4

Practical Guide to Instrumentation

INTRODUCTION

If you're like most of us, when you get a new piece of equipment, you play around with the knobs for a while, try it out on some unsuspecting circuitry, and then turn your attention to something more pressing. For items such as fancy oscilloscopes with lots of menus, you may occasionally have to look in the manual to find out where something is. Generally, though, well-designed equipment should be usable intuitively without any special training.

Even so, a number of considerations in the use of equipment and in measuring technique won't be in a manual, simply because they relate not so much to the design of the instrument as to how and when one particular instrument should be used rather than another. These instrumental considerations range from the general issues that apply to all instruments, such as the question of accuracy versus precision, to the specific, such as under what circumstances to distrust your AC power meter. This chapter offers practical guidance on general questions; it doesn't try to tell you how to use your voltmeter.

CALCULATORS AND CALCULATIONS

How Many Digits?

"I added the numbers four times, and here are the four answers."

We smirk when we hear this old chestnut, but in fact a lot of people are guilty of something similar all the time. The question is, How many digits should I write down when I make a measurement? We'll talk about accuracy versus precision in the next section about DVMs; for now, let's assume that your voltmeter is absolutely accurate, so that when it says the voltage is 15.426V you know that it is (for sure) somewhere between 15.4255V and 15.4265V.

Should you write down 15.426V? Or 15.4V? Or maybe just 15V? The answer depends on the purpose of recording the data. If you're checking that an op amp has the proper supply voltage, and you're using that op amp for general purpose only and don't care whether it has 14V or 16V as long as it works, then just writing down 15V is perfectly adequate—if indeed you bother to write it down at all. But the more common case is this: you're measuring both voltage and current, because you want to know power. Let's suppose you're using a handheld meter for the current, and it has only three digits, reading 2.02A. This means that you know the current only to three significant places—it's somewhere between 2.015A and 2.025A. Since in this case your knowledge of the current (three decimal places) is less than your knowledge of the voltage (five decimal places), the current controls your accuracy. *This means that your final answer, power, can be known only to the same three decimal places that your least accurate measurement is known to.* The correct thing to do (in this example) is to write down the voltage measurement to one more decimal place than the current, multiply, and then round off the answer to the three decimal places: $V = 15.43\text{V}$, $I = 2.02\text{A}$, so $P = 15.43\text{V} \times 2.02\text{A} = 31.1686\text{W}$, and the answer recorded should be 31.2W. Pay real close attention here!

The final result of the calculation CAN'T have more digits than the LEAST accurate measurement!

This is because the uncertainty is greatest in the least accurate measurement, and this in turn controls the accuracy with which you can know the answer.

Do I Care?

The question of uncertainty is not just one of those dull things you forgot in high school and have thereafter ignored with impunity. To rub in the point, consider how far off the calculation *could* be. If the voltage and current were *both* at the low end, even though the meters were reading correctly, $P = 15.4255\text{V} \times 2.015\text{A} = 31.08\text{W}$; and if both were at the high end, $P = 15.4265\text{V} \times 2.025\text{A} = 31.24\text{W}$, a difference of 160mW, *entirely due to the limited number of digits on the meters*. This doesn't seem like a big deal until you realize that you might be trying to measure the efficiency of an efficient converter. If the input power was 33.3W, then the efficiency is somewhere between $\eta = 31.08\text{W}/33.3\text{W} = 93.3\%$ and $\eta = 31.24\text{W}/33.3\text{W} = 93.8\%$—and this doesn't even include the inaccuracies in the input measurement! The difference between being able to report 94% efficiency rather than 93% efficiency can make or break a project. The correct thing to record is $\eta = 31.17\text{W}/33.3\text{W} = 93.6\%$, halfway between the minimum and maximum (note that the extra digit was again held over until the final calculation).

A Closely Related Problem

A closely related problem is not writing down *enough* digits. Maybe there's a little flicker in the last digit of the meter, and anyway there's not much difference between writing down 2.02A and just plain 2A, right? The example above shows that to be able to write down a number representing efficiency of an efficient converter to two digits requires writing down all the measurements to three digits. The reason here is slightly different from that given above: to calculate efficiency, two numbers that are very nearly equal are being divided, and so small inaccuracies in either one make the answer quite inaccurate—because you are presumably quite interested in whether the efficiency is 94% rather than 93%.

