

Dick Cappels' project pages <http://www.projects.cappels.org>

return to [HOME](#)

## **A Microcontroller Based Digital Lock-In Milliohmmeter**

Described are the waveform capture method, example firmware and hardware designs. This material formed the basis of an article that was first published in the January, 2004 issue of Circuit Cellar magazine.

Download: assembler source [mhm031002A.asm](#)

A milliohmmeter is just the tool for checking trace resistance on a printed circuit board, tracking down shorted traces, and measuring the contact resistance of a switch or connector. Its the kind of tool that would come in real handy occasionally, but not often enough to justify shelling out hundreds of dollars. Wanting one anyway, I set out to make my own. It turned out to be not only an exciting project, but a true adventure of discovery as it provided a window into the workings of lock-in amplifiers. With a lock-in amplifier topology, a microcontroller and a little firmware makes the venerable (if somewhat noisy) LM324 bipolar op-amp provide high gain and at the same time reduces noise. Improve analog performance with a microcontroller? Now THIS is a fun project!

My main concern in the milliohmmeter design was how to get stable resistance readings with a suitably low test current. After all, I wanted to look at a few or few tens of milliohms, but I wanted to keep the test current low so as not to exceed the current ratings of some of the parts I want to test, and so I can run it all from a battery. For starters, 25 milliamps x 1 milliohm = 25 microvolts. This meant that I would have to be able to measure and display voltages in the range of tens of microvolts per count in a stable and repeatable manor. I also wanted the circuit to be forgiving of my hand wired breadboards, shown in photo 1. That is what lead me to the lock-in amplifier.




Photo1 Because it integrates the signal over 999 cycles, the lock-in amplifier is very forgiving of my layout and wiring. The microcontroller is at the opposite end of the board from the analog circuitry. An additional transistor inverter is an RS-232 receiver left over from an earlier project. The current switching PNP and input shorting pushbutton are on a small daughter board. Also shown is a 100 milliohm test resistance.

### The Wonder Of Lock-In Amplification

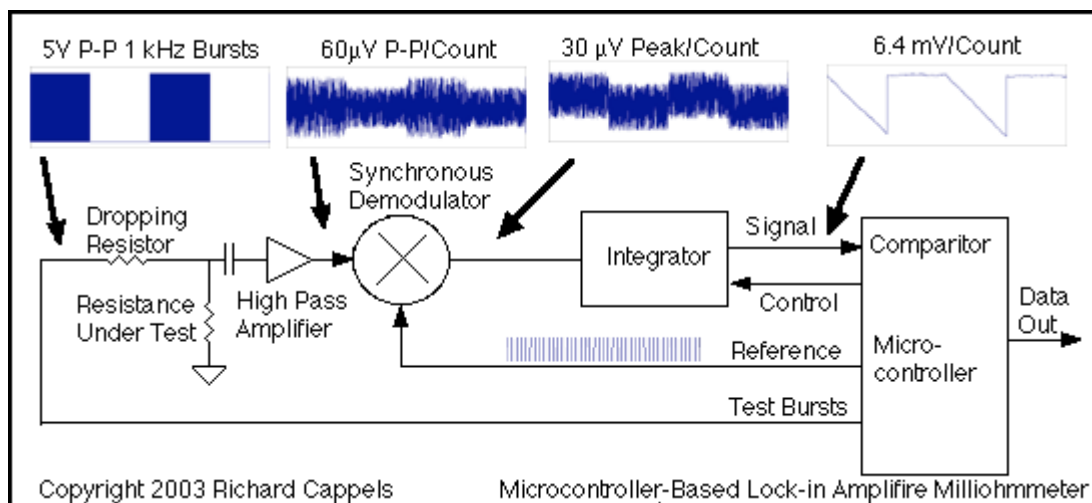


figure 1 The lock-in amplifier detects and integrates bursts of 1 kHz pulses that result from the IR drop across the resistance under test. The result is 46 db of signal gain while vastly improving the signal-to-noise ratio. The preamplifier doubles as a high-pass filter to remove the DC component of the pulses and reduce flicker burst noise from the op-amp.

