

Hobbs ElectroOptics

AN-1: Photoreceiver

Testing System

Application Note

Our methods can help you too

Phil Hobbs

Rev. 0.1.2

AN-1: Photoreceiver Testing System

Rev. 0.1.2

Phil Hobbs

Hobbs ElectroOptics

Contents

Contents	i
1 Introduction	2
2 Quick Summary	3
2.1 Test Setups	3
2.1.1 Test Hardware	3
2.1.2 Control and Analysis Software	4
2.2 Pulse Testing	5
2.3 Frequency Response and Noise Testing	6
3 Time-Domain Measurements	9
3.1 Measuring The Pulse Response	9
3.1.1 Correcting Pulse Widths and Edges	9
3.2 Light Sources for Production Testing	9
3.2.1 The LP870 Nanosecond Light Source	10
3.3 Silicon Photomultiplier Module Testing	12
3.3.1 Silicon Photomultipliers	12
3.3.2 Testing the SP55B	13
4 Noise And Bandwidth Measurements	15
4.1 Example: Noise Floor of a 1-k Ω Transimpedance Amplifier	15
4.1.1 Measurement Setup	16
4.2 Lowering The Noise Floor	17
4.3 Dark Noise And Transimpedance Bandwidth Measurements	18
4.3.1 Transimpedance Bandwidth Measurement	20
5 Recommended Equipment	24
5.1 Electronic Gear	24
5.1.1 Digital Oscilloscope	24
5.1.2 Signal Generator	25
5.1.3 Data Acq Brick	25
6 Conclusion	26
6.1 Articles & Resources	26

1 Introduction

A perennial problem in building photoreceivers is how to test them, both in development and in production. It's generally pretty simple to find the gross problems, such as oscillation, power supply crap getting into the signal, tilted pulse tops due to high- or low-frequency peaking and rolloff, or bad DC drift. Before we go into production, we need to really characterize the design to make sure it meets all the requirements, and we also need a detailed manufacturing test plan so that the production units all meet spec too.

We care about both time and frequency domains: transimpedance gain, bandwidth, dark noise, stability, and rise / fall times, as well as overshoot and other pulse-top artifacts. But no worries, right? We've got a lab full of lasers, pulse generators, fiber modulators, LEDs, oscilloscopes, analyzers, and so on and so forth.

When you actually come to try it, though, doing stable, repeatable, quick, and generally well-behaved photoreceiver calibrations isn't necessarily that easy. Production test and field service have additional challenges.

In this app note, we discuss some good ways to do detailed testing, with an eye to getting accurate and unambiguous results, both in the lab and in production. without needing some giant dedicated test stand. Turns out to be pretty doable, actually.

We'll look at some strategies for testing photoreceivers, introduce the HEO setup, and show how it makes life a lot easier for not a lot of money. We'll talk about measuring noise floors and rise / fall times, and then about a neat way of measuring the transimpedance bandwidth and noise of a highly sensitive TIA, the QL01A.

2 Quick Summary

Testing folks are busy, so we'll start off with the gist, and then get to the fun details in Section 3.

2.1 Test Setups

2.1.1 Test Hardware

A solid basic setup for testing photoreceivers needs some apparatus:

The DUT The device under test, a photoreceiver.

A quiet power supply Prevents switcher noise and other crap from getting in and spoiling the measurement.

A shield box Nothing elaborate, just something to prevent light and EMI getting in. Connect power via feedthrough caps and signals via bulkhead coax connectors. In the lab, Danish Butter Cookies tins or clean, uncoated paint cans are good. For production, maybe something with a hinged lid and a jig to hold the DUT and light source in the right positions.

Digital oscilloscope One with way more bandwidth than you need, and FFT capability. One that can take at least 1000 waveforms per second will speed things up a lot. (It isn't the update rate that matters so much, it's the FFTs and averages per second, which usually aren't in the datasheet.)

Spectrum analyzer Optional, but pretty useful, especially for finding spurious signals from pickup or power supplies. Nothing fancy required. Extra credit for good response down at baseband.

Pulse generator For controlling the light source and triggering the scope. To reduce artifacts due to cable reflections, use one with a 50 Ω output impedance.

A clean, fast light source Such as the HEO LP870. Needs to be good enough that we're just measuring the DUT, not the setup—fast edges, flat pulse tops, very low overshoot, and zero flakiness. (See Section 3.2.1.)

A very quiet amplifier Such as the HEO LA22. With its high gain, excellent accuracy, and very low noise, it makes the noisy scope and analyzer quiet enough to ignore. (See Section 4.2.)

Good cables Cable reflections can cause 2-3% lumps and bumps on your pulses, making good measurements very difficult. Get good cables, and choose the lengths so the bumps don't mess up the measurement—lots of '50 Ω ' cables are really 52 or 53 Ω . Series-terminating the DUT helps a lot because it takes two bounces for the cable mismatch to show up.