One More Problem to Avoid

A rather less common pernicious problem is the casual assumption that your measurement has unlimited precision: as we've emphasized, if a meter says 2.02A, that's what is meant, not 2.020000A! It's the same thing: your meter only has three digits, and there's no way to get a better number out of it by pushing buttons on a calculator. All you can do is get a better meter.

DVMs AND OTHER METERS

Accuracy versus Precision

In the foregoing discussion about how many digits to use in calculations, we assumed that what the meter showed was exactly right—that is, that the meter had *unlimited accuracy.* However, the meters showed only a certain number of digits (in the example, the handheld current meter showed three) and thus had a limited *precision*. But of course real meters have not only limited precision, they also have inaccuracies of various kinds, and problems of both types limit the validity of any measurements you make.

Averaging

It might be tempting to try to get a more accurate measurement in the case of a meter whose last (or several last) digit(s) are flickering by writing down the number several times and then averaging, as in Table 4.1. This idea is fine if you know for sure that the reason that the digits are flickering is due to white noise (random noise)—although in this case a better plan is to suppress the noise with a filter (see below). But if the noise is rather due to a periodic signal, such as an oscillation in the circuit at 1kHz, your meter may be reading the signal at the same point in each cycle, causing a systematic bias in the measurement, as in Figure 4.1. This is exactly the same thing as aliasing in an oscilloscope. In this case, you have an added inaccuracy, and filtering is essential; your "human averaging" is actually degrading the information.

TABLE 4.1 Sample Averaging Data

Trial 1	2.02A
Trial 2	2.06A
Trial 3	2.05A
Trial 4	2.00A
Average:	2.03A

Figure 4.1 If your DVM samples at a frequency related to an AC component of your waveform, you may get false readings.

How to Filter a DVM

To filter a DVM input, it's necessary to know what frequency noise you want to get rid of. Otherwise, you may be filtering the thing you want to measure! As an example of how to successfully filter a signal, suppose you want to filter a 100kHz noise source from a switching supply, and you need to reduce the noise by a factor of 10. This implies that the filter should have a frequency of about 100kHz/10 = 10kHz. This sets the product R times C, but not their individual values. The key here is that R is going to be in series with the input resistance of the meter. Therefore, to maintain a certain accuracy, R must be smaller than the input resistance by that amount. If the input resistance of the meter is 10MΩ, a typical value, and you are trying to achieve 0.1% accuracy, the R you choose for the filter has to be less than 10MΩ × 0.1% = 10kΩ. Completing the example, with R = 10kΩ, the capacitor (see Figure 4.2) will be

$$C = \frac{1}{2\pi\ 10\mathrm{kHz}\ 10\mathrm{k}\Omega} = 1.5\mathrm{nF}.$$

You may want to build up a few of these little circuits on perf boards with a pair of banana jacks on both ends, for easy availability when needed in the lab.

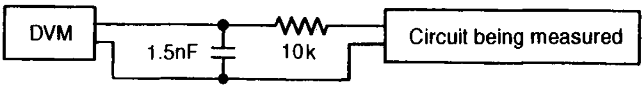

Figure 4.2 A filter for a DVM.

Measuring RMS and DVM Bandwidth

Your DVM updates its display a couple of times per second. However, this information doesn't necessarily have anything to do with its bandwidth, that is, the maximum frequency signal the instrument is able to measure. You find out the bandwidth by looking in the manual, but one thing to check is that the meter has adequate bandwidth for the measurement you're trying to make: if the bandwidth is 1kHz, the meter will certainly read a 10kHz signal as smaller than it really is.

Meters used in AC power measurements are a special case of this potential inadequacy. They frequently have very limited bandwidth, and so if you've attached, for example, a discontinuous flyback converter to the AC line, you can expect to get false readings unless you put a large capacitor in front to smooth the pulsating current.

Of course, putting a large cap in front is exactly the opposite of what you want to do to measure the power factor. In fact, measuring the input power and the power factor of an off-line converter may require two separate measurements, each with its own setup.

Another point to observe with RMS measurements is that many AC meters have a maximum *crest factor* they can tolerate; that is, if the peak (current) is very much higher than the average, they also give false readings—read the specs on this carefully.

Finally, it is worth observing that many types of meter have their internal circuitry somehow upset by high frequency noise, such as that generated by a switching power supply. (Isn't that convenient?) Almost all measurements on a converter are going to require filtering and careful attention to meter limitations.

