Ph 77 - Advanced Physics Laboratory

Department of Physics, California Institute of Technology

- Electronics Track -

Precision Measurements

This next two-week period consists of two independent one-week laboratory assignments – *Precision Measurements* and *Transmission Lines*. To make sure that equipment is available to all when needed, please do the labs in the order assigned by your Track. Then combine both labs into a single e-notebook to be submitted according to the schedule on Canvas.

Introduction

Progress in experimental physics often involves making ever more precise measurements using ever more sophisticated hardware. This trend has been ongoing for hundreds of years, and we expect it will continue well into the foreseeable future. Sometimes advances are made by measuring extremely small signals, for example doing spectroscopy at the single molecule level or observing astronomical objects out to ever farther distances. Here the experimental challenges are usually associated with improving the signal-to-noise ratio (SNR), so your choices are either increasing the signal (for example, using a larger telescope) or reducing the noise (using a better detector). Obtaining a significant improvement in the SNR is often a reliable path to interesting new discoveries, so this pursuit can be found in all active areas of experimental physics.

Another way to make progress in science and engineering is by measuring known quantities with ever greater precision. One amusing example is that high-end smartphones often include microfabricated barometers. These devices are essentially two flat plates sandwiched together with a small gap between them, forming a capacitor. When the outside air pressure changes, one of the plates flexes a bit and changes the capacitance. Once calibrated, measuring the capacitance yields an absolute measurement of the outside air pressure, turning your phone into a handy barometer/altimeter.

If you download an altimeter app and test it for yourself, you will find that the measurement precision is typically a few tenths of a meter if you average for 10-30 seconds. Assuming an atmospheric scale height of 8.5 km, this means that the barometric pressure is being measured to about 40 ppm. How this translates to capacitance depends on the (proprietary) properties of the MEMS device, but suffice it to say that having a barometer in your phone requires a precision capacitance measurement.

Accuracy and precision

For many precision measurements, it becomes important to distinguish between *accuracy* and *precision*. In the measurement world, accuracy refers to how close a measurement is to the true value of what is being measured, on an absolute scale. Precision refers to the overall measurement sensitivity, which allows you to observe small changes in the measured quantity. Unfortunately, people often use these two terms interchangeably, so "absolute accuracy" drives home the point better than just accuracy. Best to remember the concepts more than the words, because the words are often used incorrectly.

In the barometer example, an accurate measurement can be used to determine your true altitude above sea level (ignoring weather-related changes in pressure). With your smart phone altimeter, you might get an absolute accuracy of 100 meters or so, which is far worse than the 0.3-meter measurement precision. Instruments tend to drift, so measurements with high absolute accuracy are difficult. But the 40 ppm sensitivity of your smartphone barometer means you can watch your elevation change while you are hiking or even climbing stairs.

Cosmic precision

Another instructive example of precision measurements is the Cosmic Microwave Background (CMB), where adding more significant digits revealed a wealth of new physics. When the CMB was discovered in 1964, it was soon recognized as the “3-degree” black-body radiation resulting from Big Bang cosmology. As measurements improved, the mean effective temperature evolved from 3 degrees to 2.7 and then 2.73, eventually reaching its present value of 2.72548±0.00057 K.

Subtracting the mean, observers soon found the CMB dipole moment, indicating that the temperature is 0.0035 K higher in the direction of the constellation Leo and 0.0035 K lower in the opposite direction. This slight temperature shift results from us moving at 369.82±0.11 km/sec toward Leo relative to the CMB radiation, thus establishing what is effectively a rest frame of the Universe.

Subtracting both the mean and dipole moment, there remains some tiny CMB temperature perturbations at the ±0.0002 K level that vary with direction in the sky. Remarkably, cosmological models predict such perturbations, and the detailed angular structure of the CBM can be used to infer the age of the Universe with remarkable accuracy (13.787±0.020 billion years), along with the overall mass-energy composition of the Universe, consisting of 68.2% dark energy (which currently has no established physical interpretation), 26.8% dark matter (probably in the form of some unknown elementary particles), and 5% ordinary matter.

The take-away message in this example is that these remarkable cosmological discoveries were made possible by researchers over many decades developing sophisticated techniques for measuring the black-body radiation temperature with part-per-million accuracy. For this and many other reasons, the general study of precision measurement techniques plays an important role in experimental physics.

Statistical noise & Systematic errors

All physical measurements involve some level of uncertainty, and determining the accuracy of any specific measurement can be a thorny subject. It is customary to report a measurement as $X \pm \sigma_X$, but this notation is overly simplistic in the world of precision measurements. Reporting a simple standard deviation σ_X implies that a measurement is affected by some collection of random perturbations that combine (via the Central Limit Theorem) to produce measurements that are reasonably well described by a Normal (Gaussian) distribution. This is fine for most circumstances, but not all.

If the normal distribution applies, we call that *statistical noise*, which (loosely defined) means that each measurement in an ensemble is independent of the others, and averaging these measurements will improve the overall accuracy. Thus, if one makes N measurements from this distribution, then the error in the mean will be roughly $\sigma_{\langle X \rangle} \approx \sigma_X/\sqrt{N}$, where σ_X is the uncertainty in a single measurement. Importantly, this means that averaging many measurements will reveal the true value of the measured quantity to arbitrary accuracy. Ah, if only life were that simple!

In the real (not mathematical) world, physical measurements are additionally influenced by *systematic errors*, which basically just mean that the measurement is not being done correctly. If the systematic error does not fluctuate around zero, then no amount of averaging will improve the result. Of course, there is no clear divide between fluctuating systematic errors and plain statistical noise, and this is where the data analysis becomes difficult. In many research papers describing precision measurements, results are reported as $X \pm \sigma_{X,statistical} \pm \sigma_{X,systematic}$. In this notation, the known statistical errors could be improved with additional averaging, while the not-precisely-known systematic errors cannot be improved without rebuilding the experiment in some way. Researchers often work for many years to gain a good understanding of the different systematic and statistical noise sources in an experiment.

Types of Noise

Making precision measurements invariably leads to a discussion of noise sources in great detail, and it has proven useful to give names to different types of noise – if you are going to talk about something, it helps to give it a name. Each field has its own experimental challenges, but noise can usually be categorized into one of several types:

Shot noise. This type of noise comes into play when you are counting discrete objects, usually electrons or photons in physical measurements. Shot noise is well described by Poisson statistics, so counting N objects typically implies a standard deviation of about $\sqrt{N}$ (for large N), yielding a signal-to-noise ratio $SNR \approx 1/\sqrt{N}$. This is an important general rule whenever you count independent objects; if you count N of anything, then the measurement uncertainty is roughly $\sqrt{N}$.