A lock-in amplifier provides gain for the signal at a specific frequency and phase, while reducing noise. Its almost like getting something for nothing, but that's not really the case. Since you have to effectively average many measurements to improve

the signal to noise ratio, what you are really doing is trading way bandwidth and responsiveness to improve the signal to noise ratio.

In this circuit, shown in the block diagram in figure 1, a 1 kHz burst of 5 volt peak pulses is applied to a series dropping resistor to establish a pulsing current through the resistance under test. The IR drop across the resistor under test is a voltage proportional to the value of the resistance. After passing through a high pass filter to eliminate the DC component of the signal, it is rectified by a synchronous demodulator. For a period of time corresponding to the burst, the rectified signal is applied to an integrator, which charges up, and then at the end of the burst, the time it takes the integrator to discharge at a constant rate is measured to determine the size of the charge, and that is how the average amplitude of the signal applied to the integrator is determined. Since random noise is not synchronized with the switching in the synchronous demodulator, it is not rectified and it averages out to nearly zero. The longer the integration time, the less proportionate effect a given small pulse of noise has on the integrator's output, and the more gain the integrator has for the 1 kHz burst. Integration is the wonder of lock-in .amplification

### Getting The Drop

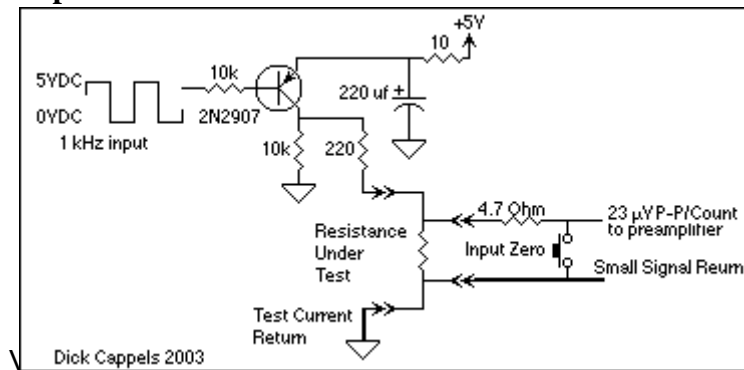


figure 2 The 1 KHz switching signal to the base of the 2N2907 causes 5 volts peak-to-peak to be applied to the 220 ohm resistor in series with the resistance under test, to cause 23 milliamps peak-to-peak to flow through it.

One critical part of measuring a very low resistance is developing an IR drop across the resistance, and measuring only that. Getting the IR drop itself was easy, a PNP switching transistor is driven into saturation with the 1 kHz burst, and the transistor's collector delivers 5 volt peak-to-peak pulses to a 220 ohm resistor in series with the resistance under test (see figure 2). The voltage across the 220 ohm resistor is virtually constant as long as the transistor is always driven into saturation and the resistance under test is very low. A 100 milliohm resistance will have a voltage drop of only 2.3 millivolts maximum, so the error introduced by using a 220 ohm resistor instead of a constant current source is only  $2.3 \text{ mv}/5\text{v} = 0.05\%$ , so I think you will agree that this approach is sufficient for a 100 count meter.

While testing the circuit I was surprised that when leaving the input open circuited, the meter gave a very low resistance measurement. I spent a long time looking over the firmware, looking for math errors or register use conflicts, before I decided to take

a look at the circuit with an oscilloscope. That's when I saw that when the meter's input terminals were unloaded, there was virtually no signal across the input terminals. Without the test resistance in series with the 220 Ohm resistor, there was no signal - the input just charged up to +5 volts and wiggled a little because of capacitive coupling between the transistor's base and collector. That's when I added the 10k resistor from the PNP's collector to ground. I should have taken my early mentor's advice, "Always look for the simplest explanation first."

To make sure that I am only measuring the resistance of the thing I am trying to measure and not the resistance of my test leads, I used a 4 wire measurement system. In this system, two wires are used to deliver the test current and two separate wires are used to measure the voltage drop. Simple enough, but how did I keep the test current ground return path separate from the voltage sense ground line? There in lies much of the art in circuit layout. Looking at figure 2, the terminal marked "Test Current Return" connects directly to the battery's negative connection to the circuit. The signal marked "Small Signal Return" connects to the rest of the analog circuits' grounds, and then joins up with the battery ground at the point in the circuit where the 1.8 volt reference voltage is generated. The overall grounding scheme can be seen in figure 3.