Measuring Efficiency: Cross-Calibration

Here is a practical method for getting rid of accuracy problems where they are most important, during converter efficiency measurements. In measuring efficiency, you have to measure an input voltage and current, and an output voltage and current. If you were to simply write down the value of a meter reading from each of these four measurements, you'd end up with quite a bit of inaccuracy in the efficiency, because each meter contributes its own piece of inaccuracy to the result. However, note that the efficiency is the *ratio* of the two voltages times the *ratio* of the two currents. Thus the absolute values are of not much concern (you don't care whether the input is 27V or 28V, etc.) If both the voltage (or current) meters are then 0.3% high, *this inaccuracy cancels out*! Thus, what you want to do is what the author calls "cross-calibration": that is, before making the measurement, you attach both DVMs being used for the voltage (or current) measurements to exactly the same point (pass exactly the same current through them, respectively); they should read exactly the same. If they don't, you can find a scale factor for one of them that makes its reading equal to that of the other. If you now use the scale factor when measuring the efficiency, the inaccuracies still cancel out.

> **Practical Note** If you wanted to do it just like calibration people do, after the cross-calibration, you would next either ensure that the meter stays on the same scale throughout the measurement or cross-calibrate it with itself on the other scale, the one it's going to be used on. But practically, this isn't necessary, cross-calibration seems to work just fine without going to this extreme. In real-life a more serious problem occurs when someone comes by overnight, uses your meters, and puts them back in a different order. Now in the morning the cross-calibration is wrong, and you don't even know it! The author doesn't have a suggestion to fix this problem.

Where to Put the Probes

> **Practical Note** Put the DVM probes on separate leads from the power leads. Don't jack them in to the load.

The idea here is that since the power leads are carrying current, they have a voltage drop that increases, the further along the wire you go. So when you're doing a precision measurement (like efficiency, again), bring out separate wires from the output for the measurement, as shown in Figure 4.3. The same should of course be done for the inputs; often on converters, attention to this matter alone can be responsible for increasing the efficiency measured by 1 or 2% (the larger number is the correct one).

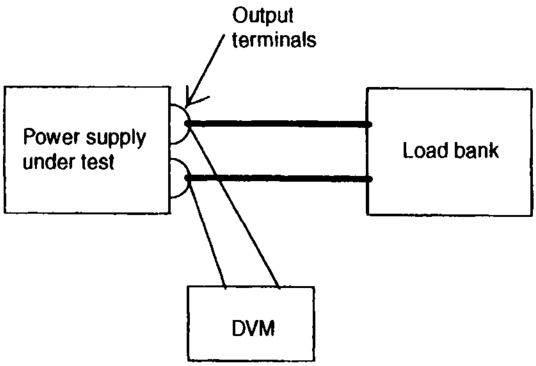

Figure 4.3 Use separate connections for carrying current and measuring voltage.

The same reasoning should be extended to the case of a small connector to which you must attach both the power connection and the meter, either input or output. Since there will be a voltage drop along the pin due to the current, the correct way to measure the efficiency is to attach the probes closest to the power supply under test, and then attach the power connections further away.

Measuring Very Low Resistances

Exactly the same technique should be used to measure a very small resistance. For example, you may need to know the resistance of a PCB trace, or the winding resistance of a transformer. Although a four-wire ohmmeter could be used, these instruments are frequently not conveniently available in a lab. The best way to do this measurement is to pass 1A (or 10A, or whatever) through the part or trace being measured, and use a DVM to measure the voltage drop—but be sure to have the DVM probes <u>inside</u> the power connections, as shown in Figure 4.4.

Using a Shunt for *I* > 10A

Most DVMs can't measure a current greater than 10A. A good choice here is to use a current shunt, and measure the voltage with a DVM. Shunts are discussed briefly in the components chapter, Chapter 3. The shunt typically has separate leads for power connections versus sense connections: use them! Although shunts are usually 1% tolerance, you can achieve better accuracy by cross-calibrating the shunt together with the DVM you're going to be using in the measurement. That is, pass (say) exactly 1A

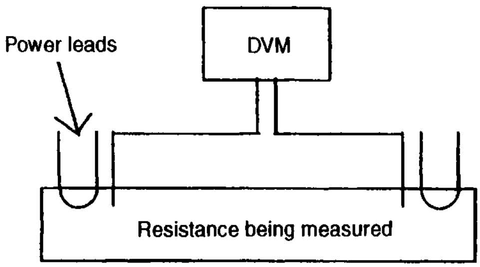

Figure 4.4 You can measure small resistances with a current source and a DVM.

through the shunt, and measure the voltage, calculating the resistance of the shunt, which amounts to the scale factor in this case. Just be sure to keep the shunt and DVM together as a pair thereafter. Otherwise you will lose the calibration.

