ref}} = \cos(\omega_{\text{ref}} \cdot t),$$

and apply it as the input to the system:

$$S_{\text{app}} = A_{\text{app}} \cos(\omega_{\text{app}} \cdot t).$$

Since the input to the PSD, S_{in} , varies in response to the applied signal, we expect that it will have the same frequency and a constant phase shift to the reference:

$$S_{\text{in}} = A_{\text{in}} \cos(\omega_{\text{ref}} \cdot t + \mathbf{d}_{\text{in}})$$

We may model this noise as uniform broadband noise: that is, the noise consists of sine waves of given amplitude, at every possible frequency and phase.

$$S_{\text{in}} = A_{\text{in}} \cos(\omega_{\text{ref}} \cdot t + \mathbf{d}_{\text{in}}) + \sum_{\omega_{\text{noise}}} A_{\text{noise}} \cos(\omega_{\text{noise}} \cdot t + \mathbf{d}_{\text{noise}})$$

Figure 2 shows, on a log scale, a typical amplitude spectrum for the input signal. One can see the large DC offset and noise component and the generally broadband nature of the noise.

We must extract only the interesting part of the input signal; to do so we rely on a simple trick from trigonometry. If we multiply two sinusoidally varying functions, the resulting signal is equivalent to a sine wave at the sum of the input frequencies, plus a sine wave at the difference of the input frequencies. To see this, recall the sum rules for cosines:

$$\cos(a+b) = \cos(a)\cos(b) - \sin(a)\sin(b)$$

$$\cos(a-b) = \cos(a)\cos(b) + \sin(a)\sin(b)$$

If we add those two equations, the last term will go away and we find

$$\cos(a+b) + \cos(a-b) = 2\cos(a)\cos(b)$$

or equivalently,

$$\cos(a)\cos(b) = \frac{1}{2} [\cos(a+b) + \cos(a-b)]$$

Let us see what happens if we multiply twice our reference signal times our noisy output signal.

$$S_{\text{mult}} = 2 \cdot S_{\text{ref}} S_{\text{in}} = 2 \cdot A_{\text{in}} \cos(\omega_{\text{ref}} \cdot t) \cos(\omega_{\text{ref}} \cdot t + \mathbf{d}_{\text{in}}) \\ + 2 \cdot \cos(\omega_{\text{ref}} \cdot t) \sum_{\omega_{\text{noise}}} A_{\text{noise}} \cos(\omega_{\text{noise}} \cdot t + \mathbf{d}_{\text{noise}})$$




Figure 2: Log Amplitude versus Frequency

$$\begin{aligned}
&= A_{in} \cos(\mathbf{w}_{ref} \cdot t - \mathbf{w}_{ref} \cdot t + \mathbf{d}_m) + A_{in} \cos(\mathbf{w}_{ref} \cdot t + \mathbf{w}_{ref} \cdot t + \mathbf{d}_m) \\
&\quad + \sum_{\mathbf{w}_{noise}} A_{noise} \cos((\mathbf{w}_{ref} + \mathbf{w}_{noise}) \cdot t + \mathbf{d}_{noise}) \\
&\quad + \sum_{\mathbf{w}_{noise}} A_{noise} \cos((\mathbf{w}_{ref} - \mathbf{w}_{noise}) \cdot t - \mathbf{d}_{noise}) \\
&= A_{in} \cos(0 \cdot t + \mathbf{d}_m) + A_{in} \cos(2 \cdot \mathbf{w}_{ref} \cdot t + \mathbf{d}_m) \\
&\quad + \sum_{\mathbf{w}_{noise}} A_{noise} \cos((\mathbf{w}_{ref} + \mathbf{w}_{noise}) \cdot t + \mathbf{d}_{noise}) \\
&\quad + \sum_{\mathbf{w}_{noise}} A_{noise} \cos((\mathbf{w}_{ref} - \mathbf{w}_{noise}) \cdot t - \mathbf{d}_{noise})
\end{aligned}$$

Since the reference and the signal had the same frequency, the difference frequency is at DC (zero frequency). If we low-pass filter the multiplied signal, we will find that only two terms survive, the DC term due to the output of the system, and the component of noise with frequency near the reference signal.

$$S_{filtered} = A_{in} \cos(\mathbf{d}_m) + A_{noise @ ref} \cos(\mathbf{d}_{noise @ ref})$$

If we adjust the phase of the reference so that $\mathbf{d}_m = 0$, then $\cos(\mathbf{d}_m) = 1$ and

$$S_{filtered} = A_{in} + A_{noise @ ref} \cos(\mathbf{d}_{noise @ ref})$$

The amplitude of the filtered signal is the amplitude of our system output, plus a (significantly reduced) noise contribution. The phase of the noise signal will not match the phase of the reference, but rather will vary randomly.