Shot noise is especially important when making measurements at high frequencies. When the measurement interval is short, then N is small and shot noise can be an important noise source. If you have enough time, then shot noise can often be reduced to negligible levels simply by increasing N.

Thermal noise. Thermal fluctuations provide another fundamental noise source that is often present in physical measurements. Interestingly, thermal noise is always associated with some form of damping via the *fluctuation-dissipation theorem*. In electronic measurements, resistors produce a thermal noise voltage $\delta V = \sqrt{4kTRB}$, where kT is the Boltzmann energy, R is the electrical resistance in Ohms, and B is the measurement bandwidth in Hertz. The thermal noise voltage is usually expressed in nV/rtHz. (pronounced "nanovolts per root Hertz"). If you short the two ends of a resistor (so δV must equal zero), then thermal fluctuations result in a noise current $\delta I = \sqrt{4kTB/R}$. A 50-Ohm resistor produces a thermal noise of about 1 nV/rtHz.

Quantum noise. This noise source traces back to the Heisenberg uncertainty principle, so you can always find an $\hbar$ somewhere in quantum noise formulations. As one example, gravity-wave interferometers use photons to measure the positions of free test masses, so one wants a large N (number of photons counted in a measurement) to reduce shot noise. However, photons impart momentum that will move the test mass, and this *radiation-pressure noise* increases with increasing N. When you do the math, shot noise dominates at low N while radiation-pressure noise dominates at high N. The combined noise is lowest at some optimal N, and this is the quantum noise limit. It is extremely difficult to even observe the quantum noise limit in any experiment, so quantum effects are negligible in all but the most sophisticated physical measurements.

White noise. Also called Gaussian noise, this type of noise is described by a frequency-independent noise spectrum. White noise is common in electronic amplifiers over a broad frequency range, although the noise must inevitably decrease above some high-frequency cutoff.

"1/f" noise. This type of noise permeates many electronic measurements, identified mainly by the noise spectrum increasing at low frequencies like 1/f, where f is frequency. Oddly, 1/f noise (pronounced "one over f noise") is not well understood physically, and it is known mainly from empirical observations. Electronic instruments tend to drift over long times, and studying this phenomenon in detail reveals a 1/f spectrum in a great many circumstances. Figure 1 shows what 1/f noise looks like in the time domain relative to basic Gaussian noise. Electronic amplifiers are nearly always affected by 1/f noise at frequencies below 10-100 Hz.

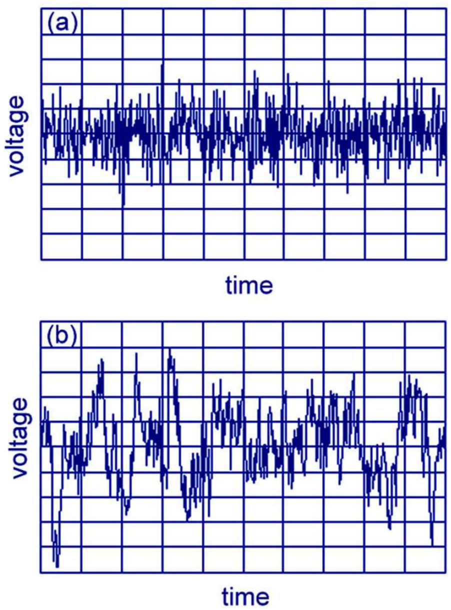

Figure 1. The upper panel (a) shows a signal as a function of time that is dominated by white noise, also called Gaussian noise. White noise typically has a flat spectrum out to some cutoff frequency. The lower panel (b) shows a signal dominated by 1/f noise. At very low frequencies, 1/f noise is essentially equivalent to a signal that slowly drifts over time. Physical measurements are often dominated by white noise at high frequencies and by 1/f noise at low frequencies.

Environmental noise. As the name implies, coupling to a noisy environment can often introduce noise into precision measurements. Examples include RF broadcast signals, 60-Hz or kilo-Hz coupling from power supplies, seismic effects, acoustic effects, etc. The solution is usually to reduce the coupling… often harder than it sounds. Environmental disturbances can be characterized as statistical noise, or systematic error, or somewhere in between.

Technical noise. This is something of a catch-all term for instruments that are not limited by fundamental physical effects. As a general rule, it is quite expensive and time-consuming to build any experiment that is limited by fundamental noise sources. In most laboratory situations, therefore, physical measurements are limited by some form of unwanted technical or environmental noise. This is certainly true for most of the electronics test equipment we use in the Ph77 lab.

Measuring electrical signals and components

Performing cutting-edge measurements in science and technology is often a long game. It may take decades of hard work by many people to produce important scientific results or change the world. Alas, we are not going to do this in the Ph77 lab. Instead, our focus is on teaching measurement strategies using generally inexpensive and robust hardware, so we tend to spend most of our time looking at extremely mundane things like measuring resistors with great precision. For proper motivation, therefore, please imagine that you are measuring something of great interest or importance like the CMB … it just happens to look like a resistor today.

Precision digital multimeters

Most modern measurements come from electronic detectors, so a good starting point in the precision measurement game is to consider some standard pieces of electronics test equipment that give you a lot of significant digits. To this end, our first stop is a look at precision digital multimeters (DMMs) like the one shown in Figure 2.

You are already familiar with basic handheld voltmeters, which are typically "3½ digit" multimeters. What this quirky notation means is that the DMM has a readout of four digits, but the most-significant digit can only have values of 1 or 2 (sometimes 3 or 4). Manufacturers do this because humans tend to focus on measurements near 1, 10, 100, etc., regardless of what is being measured (consider dollars, for example). Because your DMM must stop somewhere, it works out best to stop somewhere between the decade marks.

Depending on your needs (and budget), you can purchase DMMs with 3½, 4½, 5½, 6½, 7½, and 8½ digits, with prices ranging from about $150 to $15,000. In the Ph77 lab, we use the Keysight 34461A shown in Figure 2, which is a 6½-digit model. Having many digits is useful not only for making measurements with high absolute accuracy, but the high precision is also great for examining small time-dependent effects on large signals. For many such reasons, precision DMMs have become a staple in many laboratory situations. There is an old saying that when you have a hammer, everything looks like a nail. In the lab, a precision DMM is a remarkably useful hammer, helping you solve a plethora of day-to-day problems without having to build or program anything.