How to Use a DVM to Measure a MOSFET

Well of course you're not going to be able to actually *characterize* the MOSFET with a DVM, but you can find out if it's "busted" or not. The trick is that on their resistance settings, mains-powered DVMs produce enough voltage to turn on the gate of a FET, at least a little. (Battery-powered DVMs produce a lower voltage, so this trick won't work with them, unless you have a logic-level MOSFET.)

To check the FET, first measure the resistance from drain to source. Here, "from" means putting the positive terminal of the DVM on the drain, and "to" means putting the negative terminal of the DVM on the source. The resistance should be something like 10MΩ or more. (If it's much lower, it's busted.) Now measure the resistance "from" the gate "to" the source: it should also be 10MΩ. Then, without letting the gate touch anything, measure the drain to source again. If the MOSFET is OK, you now read something like a few ohms up to maybe 2 or 3kΩ. If it reads open again, chances are the MOSFET is dead, since measuring the resistance of the gate to source should have applied enough gate voltage to turn it somewhat on.

Note: If in this last step (measuring drain–source after having measured gate–source) the measurement shows 20 or 30kΩ, the MOSFET is probably also busted, although in a weird way that doesn't show up well except in-circuit. You should try measuring its drain–source breakdown voltage with a curve tracer; generally, such a measurement with the DVM indicates that this FET too should be thrown away, although you might want to send it to Failure Analysis so they have something interesting to do.

ELECTRONIC LOADS

Why Is My Stable Converter Oscillating?

Electronic loads are convenient. You don't have to find that unusual resistor value in a high power package, and you don't need to find a fan to cool it either. But it's wise to recognize that these instruments have some significant limitations. Otherwise you may waste a lot of time looking in circuitry for problems that really lie in the instrumentation.

An electronic load is basically a bunch of power transistors in parallel, controlled by a feedback loop or loops, and operated in the linear region so that they are dissipative. The kind of control loop determines whether they appear as a resistive load, as a current sink, or as a constant voltage.

Regardless of mode of operation, electronic loads are always being controlled in a feedback loop, and of course this feedback loop has finite bandwidth: in typical units this bandwidth is somewhere in the vicinity of 1–5kHz. If your converter has more bandwidth than your electronic load, the load will not be able to make its transistors look resistive; as a result, your converter sees some unusual type of load impedance and may well start to oscillate. If your stable converter is oscillating, think about the load!

Practical Note During converter development, try to use resistive loads if possible; reserve the use of electronic loads for the production line, after the bugs have been worked out. If it's necessary to use an electronic load, at least try to hang the anticipated load capacitance on the load's inputs, to reduce the load's impedance at the high frequency end.

Minimum Input Voltage

Electronic loads typically require some minimum input voltage to operate properly—read the manual to find out what this is if you're going to use an electronic load on a rail of less than 5V (and they usually can't stand negative input voltages at all). Typically, this minimum is in that range of 2–3V. Some types of electronic load are especially pernicious in that they will appear to operate at less load voltage than their minimum but won't have the proper impedance characteristics. You can waste a lot of time trying to debug your converter with something like this.

OSCILLOSCOPES

Aliasing

If you don't know or aren't sure what "aliasing" means, try this little experiment. Set the output from a function generator producing a 100kHz sine wave and put it into your digital oscilloscope with the scope set on 10μs/div. You see 10 cycles of the sine wave. Now start increasing the sweep time, to 100μs/div, then 1ms/div. When you get to 10ms/div. you may suddenly see a nice sine wave again. (If not, try a few more turns of the sweep speed knob.) This effect occurs because the oscilloscope can display only a finite number of points, and if it displays for a long enough period, it may end up *looking* like something it isn't. (A textbook will explain the mathematical details.)