2.1.2 Control and Analysis Software

To get into production, the test setup will also need a bit of software for instrument control, data acquisition, analysis, and report generation. It's actually pretty helpful in development too, because it saves a lot of fooling around, and you get all the data in a consistent format, including all the instrument settings so you can reproduce it if needed.

Most test instruments have a programmable interface that talks the [SCPI protocol](#). We have an HP 8566B spectrum analyzer (originally introduced in 1984) that talks an early predecessor of SCPI, and is happily controlled and read out over its GPIB port. More modern instruments tend to have USB, ethernet, or wifi connections. The data rates are faster and the cables skinnier, but the flavor is pretty much the same as 40 years ago.

Our typical photoreceiver test setup is made up of commercial test gear stitched together with a script, plus our additions such as the LP870 and LA22. That allows the entire process to run on its own. Here's some advice on doing this.

Use [PyVISA](#). "Control your instruments with Python." Great stuff for talking to newer equipment. PyVISA originally needed closed-source back-end modules, but there are

pure Python versions now too.

A fave of ours: Prologix GPIB-Ethernet Adapter On older devices that rely on GPIB, one of these adapters with [this library](#) gets you talking to them easily.

Keep the raw data It's often useful to suck in a big long scope trace and do the frequency analysis offline. Besides being more flexible, it keeps the raw data around, which is a win in itself: it lets you ask different questions afterwards. For instance, you can try different things on the same data set, such as applying different data windows before FFTing, and check for blunders, such as the DC level drifting during the measurement. (If all you have is a spectrum analyzer trace, you're stuck.)

Doing it yourself also means you know in detail what's going on, so you don't have to worry about the 2.5 dB problem,¹ for instance.

Use SciPy SciPy has a suite of powerful signal analysis tools, with examples available online. There's a good one on FFTs [here](#).

The `fft` (Fourier analysis), `signal` (signal processing), and `optimize` (root finding, optimization, and curve fitting) subsections are particularly useful for this type of analysis code.

We're happy to work with OEM and licensing customers on their testing setups—give us a call at the lab, +1-914-236-3005.

2.2 Pulse Testing

Testing rise and fall times is more familiar, so let's start with that.

Required: DUT, oscilloscope, pulse generator, and LP870 Nanosecond Light Source.

1. Connect the DUT to the scope

¹Section 4.3.1.1

2. Trigger the scope off the same generator output that drives the LP870. Use $1\text{ M}\Omega$ Z_{in} , and daisy-chain the cables to reduce reflections.
3. Connect the pulse generator to the LP870. Set it to make 0-3 V pulses about five times longer than the expected rise time of the DUT. Adjust the scope so that the pulse fills most of the display. This allows eyeball estimates of asymmetry.
4. Point the LP870 at the photodetector, and adjust the aim and separation to get a suitable pulse height.
5. Measure the rise and fall times, overshoot, and pulse-top aberrations if any. If the rise and fall times aren't $> 3\times$ those of the LP870, correct them by subtracting in RSS.²
6. Estimate the DUT's transimpedance bandwidth $B \approx 0.35/t_r \approx 0.35/t_f$ (10%–90%), or $B \approx 0.22/t_r \approx 0.22/t_f$ (20%–80%). For a single-pole response (simple real exponential) the 10%–90% times will be about $1.6\times$ longer than the 20%–80% ones.
7. Adjust the pulse width so it's long enough to see tilt, ringing, and other slower pulse-top artifacts. If they're caused by cables, very short or very long ones will help. (A 4-inch cable makes 1-ns bounces.)

2.3 Frequency Response and Noise Testing

Here we'll briefly show how to do the same sort of thing in the frequency domain, to measure the noise floor and transimpedance frequency response. The following procedure assumes you're using a scope FFT—it's easier with a spectrum analyzer. It also assumes that your photodetector has no internal gain, e.g. a PIN diode or an APD running way below breakdown.

Required: DUT, LP870 or incandescent-bulb flashlight, LA22 Lab Amplifier, scope or spectrum analyzer. (If your TIA gain is high enough, the scope may be okay by itself, apart from some sanity checking.)

²See Section 3.1.1.