Implementation

Our LIA performs lock-in detection completely in software – we read the time-varying signal, noise and all, into memory using a general-purpose data acquisition (DAQ) card. The input waveform passes through all the subsequent stages of lock-in detection as a digital record. The program handles all of the details, from extracting the resistance and other quantities to graphing and logging that data.

The program is written in National Instruments LabView, version 5.0. LabView is a visual high-level programming language for data acquisition, analysis, and simulation. “Coding” a LabView program involves laying out a data flow diagram. For example, the figure at the right shows the section of our LIA program in which the lock-in is performed.

The only hardware required is a high-gain low noise amplifier¹⁶ and a general-purpose DAQ card¹⁷. The DAQ card must have three analog inputs (temperature, applied current, and voltage) and one analog output (the reference signal) with at least 12 bits of resolution and a 20 kHz sampling rate. We used a National Instruments PCI-MIO-16E-44 multifunction I/O board, which worked excellently, although we recommend a board with 16- or 24-bit resolution for demanding applications.

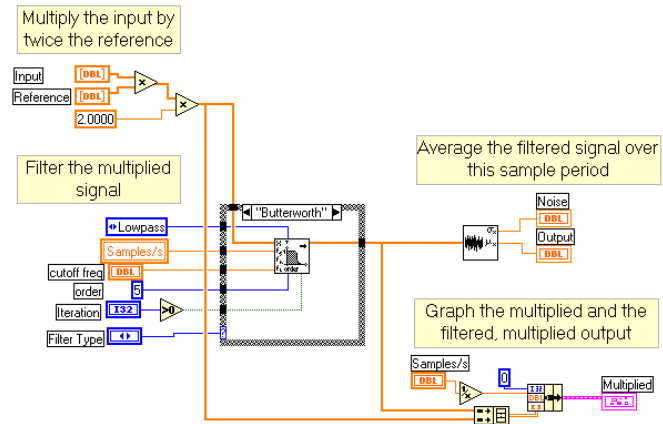


Figure 3: An Excerpt from the LIA Source Code

Since the smallest voltage our board can input is on the order of a millivolt, it is necessary to amplify the input signal. Any low-noise, high gain amplifier will work; we use a homemade two-stage device based around the Burr-Brown OPA111. Once the signal is read in by the DAQ board, no further signal degradation occurs – the signal is now a digital record and is processed exactly. Therefore, the sensitivity of the instrument is limited only by the quality of its input amplifier. We believe that with an amplifier of sufficient quality its sensitivity could rival that of commercial units.

Overview of the Program

Our program begins by initializing the DAQ card – it sets the acquisition parameters, such as the sampling rate and the voltage range of each signal, and performs other housekeeping. The program must also output the reference wave to use as the applied signal. There are many applications where we cannot use the reference to directly generate the applied signal; one must instead use the applied signal to generate the reference. We can not use the applied signal directly – any DC offset

in the reference signal will cause the LIA to pass through a proportional amount of DC noise from the input. (The excellent results our LIA achieves are due in part to the high precision of the reference signal; its only errors are due to quantization and sampling error, which show up as a negligible high-frequency component). For example, in optical systems the applied signal is modulated by mechanically chopping the light beam. In such cases, one uses a phase-locked loop to synchronize the reference with the applied signal. This would greatly increase the complexity of our implementation and so we concentrate on those applications in which we can directly modulate the applied signal.

To generate the applied signal, the program calls a routine that writes the reference waveform into an output buffer and starts its generation. In our implementation we use 1024 points to represent four cycles of the waveform. The DAQ card generates this waveform at an output rate of approximately 5500 Hz, which yields the reference frequency of 37 Hz. From then on until the program ends, the DAQ board continuously generates the waveform – no interaction with the computer is necessary.

The program also sets the data acquisition parameters and begins acquisition. The board continuously samples its inputs at the specified rate and stores the values in a temporary buffer. Our program periodically retrieves the contents of the input buffer and processes that batch of samples as an array. In practice, we trigger each acquisition to occur as the DAQ board begins generating its output waveform buffer, and we acquire an entire output period at a time (the board we use makes this straightforward). Since our acquisition is thus synchronous and continuous, we avoid any windowing errors. (If we did not sample continuously, there would be a discontinuity in the record where we acquired each array; these discontinuities would appear as noise at the frequency which we execute the loop.) Instead, we always read in an exact number of cycles and we process each buffer in real time.