Figure 2. The Keysight 34461A 6½-digit multimeter, which measures DC voltage, AC voltage, resistance, and capacitance with varying levels of accuracy and precision. The instrument will also display measurement trends and statistics via a broad range of useful software features.

Precision DMMs work by incorporating highly stable and accurate voltage and resistor references that have been developed over many decades. In the U.S., the National Institute of Standards and Technology (NIST) creates and maintains national standards for a wide range of measurement technologies like these, and there are similar government agencies in other countries. The field of *metrology* focuses on defining and improving these standards along many fronts, and such efforts comprise a substantial area of modern research. The details are beyond the scope of Ph77, of course, but metrology and related research areas intersect many STEM career paths.

Fortunately, you can use a precision DMM without knowing what's inside the box. What you do need to know (if you want to use these tools effectively) is the instrument specifications (a.k.a. specs). For the Keysight 34461A, you can find some excerpts from the spec sheet in Appendix I below. As you will see in the lab, the specifications tell you when you can trust all those digits, which is quite important for estimating uncertainties in your measurements. All electronic devices and components come with spec sheets, from the fanciest spectrum analyzer down to the lowliest resistor. If you ever find yourself trying to advance science or technology using precision electronic measurements, then you will soon become quite familiar with reading spec sheets.

A key feature of most precision DMMs is that they are designed to measure DC signals, so they are optimized for this task. An inherent problem with this is that DC signals experience considerable 1/f noise. The noise doesn't go fully to infinity as f → 0, but working at DC is often disadvantageous in the precision-measurement game.

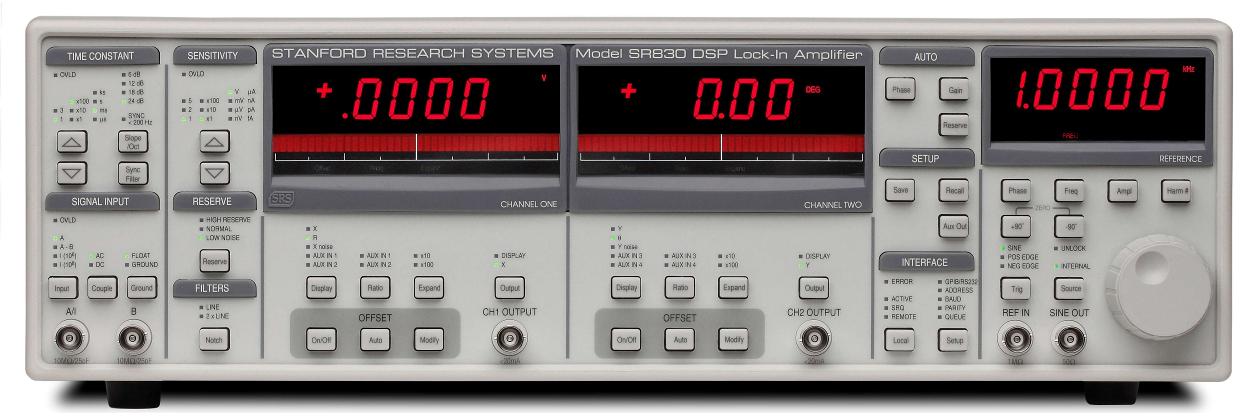

Figure 3. The SR830 DSP Lock-in Amplifier by Stanford Research Systems.

Phase-sensitive detection & Lock-in amplifiers

Avoiding 1/f noise brings us to the topic of phase-sensitive detection and lock-in amplifiers. You have already seen the basic concepts of phase-sensitive detection many times when working with oscilloscopes. If you have a small periodic signal of some sort, then you set up a synchronized (SYNC) signal with a much larger amplitude and use this to trigger the 'scope. And then you average traces to increase the SNR until you can measure the amplitude of the small signal. If the DC signal level varies slowly with time, the trace average is hardly affected, so your signal measurement is insensitive to 1/f noise. A *lock-in amplifier* expands this measurement strategy and takes it into the realm of precision instrumentation. The specific example we use in this Track is shown in Figure 3.

Before getting lost in all the knobs and dials in the lab, it is useful to outline just what this instrument does. Unlike a DMM, a Lock-in does nothing with DC signals. You always start by producing some kind of periodic signal, and the usual method is to "chop" some kind of detected signal on and off, as illustrated in Figure 4. Here we shine a blue laser on a sample material, and the theory we want to test predicts that the laser would produce a weak fluorescent signal in the red part of the spectrum. The expected red fluorescence photons are focused onto a photodetector, while the filter rejects all the blue laser light.

The important features of this measurement example are:

1) The chopped laser light means that the fluorescence signal looks like a square wave – the red photons are on when the laser is on and off when the laser is off.
2) The experimenter knows when the laser is on/off and creates an analog "Reference" signal $Ref(t)$ that is tied to the rotation of the chopper wheel. The lock-in interprets this signal and produces a digital signal $RefDig(t)$ that is equal to 1 when the laser is on and -1 when the laser is off.
3) The lock-in then produces an output signal equal to the running average

$$V_{out} = \langle RegDig(t) * V_{in}(t) \rangle_T \qquad (1)$$

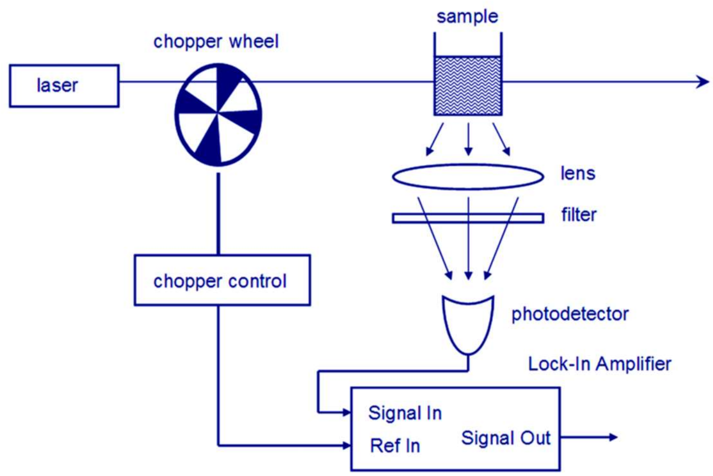

Figure 4. A standard example illustrating phase-sensitive measurement using a commercial lock-in amplifier. Here the goal is to detect weak fluorescence signal from a sample that is illuminated by a chopped laser beam. Note that the optical filter blocks any scattered laser light while letting the fluorescence light reach the photodetector.

where the averaging time T is selected by the user. The lock-in amplifier combines a precision preamplifier (similar to the SR560 you used earlier in this Track) and a specialized digital data acquisition system, all designed specifically for making precision measurements of low-amplitude chopped signals. Additionally, the lock-in can output both the in-phase and quadrature components of the signal (the latter being 90 degrees out of phase with the reference signal), or the amplitude and phase of the input signal relative to the reference signal.