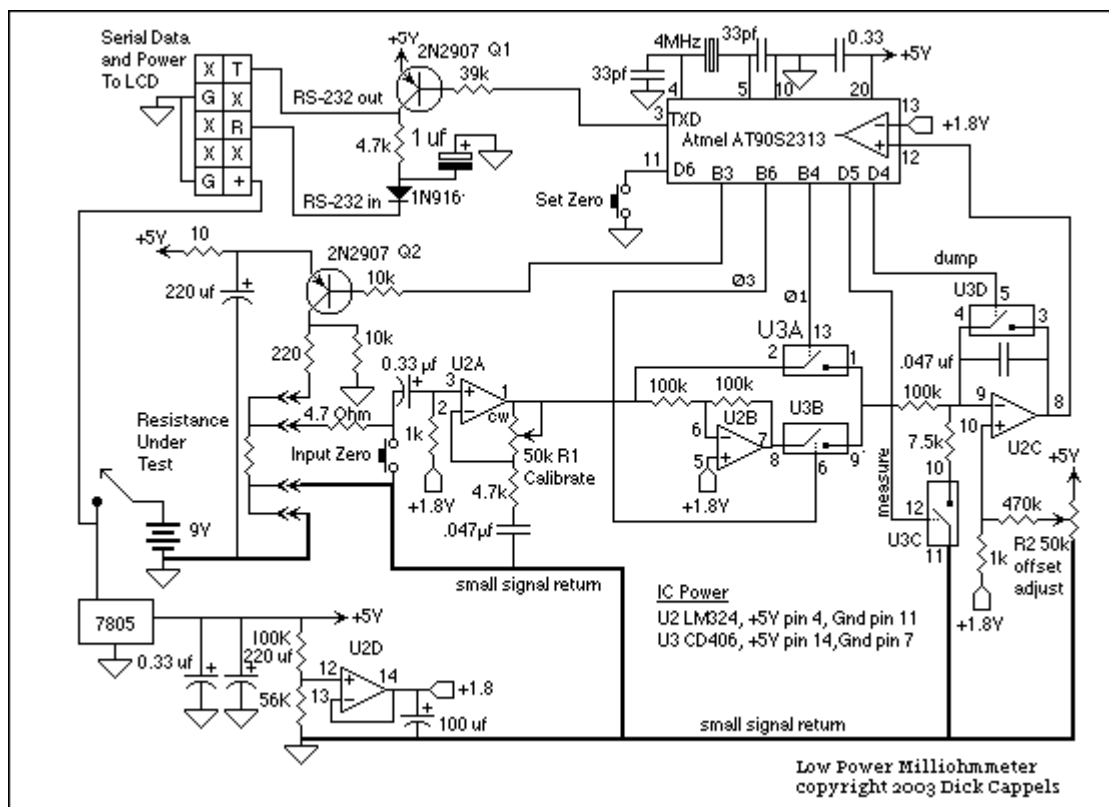


figure 3 Grounding is important to keeping offsets below the allotted maximum. Care needs to be taken to keep power supply currents from returning through the small signal grounds. If the RS-232 input is not connected, the 1N916 and 1 uF capacitor may be omitted and the lower end of the 4.7k resistor grounded.

The big trick here is to keep the various 1 kHz signals from the synchronous

demodulator and say currents from the test current circuit out of the signal from the resistance under test. The circuit is somewhat forgiving in that the unintentional signals end up affecting the integrator's output as offset and gain errors. The gain error is easily taken care of in calibration. The offset error is automatically take out by the firmware in the microcontroller.

Another challenge was finding a convenient way of shorting the input signal to zero the meter. I wanted to get an effective short of 1 milliohm or less, but FET's were out of the question for this kind of resistance, and the mechanical pushbuttons I had were coming out between 10 and 30 milliohms after a few dozen operations (one of the nice things about having a milliohmmeter is you can actually measure these parameters). I finally settled on the solution of putting a 4.7 ohm resistor in series with the signal path. If the test leads are connected to a low value resistor ( less than a couple hundred milliohms), pressing the Zero Input button will attenuate the signal to well below 23 microvolts - less than one count.