The program next enters the main loop of the program. The program begins the main loop by reading in the applied current, the voltage drop, and the temperature. The voltage drop is passed to a subroutine that performs the phase-sensitive detection (PSD). The PSD routine simply multiplies the input signal by twice the reference, as we argued above. Since the input and reference have the same frequency, the input signal amplitude appears as the zero frequency (DC) component of the multiplied signal. Everything else – including the offensive DC offset and line noise contributions – appears at higher frequencies. If we use a low-pass filter with a large time constant, we will recover only that component which matches the reference in phase and frequency.

The PSD routine next low-pass filters the result using a fifth-order Butterworth filter with a 0.5 Hz cutoff frequency. Changes in the input signal that occur much faster than the output filter time constant will not survive. Thus, the superconducting transition appears to take approximately a second, while of course its true duration is much less. However, as long as the sample warming rate is slow (as long as there is a poor thermal link between the sample and its environment), the transition temperature may be quite accurately determined.

The result is a waveform containing the voltage drop sampled at around 5500 Hz. It is inconvenient to graph or store this much data, and so we take the mean over four cycles (0.108 s) of the voltage drop. We also divide by the nominal applied current to get the resistance.

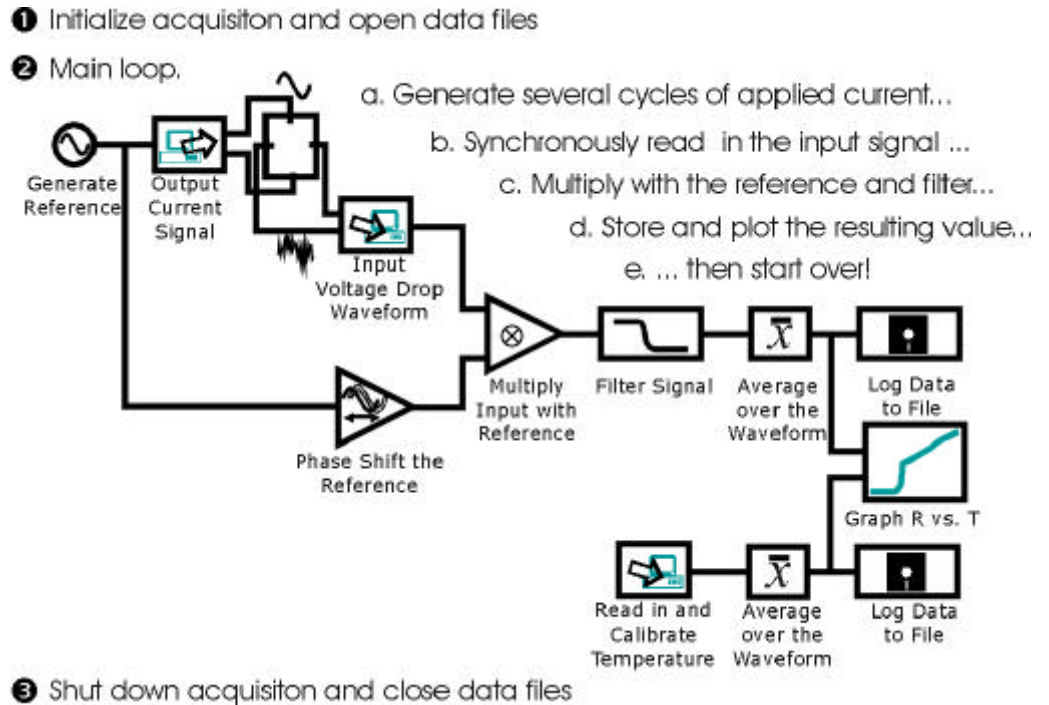


Figure 4: Block Diagram for a Lock-In Detector

The program records this average value, as well as the iteration number, applied current, and temperature, to a tab-delimited file. (It is straightforward to import a tab-delimited file to most spreadsheet applications, but for long data sets, they are slow to generate. In that case we stream to a binary file and write a separate conversion program in LabView.)