Note that the primary reason one uses this "chopping" measurement strategy is to isolate a signal that is synchronized with the Reference signal – which is we call this phase-sensitive detection. All other signals are averaged away, so all off-resonant signals will average to zero with long averaging times. For example, if the laser were chopped at 1kHz, then signals arising from ambient room lights hitting the detector would quickly average to zero, while the desired signal would average to some finite value. Obviously, chopping the laser at 60 Hz or 120 Hz would be a bad idea, given the prevalence of electronic noise at those frequencies. And one avoids $1/f$ noise by chopping at sufficiently high frequencies (usually above 100 Hz is sufficient).

Looking at this measurement in the frequency domain, the lock-in averages the noise down except in a narrow bandwidth equal to about $B = 1/T$ around the Reference frequency. Because the total noise is proportional to $\sqrt{B}$, this spectral reduction can greatly reduce the overall noise in the measurement. In principle, a lock-in can measure signals down to the amplifier limit, equal to a few nV/rtHz for a good preamp… much better than any DMM. If you average for ten seconds or more, you should be able to see nanoVolt signals (on a good day anyway).

However, the example in Figure 4 also illustrates one danger with using lock-in detection. If there is no fluorescence whatsoever, but some scattered laser light (after the chopper wheel) manages to get into the detector, then you might easily mistake this laser light for your desired fluorescence signal. The optical filter in the figure is meant to reject the blue laser light while letting through the red fluorescence light, but the filter might not be 100% effective, or some laser light might make its way around the filter. As with any precision measurement, once you reduce the known noise sources, you must still be on the lookout for spurious signals masquerading as whatever it is you are trying to measure.

Laboratory Exercises

Begin your laboratory session by turning on the Keysight DMM and locating the printed-circuit board labeled PCB-34. Figure 5 shows a circuit diagram for this board, and Figure 6 illustrates the physical layout of the board. If you look at the board itself, you should find a pair of removable "jumpers" that connect BNC1 & BNC2 to different resistors. These jumpers are short-circuiting wires that connect the posts they cover, and these posts are normally not connected. Place the jumpers on the positions connecting R2, and verify for yourself (from Figure 5 and Figure 6) that the effective circuit is now exceptionally simple circuit, with BNC1 and BNC2 connected only to R2. (This is your cue to sketch the equivalent circuit when these jumpers are in place.) Your first goal is to measure the value of R2 with the highest possible accuracy and precision.

2-wire & 4-wire resistance measurements

To measure the resistor value, connect either BNC1 or BNC2 to the DMM and select Ω2W, which is shorthand for a *2-wire resistance measurement*. This is the normal method for resistance measurement used by handheld DMMs, in which the instrument sends a known current through the resistor and measures the resulting voltage. You should see that the DMM measures a value somewhat above the

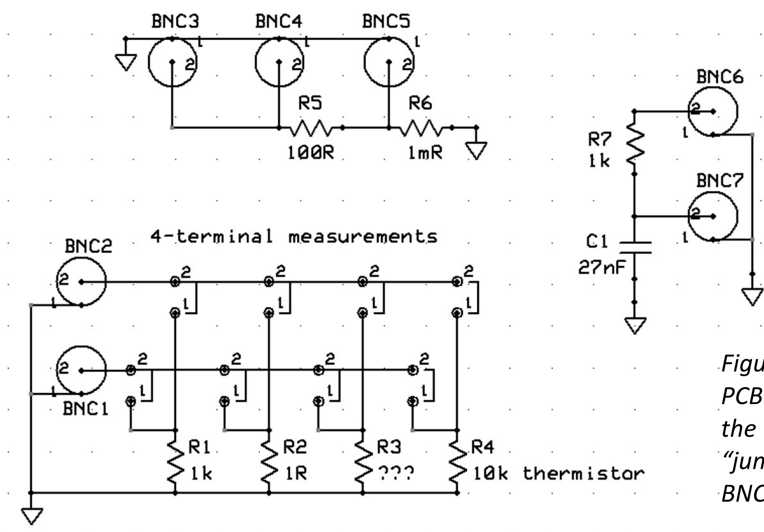

Figure 5. The circuit schematic for PCB-34. In the lower-left portion of the diagram, there are four pairs of "jumpers" that connect BNC1 & BNC2 to the four different resistors.

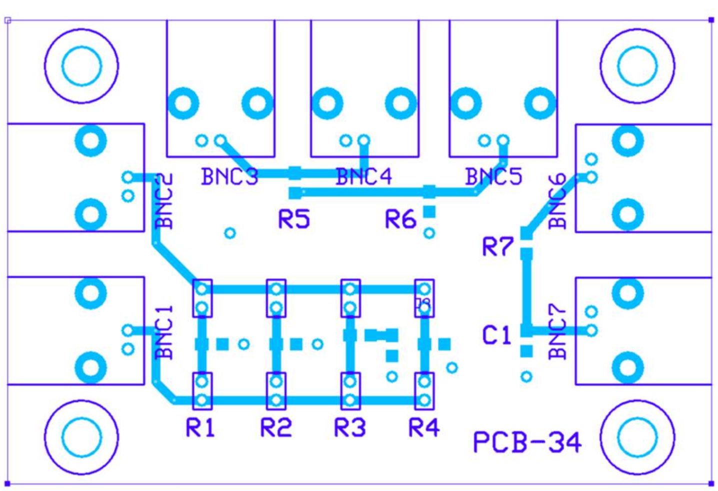

Figure 6. The physical layout of PCB-34. Compare this with the actual PCB and the schematic to understand what goes where.

expected 1 Ω, even though this resistor has a rated accuracy of 0.1%. The reason for this, as you might guess, is that the DMM is not just measuring R2, but this plus the resistance of the cable and the cable connections. The DMM sends a known current down through the cable inner conductor, through the resistor, and back up through the cable outer conductor, and it measures the resulting voltage inside the DMM. Thus the effective circuit should include R2 plus the cable and connector resistances, and these are not well determined in this setup.