It is necessary to have a resistance under test connected to the test leads while zeroing so that the offset signals are the same for the case of zeroing the meter and taking actual measurements. Perhaps a much more careful layout plus some shielding would eliminate the need to do this.

### The Preamp




figure 4 The signal from the resistor under test is AC coupled to an amplifier biased to the 1.8 volt reference supply.

The first part circuit that the voltage across the resistance under test goes to is an AC amplifier with adjustable gain that doubles as a high-pass filter. The AC coupling on the input and feedback paths are required to eliminate the DC component of the signal from the resistance under test. The high pass filter is helpful in reducing noise pickup from power lines as well as rejecting some of the 1/f and shot noise in the op-amp itself. The input capacitor and 1 k bias resistor shown in figure 4 set the high pass corner frequency for the input at  $1/(2 \text{ Pi R C}) = 480 \text{ Hz}$ . A second high pass pole at about 720 Hz is set by the 4.7k resistor and .047 uf capacitor in the feedback loop. This means that the amplifier's frequency response is fairly flat for frequencies above 720 Hz, and for frequencies below 480 Hz, the amplifier's response is reduced by a factor of 1/4 each time the input frequency is cut in half, so at 60 Hz, where power line interference in North America is a concern, the amplifier's gain would be down to 1/64 of its response at 1 kHz. Since AC coupling was necessary, positioning the poles

at these frequencies enhanced the noise performance of the circuit at no additional cost.

The circuit has a voltage gain of 1 for DC because there is no attenuation of the DC signal in the feedback path. When I first built this circuit, the 1k bias resistor was much larger, and I was disappointed when I realized that a large offset voltage on the output resulted from the op-amp's input bias current. I could have switched to a FET input amplifier, but at least half the fun in this project was to try and get impressive performance with a cheap bipolar op-amp, so I lowered the input resistor to 1k and the offset became small enough to compensate for with the offset pot. Since there is no DC gain in this stage, it was not important to try and match the input bias between the two inputs, so I didn't feel compelled to make the input resistance on the inverting input match the resistance on the noninverting input.

The AC gain is adjustable from about 1:1 to 1:10 allowed me to calibrate the milliohmmeter without worrying about the tolerance of any of the components in the circuit. Looking back, I could have made the adjustment range much smaller to make the adjustment less critical.

### Synchronous Demodulators




figure 5 The switching signal is synchronized with the incoming signal, resulting in full wave rectification of the incoming signal. Signals that are not at the switching frequency of odd harmonics of the switching frequency are not rectified.

I remember that when I was a kid, when I switched on dad's tube type car radio in the quiet of an early morning, I could hear a faint hum of a few hundred Hz from under the dashboard before sound started coming out of the speaker. It was the vibrator power supply, its oscillating electromechanical contacts performing the functions of power switch and synchronous demodulator. These days, synchronous demodulators based on FET switch, diode bridges and transconductance multipliers are used in everything from power supplies to Bluetooth receivers.

They all have the same basic function: to rectifying the incoming waveform by switching the waveform into a filter or integrator using a signal that is synchronized with the desired incoming signal. A communications application of a synchronous rectifier similar to the synchronous demodulator described here was discussed by Ed

Nilsey [1] in his Above The Ground Plane Column in the October 2003 issue of Circuit Cellar. I've seen the terms "synchronous rectifier", "synchronous modulator" and "synchronous demodulator" tend to be used interchangeably to describe this type of circuit and its operation, with the usage seeming to depend on where and how the circuit is used.

In this project, I used an inverting amplifier and two sections of an CD4066 Quad Bilateral Switch (you can also use the 74HC4066) to make the synchronous demodulator. The block diagram of the demodulator is shown in figure 5. When the incoming waveform is positive, the incoming signal itself is connected to the circuit's output, and when the input signal swings negative, the inverted version of the input signal, which is positive, is connected to the circuit's output, so the output signal is always positive. The output of the circuit drives an op-amp integrator.