Section 2: Quick Summary

1. Connect the LA22 between the DUT and scope, and power them up. With the DUT in the dark, adjust the vertical range to get 3-5 divisions p-p noise.
2. Set the sampling rate f_s to about $20\times$ your TIA's bandwidth or 250 MSa/s, whichever's greater.
3. Use enough samples per trace. Start with $N = 2500$ or so, and trade off resolution *vs.* speed to taste.³
4. Set up the FFT display: RMS V/ $\sqrt{\text{Hz}}$, ~ 20 averages during setup, ≥ 1000 for taking data. We need to make the grass on the baseline flat enough. (FFT first and average afterwards!)
5. Sanity check: Disconnect the cable from the LA22's input. Open-circuited and with a high- Z load, at 300 K it'll produce $422 \text{ nV}/\sqrt{\text{Hz}} \pm 1\%$ from 5 kHz to about 1 MHz. (Most of that comes from the $1 \text{ k}\Omega$ AC Z_{in} , so it goes like \sqrt{T} .) Try a few choices of window function (rectangular, Hamming, etc.) and check that the numbers are consistent.
6. Connect the DUT again. With it in the dark, measure the noise floor. Verify that the low frequency noise and $e_N C$ noise peak shape are what you expect.
7. Use the LP870 (with +3-5 V DC on its input) or a battery-powered incandescent-bulb flashlight to give the DUT an output of a bit over half scale. Measure its DC output, divide by the nominal transimpedance Z_m to get the photocurrent I_d , and compute the shot noise, $i_N = \sqrt{2eI_d}$, where e is the electron charge. (Shot noise in pA/ $\sqrt{\text{Hz}}$ is also given by $i_N \approx \sqrt{I_d/3.1 \mu\text{A}}$.)
8. Measure the noise floor, and verify that it's at least 10 dB above the dark noise in the bandwidth you care about. (Some fiddling with the vertical gain will be needed.) Divide by the calculated shot noise to get the transimpedance *vs.* frequency. It won't have the same shape as the dark noise.

³The frequency bins of an FFT are f_s/N wide, so 2500 points at 250 MSa/s gives 100 kHz bin widths. Actual resolution is poorer because of windowing and spectral leakage, so using 50k points is not absurd at all. This is where spectrum analyzers really win.

9. Automating this procedure is straightforward—just needs a DAQ brick and a relay or two for power switching, plus a bit of Python for control, data acq, and report generation. (See Section [2.1.2](#).)
10. You may not even need the scope; a well-calibrated filter such as our CF10 series plus a wideband true-RMS voltmeter is a time-tested way to measure noise in a given bandwidth, and that saves most of the software, besides reducing test time considerably.

3 Time-Domain Measurements

In Section 2.2 we did some rise and fall time measurements, and briefly discussed how to estimate the test stand's influence on the results. Let's dive a bit deeper into the details.

3.1 Measuring The Pulse Response

There are important applications where subnanosecond timing is vital, for instance counting photons in time-of-flight lidar and positron-emission tomography, where the timing uncertainty limits the range resolution. More often, though, we're interested in slightly longer-duration signals with higher SNR, so we need a suitable pulsed light source for testing. So what's 'suitable', exactly?

3.1.1 Correcting Pulse Widths and Edges

There's a theorem of Fourier transforms¹ that basically says that if we have pulses and responses with reasonable shapes, their widths add in RSS fashion. That is, if you're measuring a detector with a 10-ns rise time, a pulse with 5-ns edges will produce an output with $\sqrt{10^2 + 5^2}$ ns \approx 11 ns edges. So to get 10% accuracy, we don't necessarily need a super-narrow pulse—we need something that's known to be at least 2× quicker than our detector, and is stable, flexible, and convenient, with no bad habits.

3.2 Light Sources for Production Testing

In a production test setup or an instrument with built-in self test (BIST), the requirements are a bit stiffer: we want the source to also be compact, predictable, reliable, and economical, without fussy requirements.

In the lab, common approaches to this include a LED or diode laser connected to a pulse generator; a spark discharge, as from an automobile spark plug; a datacom transceiver such

¹Bracewell, R. N., *The Fourier Transform And Its Applications*, 3e, McGraw-Hill, 2000, p. 159

as a 1.2 Gb/s SFP module, driven by a pulser or data generator; or (if all else fails) the measurement setup the detector is to be used in. As when measuring noise floors, we want to be able to ignore the rise and fall times of our light source, or at least be able to correct for them accurately.

Of the usual alternatives, LEDs work OK but are often too slow; your random LED from the drawer will probably have 40-ns or slower rise and fall times. (They usually don't improve much if you try goosing the edges with a speed-up capacitor, either.) White LEDs are even slower, and also exhibit *drool*, where their edges start out fast but slow down dramatically before (eventually) settling, probably due to the fluor. Drooly pulses are useless for test sources.

Diode lasers are quicker; your average cheapo cleaved-cavity diode is about a 200-ps device, if you have a very low-inductance drive circuit and you can keep it from mode-hopping, dying from ESD, or getting its window crudded up. You need to watch out for interference fringes with diodes, as well. But as a calibration source, the main issue with lasers is that they tend to be noisy and unpredictable, unless you put more skilled engineering time into them than a one-off test source really deserves.

Datacom sources such as SFP modules are great for testing photoreceivers in the 100 MHz–1 GHz range, but the common ones are AC-coupled internally, so you get a lot of droop on the pulse tops at rep rates below 10 MHz or so.