Precision DMMs have a built-in solution to this problem, called a *4-wire resistance measurement*, which is illustrated in Figure 7. To implement this, connect the DMM *Input* and *Sense* ports to BNC1 and BNC2 (it doesn't matter which, because the circuit is symmetrical) and select Ω4W. Now the DMM sends the same known current down the Input port, but it measures the resulting voltage using the Sense port. Because the Sense port has a high input impedance, very little current flows through the Sense cable, and a low current means a low voltage drop in the cable. (From Ohm's law, the voltage drop in the cable is V = IR, where I is the Sense current, and R is the cable resistance.) This means that the 4-wire DMM measurement is largely unaffected by either of the BNC cables or any cable connections. (If this does not make sense, talk it over with your lab partner, or ask someone.)

The 4-wire measurement technique is useful whenever the cable and connection resistances are important in comparison to the resistor being measured, and the method is especially important when high accuracy is desired. This situation is common in the precision-measurement world, so all precision DMMs include the capability for making 4-wire resistance measurements.

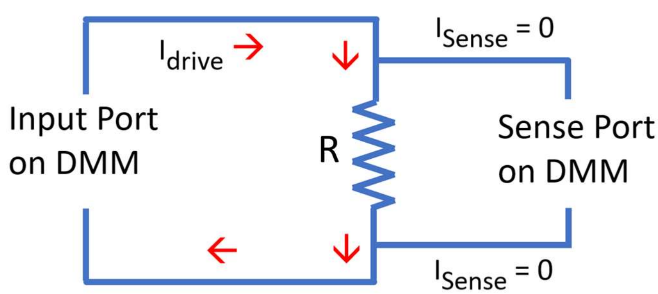

Figure 7. The geometry of a 4-wire resistance measurement. The Input Port sends a known current through the resistor, while the Sense Port measures the voltage drop across the resistor. Because the DMM Sense port has a high input impedance, essentially no current flows through the sense wires. With no current flowing, there can be no voltage drop across the sense wires. Thus the resistance of the sense wires does not factor into the 4-wire measurement. And the resistance of the input wires does not matter either, because the DMM is sending a known current down these wires.

Exercise 1. Measure R2 using both the 2-wire and 4-wire techniques and report your results in your e-notebook. The second measurement should agree reasonably well with the resistor's rated absolute accuracy of ±0.1%. Resistors are universal electronic components, so manufacturers know how to make them well.

Error budgets

One of the most important activities you encounter in precision measurements involves instrument and device specifications, often called "spec sheets". When you want to make accurate measurements, you need to understand what you are measuring, how your measurements are being made, and the quality of your instruments. A standard way to approach this problem is to make an *error budget*, which is essentially just a list of all the potential errors and uncertainties (at least those you know about) in a measurement. To see how this works, let us walk through an especially simple example of measuring a resistor.

Begin by setting up a 4-wire measurement of the 1kΩ resistor (R1) and record your result. A first obvious question concerns the intrinsic accuracy of the DMM, and you can find this information laid out for you in the DMM specifications in Appendix I below. Under the heading of resistance measurements, go to the 1kΩ line and assume a time period of one year (the age of your instrument). As you can see, the chart says 0.0100 + 0.0010, and the key tells you that these numbers mean:

± (% of reading + % of range)

In our present case, the reading and the range are both near 1kΩ, so the expected instrument accuracy is ±0.011% of 1 kΩ = 0.11 Ω. Thus, the spec sheet tells you how much you can trust the resistor measurement, assuming just the uncertainties in the instrument.

The instrument temperature coefficient is also listed as 0.0006 + 0.0001, and you should verify that this means a measurement temperature coefficient of 0.007 Ω/C, which is negligible compared to the overall uncertainty of 0.11 Ω.

Now that you know how well the instrument can do, the next question regards the expected accuracy of the resistor value. Here again, the manufacturers measure these parameters quite well, and the resistors on this PCB have the following specifications:

Resistor specifications for PCB-34

Resistance	Uncertainty	Temp Coef
0.001 Ohm	±1%	±50ppm/°C
1 Ohm	±0.1%	±25ppm/°C
100 Ohm	±0.1%	±25ppm/°C
1000 Ohm	±0.05%	±2ppm/°C
R3	±0.01%	±2ppm/°C

Putting in the numbers, verify for yourself that the intrinsic value of R1 should be 1kΩ, with overall uncertainty of ±0.5 Ω, and a temperature coefficient of ±0.002 Ω/C.

Thus you can see that the R1 uncertainty (0.5 Ω) is greater than the DMM uncertainty (0.11 Ω), so we can neglect the latter. We measured both of our PCB-34 boards using three different Keysight 34461A instruments, and in all cases the measured R1 was equal to 1kΩ to an overall accuracy of better than ±0.1 Ω. Thus these 1kΩ resistors beat spec by a healthy margin, so all is right with the world. Of course resistors are not exactly cutting-edge technology, so we expect that things should work as advertised. Of course, things are not so easy when you graduate to more complex physics experiments.

Exercise 2. Do a similar analysis using R2 (1Ω ±0.1%) and record your results in your e-notebook. Assume that room temperature can vary by ±1 C, which factors into the expected R2 value.

Exercise 3. Measure the value of R3 and report your result. Include an error budget along with your measurement, thus providing a realistic uncertainty estimate for your reported value. The goal here is not just to measure the value of an unknown resistor, but to understand how one makes a suitable error budget for a precision measurement.

If you think this kind of analysis is exceedingly dull, you are correct. However (spoiler alert), working as a professional scientist can be tedious, and a lot of modern experimental physics research involves precision measurements. In Ph77, we believe it is good to have some exposure to these concepts now, before you sign on for some CMB-scale experimental venture.

Exercise 4. As long as you have this set up, go back to your 4-wire measurement of the 1kΩ resistor (R1). Press Display/Histogram, then Clear Readings and play around with the settings to make a nice plot on the screen. Clear the readings, record for about 30 seconds, then stop the system and take a photo for your e-notebook using your phone camera. Select Math to see statistics along with the plotted distribution. The histogram feature gives you a quick look at the statistical fluctuations in the readings, and you should know that this feature exists in high-end DMMs. In the case of R1, the actual measurement uncertainties are dominated by systematic effects. But the statistical distribution tells how well you can see small changes in the resistor value over short times, which is often useful.

Exercise 5. Switch to an R1 (1000 Ω) 4-wire measurement and press Display/Trend Chart. Play around with the settings to see how this works. While the chart is running, try blowing down gently on R1. Does the result meet expectations based on the specified temperature coefficient? Make a nice plot showing a temperature rise and drop, and add a photo to your e-notebook with your analysis.