### Integrating and Measuring

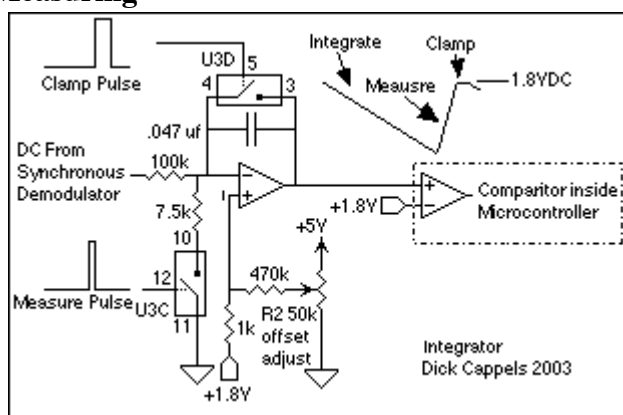


figure 6 The slope charging the integrator toward ground from the 1.8 volt reference is determined by the amplitude of the signal from the demodulator. The slope of the discharge back up to the 1.8 volt reference is constant.

The integrator shown in figure 6 is at the heart of the milliohmmeter's operation. It is used to both integrate the signal from the detector to provide gain while reducing noise, and it is used to measure the level of the signal. The signal from the demodulator causes the charge in the .047 uF integration capacitor to build up, driving the output toward ground from its 1.8 volt reference. After the 999 cycle of the 1 kHz measure burst, the synchronous demodulator is turned off by stopping the switching pulses.

To read out the total charge in the integrator, FET switch U3C is turned on to discharge the integrator capacitor with a constant current of 1.8 volts/7.5k = 240 microamps, producing a linear positive-going ramp at 240 microamps/.047 uF = 5,100 volts per second on the output of the op-amp. While the output of the op-amp is ramping back toward the 1.8 volt reference, the microcontroller, an Atmel AT90S2313 sits in a loop incrementing a counter every 1.25 microseconds, until the comparator changes state, indicating that the ramp has reached the 1.8 volt reference. The A to D conversion sensitivity is 1.25 microseconds per count X 5100 volts per microsecond = 6.375 millivolts per count.

The gain of the integrator is equal to  $(0.999 \text{ seconds} \times \text{peak input voltage}/100\text{k})/0.047 \text{ microfarads} = 212.5 \text{ volts out per volt in}$ , or 46 db. The conversion factor for the entire process, integration and measurement, is  $6.375 \text{ millivolts per count} / 212.5 \text{ volts/volt} = 30 \text{ microvolts peak per count}$ . Since the signal is AC coupled and peak-to-peak = 2X peak, the sensitivity is 60 microvolts peak-to-peak input per count.

The input amplifier sees 23 microvolts peak-to-peak per count on its input ( $23 \text{ milliamps} \times 1 \text{ milliohm}$ ), so the input amplifier would need a gain of about  $60 \text{ microvolts} / 23 \text{ microvolts} = 2.6\text{X}$ , but because the low frequency rolloff in the high pass filtering results in waveform distorting, the input amplifier needs a gain of about 4X. Thus, the preamp's adjustable gain of 1:1 to 1:10.

When the integrator finishes ramping back to the 1.8 volt reference, FET switch U3C is turned off and FET switch U3D is switched on to short the integrator capacitor and clamp the op-amp's output to the 1.8 volt reference. If for some reason the circuit were to find itself with the op-amp output more positive than the 1.8 volt reference at the end of the integration period, the circuit would latch up. This reset pulse assures that the circuit would not be in that state for more than one cycle.

### **Milliohms to Digits**

All the timing is performed with an Atmel AVR AT90S2313 microcontroller running at 4 Mhz. The presence of an analog comparator on the chip and fast program execution with low current requirements made this and similar ARV's ideal for this project. The on-chip UART, baud rate generator, interrupt timer, and EEPROM were also really nice to have. I believe this can be done with processors like the PIC, but I'm already familiar with AVR's and they well suited to this application.

The 1 kHz burst is generated by a 4 kHz interrupt that causes a 4 bit waveform image to be shifted out to the output pins 1/4 cycle at a time. I re-used code designed to produce quadrature signals for an earlier project. The firmware counts 999 cycles of the 1 kHz square waves to make a 999 millisecond burst.