Aliasing happens when your oscilloscope either can't take enough samples per second to catch everything your signal is doing or, more likely for power supplies, doesn't have enough memory to store the entire waveform occurring in a sweep period. For example, if the scope has 1000 points per channel, and you're set at 10ms/div, with 10 divisions displayed, the sweep time is 10ms/div $\times$ 10 divisions = 100ms, and since it only has 1000 points, it can only display intervals of 100ms/1000 points = 100μs; anything happening in less than this time can be aliased. That is, you may miss the event entirely, or it may appear to be at a lower frequency than it really is.

The practical point here is that just because your digital oscilloscope *says* your signal is doing something, this isn't necessarily happening. If you don't know what to expect, you can be seriously fooled. This is particularly likely to happen when you have 60Hz pickup and you're triggered in such a way and at such a time base that you don't see anything at that low frequency; you try to figure out why the waveform is different when you check it twice! The best plan when looking at a brand new signal is to sweep over a very broad range of time bases to ensure that there isn't something unexpected happening. Or, be an old fuddy-duddy: try using an analog oscilloscope for the early stages of investigating a new signal; *it* won't alias.

NETWORK ANALYZER

A network analyzer is an instrument for measuring the response of a system to a sine wave; that is, it measures the transfer function by producing a sine wave whose frequency varies slowly with time, then measuring the magnitude and phase response of the system to this signal. Network analysis can be as simple as measuring the impedance of a capacitor as a function of frequency, or as complex as measuring the closed loop response of a converter.

A network analyzer is a must-have for the loop measurements we're going to be doing in the chapter on stability, Chapter 6. Since they are complex instruments, and not nearly as familiar as oscilloscopes, we're going to give detailed instructions for operating the HP3562A, a typical instrument in the medium price range. The HP3562A is not particularly convenient for measuring impedances, but it works fine for measuring loops, which is what the instructions will be targeted for. Other popular models (such as those of Venable Instruments) are quite similar in concept. Thus, although the detailed instructions regarding which buttons to push will differ, the general procedure will be the same.

Step-by-Step Instructions

In the 10-step instruction set that follows, the hardware, or hard, buttons labeled on them on the front panel are indicated by the capitalized word BUTTON, whereas the soft (i.e., software-generated) buttons that appear on the display screen in response to pressing a BUTTON are designated by the capitalized word SELECT.

1. The analyzer takes about a minute to warm up and perform its self-checks. The very first thing to do with the HP3562A is to push the BUTTON Cal, and SELECT Auto Off. This particular machine, unlike others, will otherwise calibrate itself without warning, often splat in the middle of your measurement. This can cause your supply being tested to explode, since calibration is accompanied by the generation of signals.
2. The next step is to select the measurement mode. To measure the loop response of a converter, you want the analyzer to produce a time-varying sine wave. Push the BUTTON Meas Mode, and SELECT Swept Sine. Other models may require you to select network analyzer mode, or gain–phase mode.
3. Next we set the frequency range over which the transfer function is going to be measured. A typical sort of range for a moderate bandwidth converter might be from 10Hz to 10kHz, three decades. We also have to set how quickly the sweep is going to be made. On this machine, a reasonable compromise between signal-to-noise ratio and operator patience might be 30 seconds per decade, for a sweep time of 1.5 minutes. You can also select how much averaging is to be done (how many times the same band is averaged before going on to the next frequency), but we'll not use this in this example. On fancier machines, you may be able to select directly the windowing bandwidth, which is accomplished in this machine with the sweep speed. Thus, press the BUTTON Freq and SELECT Start Freq. On the keypad, press the BUTTONS for 10, and then SELECT Hz. Next, SELECT Stop Freq, press the BUTTONS for 10, and SELECT kHz. Finally,

SELECT Sweep Rate, and use the BUTTONS Up/Down to get to about 30s/Dc.

If the plots look scraggly, which happens usually at the lower end of the frequency spectrum, you can sweep more slowly, and/or average; but of course either ploy increases the time needed for the measurement.