3.2.1 The LP870 Nanosecond Light Source

So, back to the point: what we really want is a fast, trouble-free test source with low noise, clean edges, and no *agita*—something like the [LP870 Nanosecond Light Source](#). It uses a specially-designed LED, and as you can see from [Figure 3.2](#), it produces nice rectangular pulses with quick edges: rise/fall times of 3.75 ns/2.25 ns (10%–90%) and 2.5/1.25 ns (20%–80%). (The measurements were made with a prototype 1-GHz-class silicon APD photoreceiver. Look for a product like that soon.)



Figure 3.1: LP870 Nanosecond Light Source




Figure 3.2: Pulse response of the LP870, measured with a P400 pulser and a prototype fast APD photoreceiver. Measured rise/fall times are 3.75 ns/2.25 ns (10%–90%) and 2.5 / 1.25 ns (80%–20%).

3.3 Silicon Photomultiplier Module Testing

Among HEO’s most popular offerings are our custom and semicustom SiPM modules. We’re now offering these as standard products, beginning with the SP55B, which we of course have to test. Let’s start out with a look at SiPMs to see what we should expect.

3.3.1 Silicon Photomultipliers

Silicon photomultipliers (SiPMs) are more or less what they sound like—silicon photodetectors whose gain can be varied over a wide range (like 10^0 – $10^7\times$) by changing their bias voltage. However, their detailed properties are different enough from a photomultiplier tube’s that they’re worth a word or two.

SiPMs are actually arrays of thousands of tiny elements connected in parallel. Each element is made of a single-photon avalanche diode (SPAD) in series with its own quench resistor. Their low-gain and bright-light behavior can be complicated, but at higher gains in dimmer light, it’s pretty simple.

A photon detection event knocks an electron loose in some SPAD. It accelerates rapidly under the strong applied \mathbf{E} field, knocking other electrons loose and forming an avalanche. The avalanche discharges the SPAD’s capacitance until \mathbf{E} drops enough that the avalanche dies out. Thus each detection event dumps a teacup’s worth of charge, and the gain equals the number of electrons in the cupful.

Since the diodes are fully depleted in the normal operating voltage range, their capacitance is very nearly constant. The endpoint voltage doesn’t vary much, because the avalanche dies out very rapidly once the impact ionization rate drops below the recombination rate. Thus the SiPM’s gain goes linearly with the overvoltage in this region, because $q = CV$.

After each detection event, the element recharges via its built-in quench resistor until it’s ready for another one. The recharge current is what’s visible to the external circuit, and it follows the normal decaying exponential RC curve, usually with $\tau = 5$ ns–60 ns, depending on the device. The 90%–10% fall time is 2.2τ as usual, say 10–130 ns.

The rise time of this current is set by the speed of the avalanche and by circuit strays, so it's fast: normally under 1 ns, and often less than 300 ps. (Whether the TIA is fast enough to reproduce that is another matter entirely.)

You only see that super-sharp edge with very brief pulses, however. A nonzero-width light pulse creates avalanches in many elements, which overlap in time, so that the result is the convolution of the light pulse with the SiPM's recovery response. When the pulse is longer than a few times τ , you get symmetrical exponential edges, just as with a normal one-pole RC lowpass filter. (The RC impulse response has the same shape as the SiPM's, of course, so this should be no surprise.)

3.3.2 Testing the SP55B

Recently we wanted to test the edge response of our forthcoming SP55B Silicon Photomultiplier Module which is a medium-speed, high sensitivity photoreceiver that's well-behaved at gains up to 7.5 million, with low dark count rate and good detection probability. Its recovery time constant is about 50 ns, so we expect 10%-90% rise times around 110 ns.

We followed the work flow in Section 2.2, and got the results shown in Figure 3.3.

The edges are very nice looking, and the speed is close to what we expected, with a 100-ns rise and 116-ns fall. The 100-ns edges are from the SP55B alone: the worst-case contribution of the LP870 is only

$$\Delta t_r = \sqrt{(100 \text{ ns})^2 + (3.75 \text{ ns})^2} - 100 \text{ ns} = 0.07 \text{ ns}, \quad (3.1)$$

so that 16-ns asymmetry is real.




Figure 3.3: SP55B SiPM Photoreceiver Pulse Response

4 Noise And Bandwidth Measurements

From this point, we'll be going deeper into explaining what's going on.

4.1 Example: Noise Floor of a 1-kΩ Transimpedance Amplifier

A typical TIA for brighter-light situations might be just a photodiode, low-noise op amp, and 1-kΩ feedback resistor, as shown in Figure 4.1.

Ideally we want the noise of our TIA to be limited by the Johnson (thermal) noise of the feedback resistor when there isn't any light, and by the photocurrent shot noise when there is.