Next switch to measuring R4 on the PCB, which is a 10 kΩ *thermistor*. A thermistor is much like any other resistor, except the temperature coefficient was made as *high* as possible… just the opposite of what you want for most resistor applications. Thermistors are used as temperature sensors, and this particular example has a temperature coefficient of about -4.5%/C. The resistance is about 10 kΩ at room temperature, decreasing about 4.5% per degree as the temperature increases. Thermistors are best suited for measuring temperatures from about -50C to 100C, which is convenient for human environments. If you try blowing down on R4, you can easily see the resistance changes. Thermistors are great for measuring small temperature changes, but their absolute accuracy is rarely better than 0.1 C.

If you watch a Trend Chart, you can see that pairing a thermistor with a 6½ precision DMM makes for a *very* sensitive thermometer! In fact, the room temperature is simply not stable enough to see statistical fluctuations in your R4 measurement; the measurements simply reflect small environmental changes. This illustrates a nice use of a precision DMM to record temperature changes in the lab or in an experiment. The DMM will save all the numbers in memory (up to 10,000 samples with this instrument) and you can download the data to a USB drive for later analysis.

Exercise 6. Go back to R3 to see statistical fluctuations in the measurement using the Histogram feature. Measure the standard deviation δR using the Math display, and assume δR/R is the same for an R4 measurement. For a single data sample, what is the temperature sensitivity δT you can obtain with R4 and this DMM?

Exercise 7. Say you want to measure a very small resistance like 1 mΩ. Make an error budget for this case, which clearly shows that the Keysight DMM is a very good tool for this job.

Before moving on from the DMM, note that you always need to apply some judgment when it comes to measurement uncertainties. A DMM will measure R, and you can measure the statistical σ_R with a few button presses. But this strategy does not tell you anything about the instrument's overall accuracy, and that is where error budgeting comes into play. And you are often on your own when it comes to assessing the origins of systematic errors. One common feature of cutting-edge science and technology, true across many fields, is that you are never exactly sure what you are doing. But today's unknown is tomorrow's discovery. If you just keep at it, something good might happen.

Phase-sensitive detection using a Lock-in Amplifier

Lock-in amplifiers are general-purpose instruments that have been optimized for measuring small signals. To see the general principles in action, start by setting up the circuit shown in Figure 8. Set up the signal generator to produce a 1 kHz sine-wave signal at 1 Vrms, using the latter because lock-ins generally work with rms voltages. And let us walk through some lock-in settings:

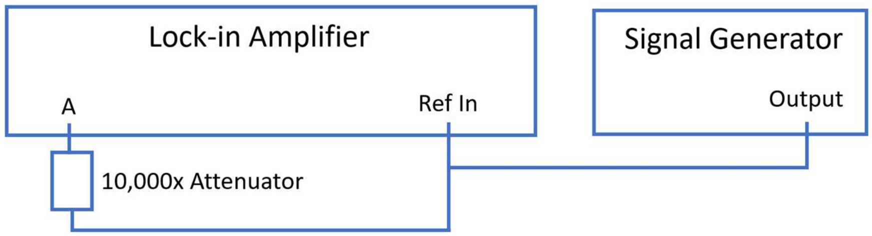

Figure 8. A simple set-up for examining the main functions of a lock-in amplifier. Here a large sine-wave signal provides the Reference, while a much smaller derivative signal provides the voltage being measured.

1) Set the Time Constant to 100 msec. This is fast enough compared to human reaction times, but slow enough to do some signal averaging to reduce noise.
2) Set the Signal Input to A/AC/Float… then the instrument looks at channel A only, with AC coupling, and ground side of the BNC is not grounded internally in the instrument.
3) Set the Sensitivity to 200 μV. The Sensitivity is equal to the maximum input signal before saturating the instrument's A/D converter.
4) Reserve = Low Noise
5) Set the Displays to X (the signal component in phase with the Reference signal) and Y (the quadrature component of the input signal).
6) Set the Outputs to Display. With this, the maximum number of "bars" on the display produces and output voltage of 10 V.

Given that the attenuation value of 10,000x is only accurate to 1%, verify that the observed signal matches expectations. If you switch from X/Y to R/θ, you can see the signal amplitude R and the phase difference θ between the Input and Reference signals.

Exercise 8. Connect the X signal output to the DMM and measure a histogram of the data. Take a picture of the histogram for your e-notebook. This shows the statistical noise on the measurement, and you should see a SNR of better than 500ppm. So measuring a 100 μV signal is super easy with a lock-in!

When you are using this lock-in, it is quite easy to produce meaningless results, as the instrument does not have a good "granny mode" (this being characterized by lots of self-scaling, auto-adjusting features). As the operator of this instrument, you have to pay attention to what you are doing. Here are some settings and indicators to watch for:

1) **Reference Unlock**. If this small red LED is on, then the Lock-in is not locked onto a reference signal, so all measurements will be meaningless. To fix this, supply a proper reference signal to the Ref IN port.
2) **Signal Overload**. If you see a small illuminated "OVLD" indicator in either of the display panels, then the input port is overloaded, so all measurements are again meaningless. To fix this, set the Sensitivity to a higher maximum voltage.

3) **Output bars**. If you do not see any bars on the outputs, then you should decrease the maximum Sensitivity voltage. If you see full bars, then you probably have an overload condition, or are close to it. Best to see an intermediate number of bars on at least one of the outputs. Remember that full bars (10V on the output) means the input signal is equal to the Sensitivity setting.
4) **Output Offsets**. If you see an illuminated "Offset" indicator in either of the display panels, then some nonzero voltage is being added to the output, which you probably do not want. Toggle the *Offset On/Off* button to remove the offset.
5) **Display Channel**. Make sure both displays are set to what you want, usually to X, Y, R, θ, etc.
6) **Phase**. Usually the phase (relative to the Reference signal) should be set to zero, which you can do by pressing the +90 and -90 buttons simultaneously.
7) **Time Constant**. This describes the input low-pass filter, and we recommend using a human-friendly value between 0.1 and 1 seconds.

Exercise 9. Add a second 10,000x attenuator to the input, so now your Reference signal is being attenuated by a factor of 100 million. Set the Time Constant to 1 second, so now your measurement has a 1-Hz bandwidth. Adjust the Sensitivity until you see some bars, and send the Output to the DMM for a histogram measurement. To avoid noise from cable contacts, do not touch any of the cabling when you are collecting a histogram. How does your measured signal compare with expectations? Document your measurement, including a photo of your histogram and a measured noise level expressed in nV/rtHz. [Recall from earlier in this Track that a state-of-the-art room-temperature amplifier exhibits an input noise of about 1 nV/rtHz, limited mainly by thermal noise from the internal resistors. When you add electrical protection circuity (to deal with static electric shocks and careless users), high-quality test equipment (like the SR830 Lock-in and the SR560 amplifier you used previously) have input noise levels of a few nV/rtHz under good conditions.]