It counts 999 1 kHz cycles per burst for the integration period, then allocates 250 microseconds for the measurement, using a timing loop of 1.25 microseconds long. The firmware allow a maximum of 190 counts (237.5 microseconds) so the measuring routine will fit within a 250 microsecond interrupt. After the measurement, it applies the clamp pulse for 750 microseconds before starting the next measurement cycles.

The firmware keeps track of odd and even measurement cycles. On odd measurement cycles, the test burst is not sent to the PNP transistor, though the demodulator operates and the integration and measurements take place. The measured value during the measurement cycle without the test burst, is very a very low count and is stored to be used as the measurement offset value.

On even measurement cycles, the burst is sent to the PNP transistor to put 5 volt peak-to-peak across the 220 ohm resistor and the resistance under test, while the resulting IR drop is integrated then measured. The measured offset value is subtracted from the measured resistance, then the value obtained with the input shorted is



subtracted. The result is converted from binary to ASCII digits and sent via the on-chip UART to be displayed.

Since the maximum A/D count is 190 milliohms, the A/D dynamic range has to be shared by the zero input value, the A/D zero value, and the resistance under test. The A/D is adjusted with the 50k offset potentiometer in the integrator circuit, and I set it to about 10 counts so it can drift a little either way and not cause clipping of the measured signal. This is the offset that is automatically measured and removed in every pair of measurement cycles, so the drift of this parameter will not affect the displayed resistance unless the offset goes below zero ohms or causes the measurement to exceed the A/D converter's maximum count of 190.

The input zero value is the measurement taken by the A/D converter with the measurement leads connected to a low value resistor and the input zero button pressed. This value is stored in the on-chip EEPROM so it is retained when power is removed.

The output data display, illustrated in figure 7, is composed of two lines that can be displayed on an RS-232 ASCII terminal, or in my case, on a two line by 16 character LCD with RS-232 input.

mOhms	Offset
101	8 [13]

figure 7 The offsets are displayed to aid in zeroing the circuit.

The top line shows the labels for measured milliohms and the offsets. The lower line displays the resistance reading in milliohms, the A/D converter offset count, which is controlled by the offset pot and zeroed out one each pair of measurement cycles, and the input zero offset, in square brackets, which is stored in EEPROM.

Since it takes two measurement cycles to measure the offset then measure the resistance under test, and each measurement takes 1 second, I added an LED to indicate when an actual measurement is being taken. Its particularly useful while zeroing the input. Hold down the input shorting button, momentarily press the zero button, and wait for two full cycles of the LED before letting the input shorting button up.

I set this up so that the serial LCD gets its power and data from a single connector on the milliohmmeter board, and the whole thing runs on a single 9V battery. It can set in on the shelf most of the time, and its a snap to move to the bench and use whenever I need it. Its surprising how often I've used this little meter now that I have it. around.

[HOME](#)

Contents ©2003 Richard Cappels All Rights Reserved.

<http://www.projects.cappels.org/>

Dick Cappels' web version first posted in January, 200.3.

You can send email to me at [projects\(at\)cappels.org](mailto:projects@cappels.org). Replace "(at)" with "@" before mailing.

**For commercial license, [click here](#).**



**Liability Disclaimer and intellectual property notice**

**(Summary: No warranties, use these pages at your own risk. You may use the information provided here for personal and educational purposes but you may not republish or use this information for any commercial purpose without explicit permission.)**  
I neither express nor imply any warranty for the quality, fitness for any particular purpose or user, or freedom from patents or other restrictions on the rights of use of any software, firmware, hardware, design, service, information, or advice provided, mentioned, or made reference to in these pages. By utilizing or relying on software, firmware, hardware, design, service, information, or advice provided, mentioned, or made reference to in these pages, the user takes responsibility to assume all risk and associated with said activity and hold Richard Cappels harmless in the event of any loss or expense associated with said activity. The contents of this web site, unless otherwise noted, is copyrighted by Richard Cappels. Use of information presented on this site for personal, nonprofit educational and noncommercial use is encouraged, but unless explicitly stated with respect to particular material, the material itself may not be republished or used directly for commercial purposes. For the purposes of this notice, copying binary data resulting from program files, including assembly source code and object (hex) files into semiconductor memories for personal, nonprofit educational or other noncommercial use is not considered republishing. Entities desiring to use any material published in this pages for commercial purposes should contact the respective copyright holder(s).