A resistor has an open-circuit noise voltage density

$$e_N = \sqrt{4kTR}. \tag{4.1}$$

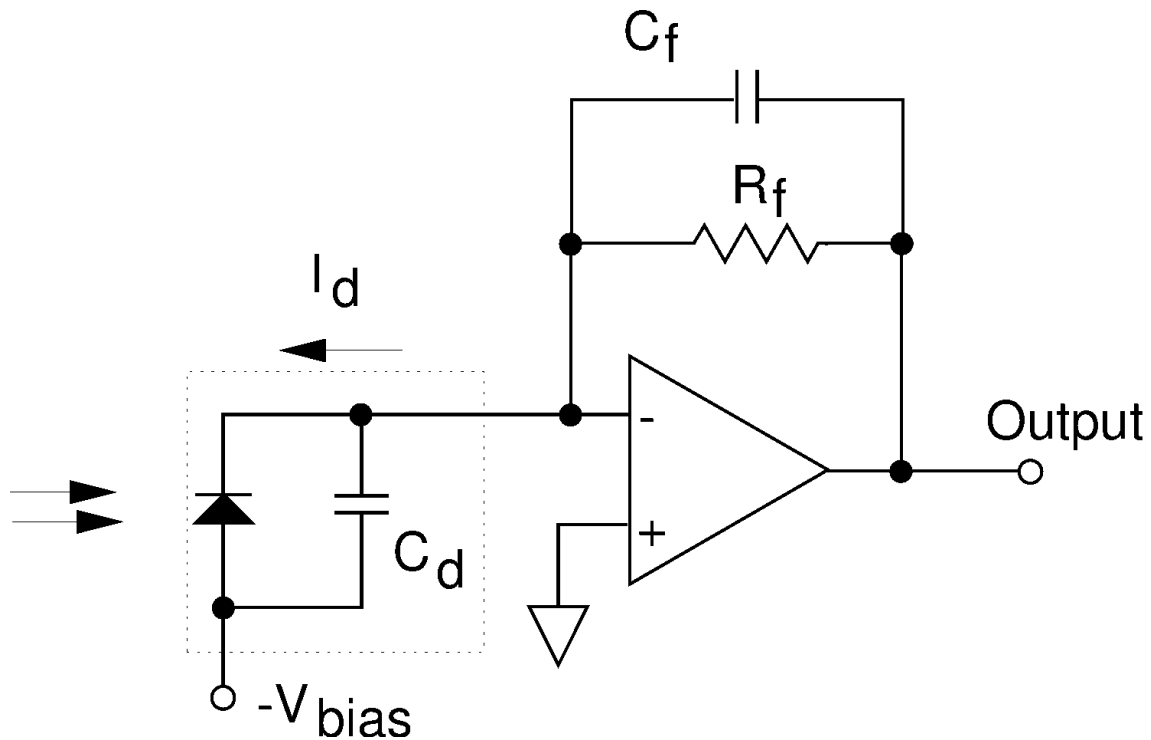


Figure 4.1: Generic op amp TIA

At 300 K, this is

$$e_N \approx \sqrt{R/60.4\Omega} \text{ nV}/\sqrt{\text{Hz}}. \quad (4.2)$$

The 1-k Ω feedback resistor thus has a Johnson noise of 4 nV/ $\sqrt{\text{Hz}}$, so we need a quiet op amp. A suitable one is the OPA211, whose flatband noise is 1.1 nV and 1.7 pA in a 1-Hz bandwidth.

Since the three noise contributions are uncorrelated, we can calculate the TIA's output noise using the usual root-sum-squares (RSS) formula:

$$e_N = \sqrt{(1.1 \text{ nV})^2/\text{Hz} + (1 \text{ k}\Omega \cdot 1.74 \text{ pA})^2/\text{Hz} + (4 \text{ nV})^2/\text{Hz}} = 4.6 \text{ nV}/\sqrt{\text{Hz}}. \quad (4.3)$$

That's 1.2 dB above the Johnson noise alone, not too bad. (We could equally well put the resistor's contribution in as a current noise of 4 pA/ $\sqrt{\text{Hz}}$, but [unlike the op amp] the resistor has only the one noise source, so we only include it once.) So how do we test this thing?

4.1.1 Measurement Setup

If we have a nice baseband spectrum analyzer such as an HP 89441A, it will usually have a high-impedance input available. However, the 89441A has a flatband noise of about 8 nV/ $\sqrt{\text{Hz}}$, making accurate noise floor measurements difficult. With our 4.6 nV/ $\sqrt{\text{Hz}}$ amplifier, we can calculate the indicated noise using (4.3): it would be about $\sqrt{8^2 + 4.6^2} = 9.2$ nV/ $\sqrt{\text{Hz}}$, which is 6 dB too high. When we turn on the TIA, the only visible effect on the analyzer screen is a tiny (1.2-dB) shift of the noisy baseline, from 8 to 9.2 nV. That's too small to see unless we're doing many averages to flatten out the grass on the baseline.