Exercise 10. Use the lock-in to measure the value of the 1mΩ resistor on PCB-34. Include a sketch of your circuit diagram. Note that the signal generator has an output impedance of 50Ω, so connecting a 1 Vrms signal to BNC4 will pull the voltage at that point down to 0.667 Vrms (why?). Measure the value of the 1mΩ resistor and report this with the statistical noise expressed in Ohms/rtHz. Include a photo of your histogram, and document your work. [Hint: although the 1mΩ resistor is specified as ±1% in the above table, this does not include the resistance of the solder connections and wiring on the PCB, which can be significant at the mΩ level.]

Exercise 11. Use the lock-in to measure the capacitor on PCB-34. A good strategy here is to observe R and θ on the displays, and then adjust the drive frequency until θ = 45 degrees. Then use theory to calculate the capacitance C. You will see that this strategy gives a measurement of C that is independent of the input voltage, which is good. This particular capacitor is specified at ±5%, worse than the resistors because it is difficult to manufacture accurate capacitors. Is the capacitor in spec?

As you can surmise, precision DMMs and Lock-in amplifiers can be useful general-purpose instruments in the lab. Both allow you to make high-quality measurements without having to build any hardware or develop any data-acquisition software. If your career takes you into experimental physics research or industrial hardware development, you may well encounter these tools again.

Appendix 1
A selection of specifications for the Keysight 34461A Digital Multimeter

Specifications 34461A

- 34461A accuracy specifications: ± (% of reading + % of range) [1].
- These specifications are compliant to ISO/IEC 17025 for k = 2.

Range [2]/frequency		24 hours [3] T_{CAL} ± 1 °C	90 days T_{CAL} ± 5 °C	1 year T_{CAL} ± 5 °C	2 years T_{CAL} ± 5 °C	Temperature coefficient/°C [4]
DC voltage						
100 mV		0.0030 + 0.0030	0.0040 + 0.0035	0.0050 + 0.0035	0.0065 + 0.0035	0.0005 + 0.0005
1 V		0.0020 + 0.0006	0.0030 + 0.0007	0.0040 + 0.0007	0.0055 + 0.0007	0.0005 + 0.0001
10 V		0.0015 + 0.0004	0.0020 + 0.0005	0.0035 + 0.0005	0.0050 + 0.0005	0.0005 + 0.0001
100 V		0.0020 + 0.0006	0.0035 + 0.0006	0.0045 + 0.0006	0.0060 + 0.0006	0.0005 + 0.0001
1000 V		0.0020 + 0.0006	0.0035 + 0.0010	0.0045 + 0.0010	0.0060 + 0.0010	0.0005 + 0.0001
True RMS AC voltage [2, 5, 6]						
100 mV, 1 V, 10 V, 100 V, and 750 V ranges						
3 Hz to 5 Hz		1.00 + 0.02	1.00 + 0.03	1.00 + 0.03	1.00 + 0.03	0.100 + 0.003
5 Hz to 10 Hz		0.35 + 0.02	0.35 + 0.03	0.35 + 0.03	0.35 + 0.03	0.035 + 0.003
10 Hz to 20 kHz		0.04 + 0.02	0.05 + 0.03	0.06 + 0.03	0.07 + 0.03	0.005 + 0.003
20 kHz to 50 kHz		0.10 + 0.04	0.11 + 0.05	0.12 + 0.05	0.13 + 0.05	0.011 + 0.005
50 kHz to 100 kHz		0.55 + 0.08	0.60 + 0.08	0.60 + 0.08	0.60 + 0.08	0.060 + 0.008
100 kHz to 300 kHz		4.00 + 0.50	4.00 + 0.50	4.00 + 0.50	4.00 + 0.50	0.200 + 0.020
Resistance [7]	**Test current**					
100 Ω	1 mA	0.0030 + 0.0030	0.0080 + 0.0040	0.0100 + 0.0040	0.0120 + 0.0040	0.0006 + 0.0005
1 kΩ	1 mA	0.0020 + 0.0005	0.0080 + 0.0010	0.0100 + 0.0010	0.0120 + 0.0010	0.0006 + 0.0001
10 kΩ	100 µA	0.0020 + 0.0005	0.0080 + 0.0010	0.0100 + 0.0010	0.0120 + 0.0010	0.0006 + 0.0001
100 kΩ	10 µA	0.0020 + 0.0005	0.0080 + 0.0010	0.0100 + 0.0010	0.0120 + 0.0010	0.0006 + 0.0001
1 MΩ	5 µA	0.0020 + 0.0010	0.0080 + 0.0010	0.0100 + 0.0010	0.0120 + 0.0010	0.0010 + 0.0002
10 MΩ	500 nA	0.0150 + 0.0010	0.0200 + 0.0010	0.0400 + 0.0010	0.0600 + 0.0010	0.0030 + 0.0004
100 MΩ	500 nA || 10 MΩ	0.3000 + 0.0100	0.8000 + 0.0100	0.8000 + 0.0100	0.8000 + 0.0100	0.1500 + 0.0002
DC current	**Burden voltage**					
100 µA	< 0.011 V	0.0100 + 0.0200	0.0400 + 0.0250	0.0500 + 0.0250	0.0600 + 0.0250	0.0020 + 0.0030
1 mA	< 0.11 V	0.0070 + 0.0060	0.0300 + 0.0060	0.0500 + 0.0060	0.0600 + 0.0060	0.0020 + 0.0005
10 mA	< 0.05 V	0.0070 + 0.0200	0.0300 + 0.0200	0.0500 + 0.0200	0.0600 + 0.0200	0.0020 + 0.0020
100 mA	< 0.5 V	0.0100 + 0.0040	0.0300 + 0.0050	0.0500 + 0.0050	0.0600 + 0.0050	0.0020 + 0.0005
1 A	< 0.7 V	0.0500 + 0.0060	0.0800 + 0.0100	0.1000 + 0.0100	0.1200 + 0.0100	0.0050 + 0.0010
3 A	< 2.0 V	0.1800 + 0.0200	0.2000 + 0.0200	0.2000 + 0.0200	0.2300 + 0.0200	0.0050 + 0.0020
10 A [8]	< 0.5 V	0.0500 + 0.0100	0.1200 + 0.0100	0.1200 + 0.0100	0.1500 + 0.0100	0.0050 + 0.0010
Capacitance [15]						
1 nF		0.50 + 0.50	0.50 + 0.50	0.50 + 0.50	0.50 + 0.50	0.05 + 0.05
10 nF		0.40 + 0.10	0.40 + 0.10	0.40 + 0.10	0.40 + 0.10	0.05 + 0.01
100 nF		0.40 + 0.10	0.40 + 0.10	0.40 + 0.10	0.40 + 0.10	0.05 + 0.01
1 µF		0.40 + 0.10	0.40 + 0.10	0.40 + 0.10	0.40 + 0.10	0.05 + 0.01
10 µF		0.40 + 0.10	0.40 + 0.10	0.40 + 0.10	0.40 + 0.10	0.05 + 0.01
100 µF		0.40 + 0.10	0.40 + 0.10	0.40 + 0.10	0.40 + 0.10	0.05 + 0.01
True RMS AC current [2, 6, 9]		**Burden voltage**				
100 µA, 1 mA, 10 mA, and 100 mA ranges		**< 0.011, < 0.11, < 0.05, < 0.5 V**				
3 Hz to 5 kHz		0.10 + 0.04	0.10 + 0.04	0.10 + 0.04	0.10 + 0.04	0.015 + 0.006
5 kHz to 10 kHz (typ)		0.10 + 0.04	0.10 + 0.04	0.10 + 0.04	0.10 + 0.04	0.030 + 0.006