4. Next we set up the display. We're trying to get a Bode plot, so we need two traces: one for magnitude, the other for phase. Other machines may do this automatically in gain–phase mode. Start by selecting BUTTON A, for the A trace. Then push BUTTON Coord, and SELECT Mag(dB). Then push BUTTON Meas Disp, and SELECT Freq Resp, since we are sweeping frequency.
5. Now we can do the same for the second trace: select BUTTON B, push BUTTON Coord, and SELECT Phase. Then push BUTTON Meas Disp and SELECT Freq Resp.
6. To turn both traces on simultaneously, push BUTTON Active Trace A/B; this produces the magnitude trace on the upper half of the screen and the phase trace on the lower half, the usual display style for a Bode plot. Other machines produce a single screen with both gain and phase on a single plot, with two axis scales.
7. For convenience, it is helpful to let the machine control the vertical axes scales (gain and phase), rather than worrying about it yourself. Push BUTTON Scale and SELECT Y Auto Scale. If you have special display requirements (e.g., not showing what happens below 0° phase, as we do in the stability chapter, Chapter 6), you can turn off this autoscaling on either trace independently by pushing BUTTON A or BUTTON B, and then SELECTING Y Fixed Scale, and pushing BUTTONS on the keypad for the minimum and maximum display range; the analyzer clips any measurement outside the selected range.
8. Now we set up the characteristics of the signal output that's going to be used to drive the system. For a closed loop measurement, a level of about 100mVAC is suggested as a starting point. Push the BUTTON Source, and SELECT Source Level, then press the keypad BUTTONS for 100 and SELECT mV; or you can use the BUTTONS Up/Down, or use the dial to accomplish the same thing. This machine powers up with the DC Offset = 0, but for the open loop measurement, this can be changed if desired.

   The most common cause for scraggly looking plots at moderate frequencies is insufficient source drive level; try increasing it until you get a nice smooth plot. However, you need to observe the cautions below.
9. Finally, to turn on a cursor (useful not only for your measurement, but also for displaying information in a presentation), push the BUTTON X.
10. Some machines (although not the HP3562A) require you to set the input impedance and the input attenuation for each input channel. Input attenuation can be set to any value (usually either 0dB or 20dB) as long as the same one is used for both channels. An "input overload" warning that appears while the sweep is being run indicates that the input attenuation should be increased. The input impedance should be set to 1MΩ for both channels. Occasionally you will come across an instrument that has only a 50Ω input; this type CAN'T BE

USED for loop measurements, because the 50Ω load disturbs the loop components too much.

The network analyzer itself is now ready to go. But there are a few more things to know before you begin operation. The most important is to avoid thinking that the loop can be temporarily closed without injecting a signal by turning off or disconnecting the Source. **Don't do it!** With the mixer method (see the chapter on stability, Chapter 6) disconnecting the Source causes the output voltage of the mixer to double, causing the output of your supply to double! The correct way of removing the AC signal is to set it to 0V with the BUTTON Source. This caveat is irrelevant if you are using the transformer method.

Another step you'll certainly want to take is to monitor the gate drive signal with an oscilloscope while the analyzer is sweeping. Although the output of the analyzer is a constant amplitude, the response of the converter to this signal is not a constant; it happens not infrequently that what seemed a reasonable drive level at low frequencies causes the duty cycle to fluctuate wildly, and to go to zero at frequencies near the converter's bandwidth. *You have to avoid this*, because if the duty cycle ever goes to zero, the converter isn't operating, and you're no longer measuring the transfer function. This is the reason for watching the gate's duty cycle while sweeping, to ensure that the duty cycle is OK. A symptom of duty cycle collapse (if you're looking at someone else's measurement) is a sudden, discontinuous jump in the Bode plot.

Practical Note A good starting point for signal amplitude (source level) is to set the oscilloscope time base so that you're looking at a single cycle of the duty cycle (and trigger on the rising edge, as shown in Figures 4.5 and 4.6) and then increase the source level until you can just make out a little movement on the following edge. The figures show good and bad levels of drive. Figure 4.5 shows the gate drive of a converter being driven properly: the dither on the falling edge is caused by the modulation of the loop. Figure 4.6 shows the gate drive being overdriven: not only is the dither enormous, but the duty cycle goes down to 0% during a portion of the low frequency cycle.

With all of this set up, you can now turn the system on, in this order: mixer power supply, converter input power, and then analyzer sweep, by pushing the BUTTON Start. If you stop the sweep (by pushing the BUTTON Pause), reverse the order, by first turning off the converter input power and then the mixer's supply. Rigidly following this order prevents the converter from operating in open loop mode, where it might very well self-destruct, since the mixer is in the feedback loop, and it is of course an open when unpowered.