In principle, we could measure the instrument noise and subtract it in RSS, but we don't have much to work with here, just that 1.2 dB baseline shift. We'll be averaging for quite awhile to get that working, if it works at all—and how can we be sure it does? Corrections like that are easy to make if they're small, and almost impossible if they aren't.




Figure 4.2: LA22 Lab Amplifier

It gets harder quadratically as we go to lower noise levels. If the TIA's noise is $10\times$ lower than the analyzer's, all we'd get would be a 0.5% shift in the noise floor: $\sqrt{1 + 0.1^2} = 1.0050$ (0.043 dB). So a barefoot analyzer isn't a great fit for this measurement. What to do?

4.2 Lowering The Noise Floor

At HEO, we've run into such problems more than enough times. Thus we built ourselves an amplifier that's good for just about anything of that sort: the [LA22 Lab Amp](#), shown in Figure 4.2. The LA22 has an accurately calibrated gain of $100 \pm 0.5\%$, and a flatband noise of $1.1 \text{ nV}/\sqrt{\text{Hz}}$.

If we put that between the TIA and analyzer, we effectively get an analyzer with 1.1 nV noise in 1 Hz , so that our TIA's 4.6 nV becomes a measured 4.72 nV . The LA22 + analyzer adds less than 3% (0.24 dB) to the TIA's noise floor, which is easy to remove accurately—that quadratic is working in our favor now.

RF spectrum analyzers have $50\text{-}\Omega$ inputs, and are usually somewhat noisier than baseband ones, but that hardly matters: the LA22's high gain and ability to drive $50 \text{ }\Omega$ means that

we'll get an excellent measurement all the same.

This technique works with oscilloscope FFTs, too. A scope's 8-bit digitizer doesn't have much dynamic range—only 50 dB or so—but that's enough to do a good job on a flattish noise spectrum, once we override the input noise with the 40-dB gain of the LA22.

You can see the LA22's good manners in Figures 4.3 and 4.4. *Note:* These plots were made with a 12.5- Ω source resistance, which contributes 0.45 nV/ $\sqrt{\text{Hz}}$ of Johnson noise. Figure 4.3 shows the noise floor of the combination as 1.20 nV, but the noise of the LA22 itself is actually

$$e_N = \sqrt{(1.20 \text{ nV})^2/\text{Hz} - (0.45 \text{ nV})^2/\text{Hz}} = 1.11 \text{ nV}/\sqrt{\text{Hz}} \quad (4.4)$$

(We care about that 0.7 dB.)

Figure 4.4 shows the LA22's pulse response, measured with the same setup—that 12.5 Ω was a 20-dB BNC attenuator followed by three 50- Ω terminators in parallel, allowing us to apply a good-quality step to its input using a Highland Technologies P400 Digital Delay Generator, which makes fast, clean, jitter-free pulses.

So overall, the LA22 Lab Amp makes it easy to do low-level measurements such as photoreceiver characterization, using ordinary lab equipment. We use them all the time.

4.3 Dark Noise And Transimpedance Bandwidth Measurements

Measuring the bandwidth and frequency response of a circuit is generally straightforward. Often we'll attach a swept sine generator and look at the envelope of the sine wave output with an oscilloscope. It's not so easy with photoreceivers, because there's no good place to attach the sine generator—it needs to be optical, for one thing, and for another, the photoreceiver's circuitry may not be easy to probe. The QL01 Nanowatt Photoreceiver uses a 10-M Ω feedback resistor, and has a bandwidth of a bit over 1 MHz, so the effective capacitance across the resistor is only about 0.013 pF. It's tough to get a probe on that without screwing it up, so we aren't going to be injecting any sine waves there.

Section 4: Noise And Bandwidth Measurements




Figure 4.3: Noise floor of the LA22 Lab Amplifier with a 12.5- Ω source.




Figure 4.4: Time-domain response of the LA22, showing clean, symmetrical 20-ns edges.

Using an optical source would be easier, but getting one that can be modulated cleanly over a wide range is surprisingly hard.

4.3.1 Transimpedance Bandwidth Measurement

Fortunately, there’s a very useful hack: we can just use the shot noise.

For instance, take whatever setup you have and shine a battery-powered, incandescent-bulb flashlight on it.¹ Place the flashlight so as to get a photocurrent a bit below full scale. This will give you a beautifully calibrated white noise source (the shot noise of the photocurrent) of spectral density $i_N(1\text{Hz}) = \sqrt{2eI_d}$, delivered right to your photodiode.

Assuming that your TIA is quiet enough to be shot-noise limited, measuring the noise level in the dark (circuit noise) and with the flashlight (circuit noise + shot noise) will give you the gain vs frequency and noise vs frequency of your measurement system. This will also allow you to compare the noise of other light sources (such as lasers) to the shot noise. That’s super helpful in letting you know how you’re doing.