Range [2]/frequency	24 hours [3] T_{CAL} ± 1 °C	90 days T_{CAL} ± 5 °C	1 year T_{CAL} ± 5 °C	2 years T_{CAL} ± 5 °C	Temperature coefficient/°C [4]
1 A range	**< 0.7 V**				
3 Hz to 5 kHz	0.10 + 0.04	0.10 + 0.04	0.10 + 0.04	0.10 + 0.04	0.015 + 0.006
5 kHz to 10 kHz (typ)	0.10 + 0.04	0.10 + 0.04	0.10 + 0.04	0.10 + 0.04	0.030 + 0.006
3 A range	**< 2.0 V**				
3 Hz to 5 kHz	0.23 + 0.04	0.23 + 0.04	0.23 + 0.04	0.23 + 0.04	0.015 + 0.006
5 kHz to 10 kHz (typ)	0.23 + 0.04	0.23 + 0.04	0.23 + 0.04	0.23 + 0.04	0.030 + 0.006
10 A range [8]	**< 0.5 V**				
3 Hz to 5 kHz	0.15 + 0.04	0.15 + 0.04	0.15 + 0.04	0.15 + 0.04	0.015 + 0.006
5 kHz to 10 kHz (typ)	0.15 + 0.04	0.15 + 0.04	0.15 + 0.04	0.15 + 0.04	0.030 + 0.006
Continuity					
1 kΩ	0.0020 + 0.0300	0.0080 + 0.0300	0.0100 + 0.0300	0.0120 + 0.0300	0.0010 + 0.0020
Diode test [10]					
5 V	0.0020 + 0.0300	0.0080 + 0.0300	0.0100 + 0.0300	0.0120 + 0.0300	0.0010 + 0.0020
DC ratio (typ)					
	(Normalized input accuracy) + (Normalized reference accuracy)				
Temperature [11]					
PT100 (DIN/IEC 751)	Probe accuracy + 0.05 °C				
Thermistor	Probe accuracy + 0.1 °C				
Frequency: specification ± (% of reading) [12, 13]					
100 mV, 1 V, 10 V, 100 V, and 750 V ranges [14]					
3 Hz to 10 Hz	0.100	0.100	0.100	0.100	0.100
10 Hz to 100 Hz	0.030	0.030	0.030	0.030	0.035
100 Hz to 1 kHz	0.003	0.008	0.010	0.010	0.015
1 kHz to 300 kHz	0.002	0.006	0.010	0.010	0.015
Square wave [15]	0.001	0.006	0.010	0.010	0.015
Additional gate time errors ± (% of reading) [13]					
Frequency	**1 second**	**0.1 second**	**0.01 second**		
3 Hz to 40 Hz	0	0.200	0.200		
40 Hz to 100 Hz	0	0.060	0.200		
100 Hz to 1 kHz	0	0.020	0.200		
1 kHz to 300 kHz	0	0.004	0.030		
Square wave [15]	0	0	0		

1. For DC: Specifications are for 90-minute warm-up, aperture of 10 or 100 NPLC, and auto zero on. For AC: Specifications are for 90-minute warm-up, slow AC filter, sine wave.
2. 20% over range on all ranges, except 1,000 V DCV, 750 ACV, 10 A DC, 3 A Current, 10 A Current, and diode test.
3. Relative to calibration standards.
4. Add this for each °C outside T_{CAL} ±5°C.
5. Specifications are for sine wave input > 0.3% of range and > 1 mVrms. 750 ACV range limited to 8 x 10^7 V.Hz.
6. Low-frequency performance: three filter settings are available: 3 Hz, 20 Hz, 200 Hz. Frequencies greater than these filter settings are specified with no additional errors.
7. Specifications are for 4-wire ohms function or 2-wire ohms function using math null for offset. Without math null, add 0.2 Ω additional error in 2-wire ohms function.
8. The 10 A range is only available on a separate front-panel connector. Add 2 mA base per amp for inputs > 5 Arms.
9. Specifications are for sine wave input > 1% of range and > 10 µA AC.
10. Specifications are for the voltage measured at the input terminals. The 1 mA test current is typical. Variation in the current source will create some variation in the voltage drop across a diode junction.
11. Actual measurement range and probe errors will be limited by the selected probe. Probe accuracy adder includes all measurement and ITS-90 temperature conversion errors. PT100 Ro settable to 100 Ω ± 5 Ω to remove the initial probe error. Thermistor type: 2.2 kΩ (model number 44004), 5 kΩ (model number 44007) and 10 kΩ (model number 44006).
12. Specifications are for 90-minute warm-up and sine wave input unless stated otherwise. Specifications are for 1-second gate time (7 digits).
13. Applies to sine and square inputs ≥ 100 mV. For 10 mV to < 100 mV inputs, multiply % of reading error x10.
14. Amplitude 10% to 120% of range and less than 750 ACV.
15. Square wave input specified for 10 Hz to 300 kHz.