Nyquist Plots

The HP3562A can also be used to generate a Nyquist plot from the data. For this function, you need only press the BUTTON Coords, and SELECT Nyqust (sic). Since it is a single graph, it's best to turn on just A. Otherwise you will have two half-sized plots of the same thing. Not all analyzers have the ability to display Nyquist plots.

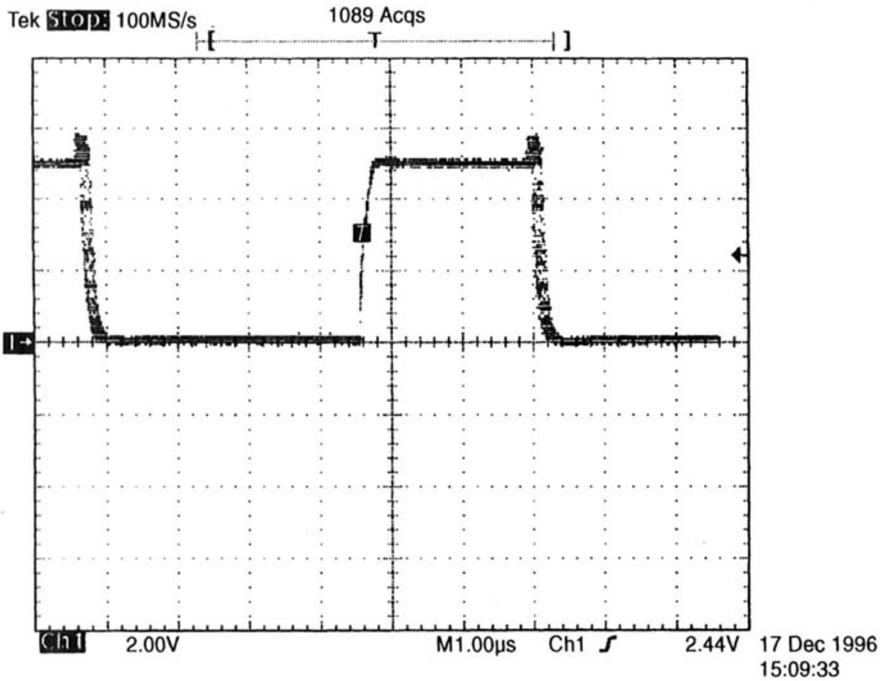

Figure 4.5 Gate waveform, showing proper network analyzer drive level.

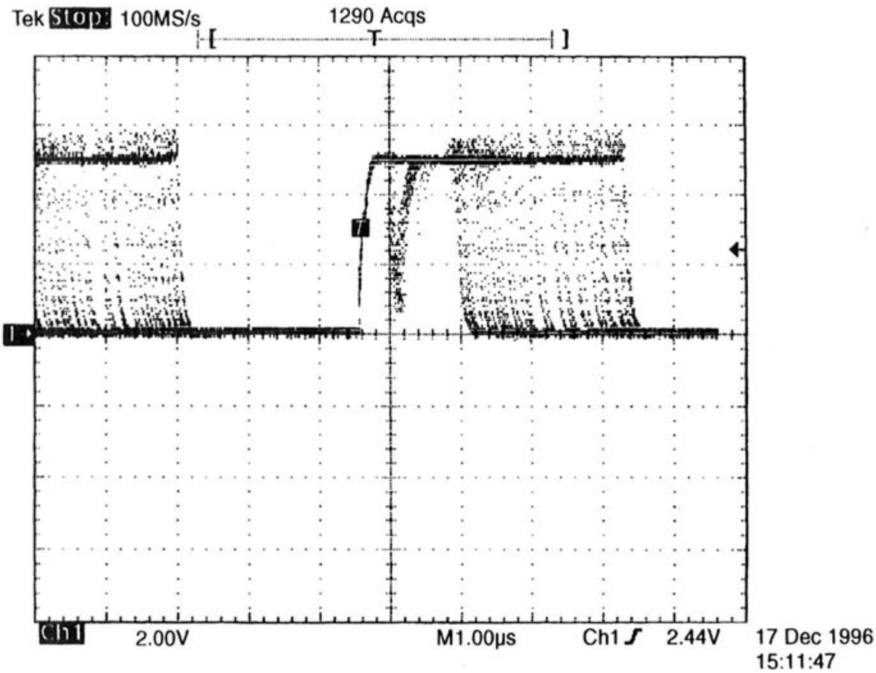

Figure 4.6 Gate drive waveform, showing too high a drive level.