¹Yes, the LP870 will work very well for this, too—just connect its input to +3-5 V and move it to a convenient distance.

In photoreceivers, the shape of the dark noise floor and the transimpedance are generally quite different. On account of the $e_N C$ noise current (amplifier voltage noise differentiated by the photodiode capacitance), there's a tendency for the dark noise to rise with frequency right out to the limit set by the amp's gain-bandwidth product (GBW). The noise gain is the magnitude of the noninverting gain of the stage, which depends almost entirely on the ratios of the input and feedback impedances, more or less independent of their magnitude. On the other hand, the transimpedance depends on both the GBW and the magnitude of the feedback impedance, but not so much on the noise gain.

For instance, here's a noise floor plot of one version of our QL01 Nanowatt Photoreceiver for various photocurrents between 0 and $1\mu\text{A}$, together with the idealized values. The bottom trace shows a pronounced $e_N C$ noise peak, which gives it a shape quite different from the shot noise floor of the brighter-light traces.

At very low frequency, the dark noise nearly reaches the Johnson noise of the 10-M Ω feedback resistor; the small offset is due to the shot noise of the ~ 2 nA photodiode leakage. Above about 200 nA, the photocurrent shot noise dominates completely so the shape of the noise floor gives the transimpedance gain curve of the TIA + photodiode combination. At intermediate currents, you can see the noise floor gradually shift between the two.

4.3.1.1 Power Gain Measurement

The beauty of this is that if we're using a quiet light source such as an incandescent flashlight or a LED driven from a very quiet power supply, the photocurrent will have exactly full shot noise.² The shot noise PSD is constant with frequency, and its RMS amplitude is given exactly by the shot noise formula. There's no other optical measurement where getting an excellent calibration is so easy—measure the DC output voltage with a multimeter, divide by the feedback resistance to get the photocurrent I_d , plug it into the shot noise formula, and you know the additive noise current to four significant figures.

²There are very special situations in which this is not quite true, but they sure don't happen by accident.

Section 4: Noise And Bandwidth Measurements




Figure 4.5: Noise floor of the QL01 Nanowatt Photoreceiver for various photocurrents from 0 to 1 μA .

Having a known, flat input noise spectrum that's well above the dark noise means that we can use a spectrum analyzer or scope FFT, average awhile to get the grassy baseline to smooth out a bit, and then compute our transimpedance gain vs. frequency by dividing the output noise voltage density by the shot noise,

$$|Z_m| = e_{Nout} / \sqrt{2eI_d}. \quad (4.5)$$

Doesn't get much slicker than that.³

Sharp-eyed readers may have noticed that in Figure 4.5, the low-current noise curves are above the expected shot-noise value, but the 200 nA and 1 μ A curves look right on the money, within a small fraction of a decibel. A still sharper-eyed one might guess that the 3-dB bandwidth of the 200 nA curve looks slightly wider. The reasons for this are interesting.

At low frequency, the dark noise is 22 dB below the 1- μ A shot noise and 15 dB below the 200-nA shot noise. It thus lifts the curves by $10 \log(1 + 10^{-2.2}) = 0.027$ dB and $10 \log(1 + 10^{-1.5}) = 0.14$ dB, respectively. At higher frequencies, the $e_N C$ contribution starts to lift the dark noise; at 1 MHz, dark noise contributes 0.3 dB at 1 μ A and 1.5 dB at 200 nA. That's why the 200 nA curve looks a little bit fatter.⁴

That's an example of how useful the shot noise test is, at least with a quiet photoreceiver.

Of course, there are a few assumptions involved that need to be verified. The main one is linearity; if the transimpedance isn't the same at 1 μ A as at zero, life gets more complicated. Compression of the photodiode responsivity at high current (due to enhanced recombination or local forward-biasing of the diode junction) is one way this can happen, but it's unlikely to be an issue with a reverse-biased photodiode at only 1 μ A. Another thing to watch out for is $1/f$ noise from a thick-film feedback resistor—if your light source is quiet, it won't have any $1/f$ noise to speak of.

³Remember to check that your analyzer or scope is computing the RMS noise floor correctly. Many compute the average of the logs, rather than the log of the average, which makes them read 2.5 dB too low. See Keysight AN-150, P. 82ff in the reference list below.

⁴In order to correct the lower-current plots, it's helpful to go all the way back to nanovolts in 1 Hz, because it's easy to get the denominator of the fraction wrong when the total noise isn't close to the shot noise.

5 Recommended Equipment

We talked about the general setup in Section 2.1. Here’s a bit more on what you should be looking for.

5.1 Electronic Gear

Here at HEO, we’re big fans of boat anchors, *i.e.* 20- to 30-year old top-of-the line equipment, usually intended for rack mounting. It’s still as good as it ever was, and mostly sells for a few cents on the dollar at auction or on eBay. You can pull together an amazing amount of capability on the cheap, but you do have to know what you’re looking for.