Each loop also updates several graphs that display, in real time, the status of the system and the detector. For example, at right is the “airline cockpit” view of the program, used for development and debugging. It shows (clockwise from top right) the input and applied waveforms; their power spectrum; the multiplied and multiplied-filtered signals; the resistivity versus time, versus temperature, and log resistivity versus temperature.

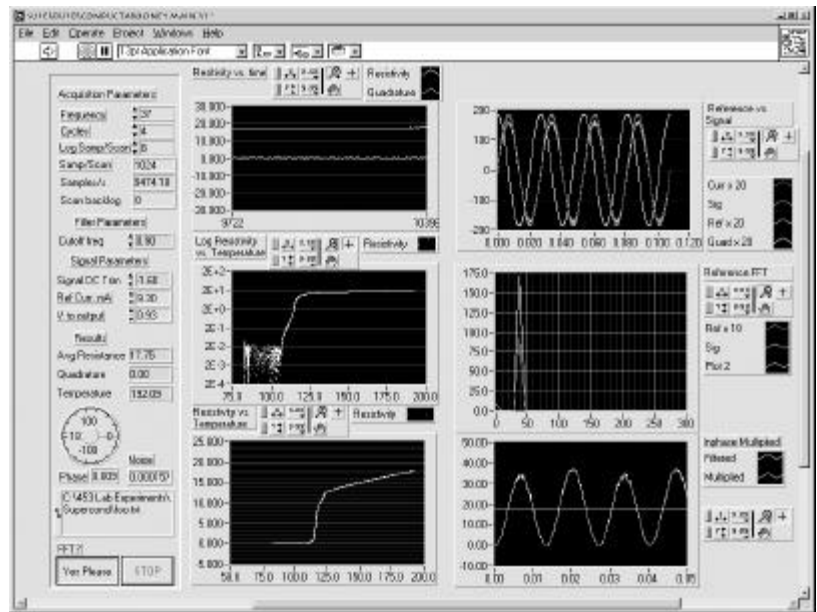


Figure 5: The LIA Front Panel

A student performing this experiment would only need to see the resistivity versus temperature graphs. However, when introducing students to the lock-in analyzer, it is extremely useful to display the waveform as it proceeds through the program, and even more so to display its spectrum. The student can directly observe the effects of broadband or line noise, and can see the LIA filter their contributions from the input. This visualization is impossible to perform with a “black box” commercial unit.

The program then repeats the main loop. This process continues until the user requests the program to stop, at which point the program stops the output generation and the acquisition, performs housekeeping such as closing its data file, and exits.

Results

The main graph in Figure 6 shows a data set recorded as the sample warmed up from approximately 80 to 200 K. The applied current was 10 mA at 37 Hz, and the output filter was set at 0.5 Hz. Each point in the graph is the average over four cycles (0.108 s) of the lock-in output. The inset graph shows the initial portion of the transition on an expanded y-axis.

One can clearly observe the transition at 115 K. The transition is not sharp: this is a general characteristic of high-transition temperature superconductors. Above the transition, we can see that the resistance increases linearly with temperature. Below the transition, we find that the resistance is 0 ± 0.02 microohms. The noise appears to be smaller between 90 and 100 K because the thermoelectric offset disappears, then reappears in this region; this reduces the noise somewhat.

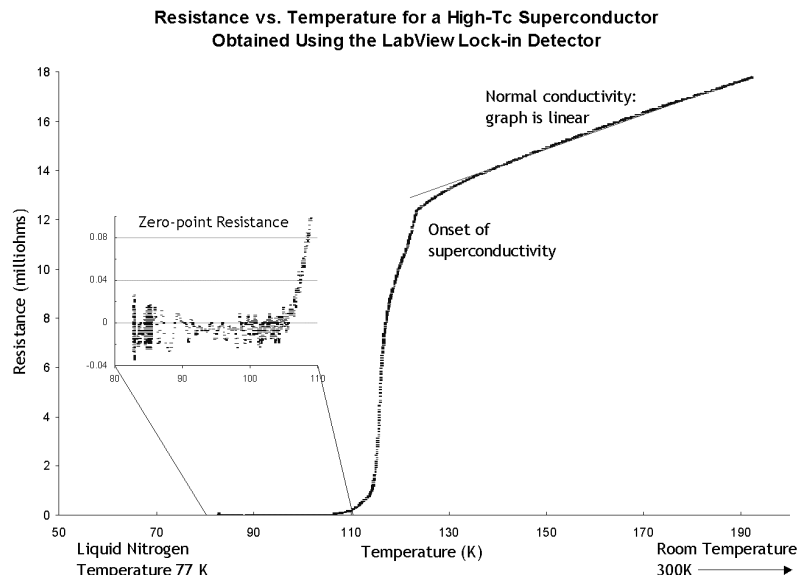


Figure 6: A Sample Resistance versus Temperature Graph

Conclusions

The lock-in analyzer we have described is a versatile tool. Although we have specifically discussed a superconductivity experiment, it would be straightforward to adapt the LIA to any small signal that may be periodically modulated. It is also low-cost in the sense that the equipment required (computer, DAQ board, and amplifier) is readily available and is used for many different experiments. The program not only replaces the lock-in amplifier but also handles the thermocouple calibration and records and charts the data. It is quite sensitive, and its sensitivity can be further improved with an appropriately expensive input amplifier and ADC. Its primary advantage over a commercial unit, however, is as a teaching tool. The ability to display both time series and frequency spectrum displays of the signal – to have an arbitrary number of spectrum analyzers and oscilloscopes connected to the signal path – is of great utility. Furthermore, LabView has powerful simulation capabilities: it can not only implement the LIA but also can simulate the entire experiment. In fact, in order to acquaint our students with LabView and lock-in detection, we give each student a partial simulation to complete before lab. That is, we can not only place a lock-in analyzer on every PC in our lab, we can give each student one to take home! These features make the PC-based digital lock-in analyzer an excellent tool for the junior- or senior-level physics lab.

References

1. More information, including source code for the lock-in detector program and a careful discussion of its limitations, is available at our course home page (<http://wham.ph.utexas.edu/ModPhy/LIA/>) or by contacting the author (flip@physics.utexas.edu).
2. C. Hays is now at the Dept. of Materials Science, California Institute of Technology.
3. We would like to thank Ravi Marawar and National Instruments for software and support, Dr. M.E.L. Oakes for helpful comments on the draft of this paper, and Dr. J. Markert for his advice on the implementation.

Sources of, and defenses against, noise:

4. P. Horowitz and W. Hill, *The Art of Electronics*. Cambridge University Press, New York, 1980.
5. *Low Level Measurements Handbook*, ed. J. Yeager and M.A. Hrusch-Tupta. Keithley Instruments, Cleveland, OH, 1998. An excellent introduction to precision measurement. Freely available on request.
6. "Signal Enhancement" Application Note. Stanford Research Systems, Sunnyvale, CA, 1999. Available at <ftp://ftp.batnet.com/pub/wombats/srsys/ftp/PDFs/Cat98-9.pdf>. A summary of fundamental noise sources.
7. S.J. Shah, "Field Wiring and Noise Considerations," Application Note 25. National Instruments Corp., Austin, TX, 1994. Available at <http://digital.natinst.com/appnotes.nsf/web/index>, #25.

Lock-in detection:

8. M. Stachel, "The Lock-in Amplifier: Exploring Noise Reduction and Phase," <http://www.lockin.de/>. An excellent web-based introduction to lock-in detection, complete with Java simulations.
9. P. Temple, *Am. J. Phys.* **43**(9), 801 (1975).
10. "About Lock-in Amplifiers" Application Note. Stanford Research Systems, Sunnyvale, CA, 1999. Available at http://www.srsys.com/html/application_notes.html. A functional description of lock-in amplifiers.
11. *Lock-in Applications Anthology*, ed. Douglas Malchow. EG&G Princeton Applied Research, Princeton, NJ, 1985. A freely available guide to applications of the lock-in analyzer.
12. D.W. Preston and E.R. Dietz, *The Art of Experimental Physics*. John Wiley & Sons, New York., 1991. Discusses lock-in detection on pp. 367-375.

Superconductivity:

13. G.C. Brown, J.O. Rasure, and W.A. Morrison, *Am. J. Phys.* **57**(12), 1142-1144 (1989).
14. M.J. Pechan and J.A. Horvath, *Am. J. Phys.* **58**(7), 642-644 (1990).
15. Semiconductor samples available from Colorado Superconductor, 1623 Hillside Drive, Fort Collins, CO 80524, <http://www2.csn.net/~donsher/>.
16. Burr-Brown OPA111, Burr-Brown Corp., PO Box 11400, Tucson, AZ 85734, <http://www.burr-brown.com/>.
17. National Instruments PCI-MIO-16E-4 multifunction I/O board. National Instruments, 11500 N. Mopac Expressway, Austin, TX 78759, <http://www.natinst.com/>.