For production test, though, you want inexpensive stuff that’s going to be in production for awhile so you can get it calibrated regularly and maintain spares to avoid downtime. We assume that you have things like a good true-RMS multimeter and so on.

5.1.1 Digital Oscilloscope

Next you’ll need a digital oscilloscope with more bandwidth than you think you’ll ever use. 100 MHz, 1 GSa/s is the slowest you should consider. You’ll want one with FFT, waveform math, and delayed sweep. Two channels is probably enough to start with, but for more complicated measurements, four channels make life a lot easier.¹

Especially for FFT measurements, it’s vitally important that your scope have a fast enough sampling rate that you don’t alias high-frequency noise down into the band you care about. Due to e_{NC} noise peaking, there’s often quite a bit more of it out there, so aliasing is no joke.

There’s lots of competition in oscilloscopes these days. Whichever one you pick, make sure it has a 1-mV or 2-mV per division scale that works up to the maximum bandwidth—some cheap noisy ones limit the bandwidth to 20 MHz on the most sensitive scale.

¹A seldom-discussed advantage is that if you blow up one channel, you still have three working ones—a single-input scope is a lot less useful. (Hey, it happens.)

There are a couple of seldom-considered issues with digital scopes that you should think carefully about: its number-crunching oomph, specifically how fast it can compute and average medium-length FFTs (1k–50k points or so), and how bad the jitter is on the external trigger input.

The FFT averaging rate can easily range 1000:1 among scopes in the same class, depending mostly on processing power, and it isn't necessarily the same as the waveform update rate. That matters when you're averaging many traces, especially with FFTs.

The trigger jitter is a bit more subtle. Many inexpensive digital scopes can trigger with jitter well below 1 period of the sampling clock, but only on the vertical channel inputs; if you use the external trigger input, they'll only trigger on a clock edge, giving a minimum jitter of 1 clock p-p.

5.1.2 Signal Generator

Besides a scope, you need a signal source. There are cheap-and-nasty options, but for real use, a decent pulse/function generator will save time and improve your life. For purposes of this app note, any vaguely decent pulse generator should be fine.

5.1.3 Data Acq Brick

For test setups, it's often convenient to get a basic data acquisition and control brick that speaks USB and is cheap enough that it can be left connected to the test jig. We've recently been using the LabJack U3 for this with our bed-of-nails test stands. Signal integrity on these setups is rarely world class, so we primarily use these for DC and near-DC measurements. The digital output of a LabJack or similar can also be used to trigger the LP870 if a function generator isn't used. (The LP870's voltage gain cleans up the pulse edges, but watch out for cable reflections and jitter if you do this.)

6 Conclusion

We've gone over the main points of testing photoreceivers using HEO's LA22 Lab Amplifier and LP870 Nanosecond Light Source, and shown how you can get good measurements quickly and easily. The LP870's clean, fast light pulses plus a decent oscilloscope solve the time-domain problems pretty handily, so that what you see on the scope is the performance of your instrument, not the test stand. The LA22 plus a scope or spectrum analyzer does the same in the frequency domain, and the *ju-jitsu* trick of using the shot noise as a perfectly-calibrated signal source rounds out the ensemble by measuring the transimpedance gain *vs.* frequency.

At Hobbs ElectroOptics / ElectroOptical Innovations, we're always happy to do customized versions of our products, as well as fully custom instrument designs. You can see some of what we've done at our sister website, [ElectroOptical Innovations LLC](#). We look forward to hearing about your latest application.

6.1 Articles & Resources

- [Keysight AN-150: Spectrum Analysis Basics](#)
- [Tektronix XYZs of Oscilloscopes Primer](#) or a less dumbed-down [1993 version](#).
- [Rigol FFT](#)
- [Rigol Pass/Fail Testing](#)
- [EDN FFTs and Oscilloscopes a Practical Guide](#)
- [Keysight 8 Hints For Better Spectrum Analyzer Measurements](#)
- [EDN Pass Fail Testing Using an Oscilloscopes Signal Processing Capabilities](#)
- [EDN Perform Pass Fail Tests With an Oscilloscope](#)
- [Rigol Oscilloscope Remote Control VISA SCPI](#)
- [EDN Analyze Noise With Time Frequency and Statistics](#)
- [EDN Oscilloscope Rise Time and Noise Explained](#)

Section 6: Conclusion

- [EDN Measure Frequency Response on an Oscilloscope](#)
- [EDN 10 Tricks That Extend Oscilloscope Usefulness](#)
- [EDN Measuring Small Signals Accurately a Practical Guide](#)

Hobbs ElectroOptics
<https://hobbs-eo.com>
info@hobbs-eo.com

160 North State Road, Suite 203
Briarcliff Manor, NY 10510
(914) 236-3005