Tutorial

Optical Meters and Detectors

Photodiode Basics

When a photon hits the photodiode material it may generate an electron-hole pair, depending on the quantum efficiency of the device. Quantum efficiency is dependent on many factors, but in general if the energy of the photon, $E = h\nu$, is less than the energy gap of the device, then the photon will pass through the device without being absorbed. On the other hand, if the energy of the photon is greater than the energy gap of the device, these photons will be absorbed very near the surface where the recombination rate is high and will not contribute to the photocurrent. It is the quantum efficiency that is responsible for the wavelength dependency of the photodiode's spectral response. Semiconductor materials such as silicon and InGaAs possess different energy gaps; consequently they exhibit different quantum efficiencies at different wavelengths, resulting in spectral responsivity profiles unique to the specific material type.

Semiconductor photodiodes are ideal for making measurements of low-level light due to their high sensitivity and low noise characteristics. Most photodiode manufacturers specifically design their diodes to be used in either the photoconductive (reverse biased) or the photovoltaic (no bias) mode, both having advantages and disadvantages. Newport's low-power **818 and 918 Series** photodiodes are used in the *photovoltaic* mode to take advantage of the reduced noise performance.

The two primary noise sources from the diode alone are Johnson noise and shot noise. In the photovoltaic mode with no light striking the photodiode surface, the photodiode is in thermal equilibrium producing random thermal noise known as Johnson current noise, given by

$$I_{\text{Johnson}} = \sqrt{\frac{4kTB}{R_{\text{sh}}}} \quad [A],$$

where k is Boltzman's constant, T is the temperature in Kelvin, B is the bandwidth of the detector/amplifier, and R_{sh} is the shunt resistance of the photodiode. It can also be seen from this equation that a photodiode with a high shunt resistance is desired to reduce the Johnson noise.

Shot noise is the noise produced by the flow of current in the diode and is given by,

$$I_{\text{shot}} = \sqrt{2qB\left(I_{\text{dark}} + I_{\text{photo}}\right)} \quad [A],$$

where q is the charge of an electron, I_{dark} is the dark current, and I_{photo} is the photocurrent. When a photodiode is used in the *photovoltaic* mode the voltage across the diode is kept at zero volts. Consequently this almost eliminates the dark current altogether. Because there is negligible dark current, the shot noise contributed by the dark current is also eliminated. To put these effects in perspective, if a detector were biased as in the photoconductive mode, the dark current would be about three decades larger than the noise equivalent current of an unbiased detector.

The photocurrent produced by the photodiode is measured directly by the power meter using an operational amplifier circuit known as a *transimpedance amplifier*. Typically measurements can be made down to the sub-picoampere regime with good reproducibility, even at room temperatures. An exception to this rule is when the shunt resistance of the photodiode is small as with the Germanium photodiode (818-IR). Because of its low shunt resistance (50 kΩ typical), tens of picoamperes can be resolved at best.

Thermopile Basics

A thermopile consists of an array of thermocouples connected together in a series in order to increase the voltage output to more easily measured levels (millivolts vs. microvolts). It may be formed by a dissimilar metal junction, or from semiconductor (Peltier) junctions. In use, absorption of a laser beam creates heat at the receiver area of the power probe. The amount of heat, and thus the temperature rise at the receiver, is determined by the power level of the laser beam. The heat flows to either an air- or water-cooled radiator which is nearly at a constant ambient temperature (unless the probe power rating is greatly exceeded!). The radiator will be referred to as a heat sink. The thermopile *hot* junctions are located at or near the laser receiver, while the *cold* junctions are located at the heat sink. The probe's voltage output results from the temperature rise between the hot and cold junctions produced by the absorbed laser energy. The two types of thermopile probes in use today are distinguished by using *radial* or *axial* heat flow paths from receiver to heat sink.

Probes constructed with dissimilar metal junctions use a *radial flow* heat path from the center to the outside of a thermal disk assembly. The heat sink is attached to the outside circumference of a metal disk which has the laser absorbing element attached (or coated) in its central area. This approach gives a response time on the order of one second to register changes of laser input power. It is also easy to scale up to kilowatt power levels by using a thicker disk and/or going to copper instead of aluminum as the disk material. Dissimilar metals are vacuum deposited on the back of the disk in a pattern such that the hot junctions form a closed ring around the laser absorbing area, and an equal number of cold junctions are located around the outside of the disk. If thermopile junctions, thermal disk and mechanical interface to the heat sink are uniform, the voltage responsivity (mV/W) is near constant regardless of where the laser is positioned on the disk. In some older thermal disk designs, non-uniformity was inherent, thereby forcing the user to carefully center the laser beam for any measurement.

Semiconductor based thermopile probes employ a more traditional *axial flow path* for absorbed heat. In this case a layered construction is used to assemble the probe. The heat sink is positioned behind a disc-shaped laser absorbing element, and Peltier junctions are sandwiched between the absorber and heat sink. The junctions are tiled together to match the size of the absorber. The heat flows straight back (axially) from the absorber as opposed to a radial (outward) direction. This has the advantage of scaling easily to very large absorber areas. However, a slow time constant (on the order of 30 seconds) results from the thermal impedance of the junction assemblies.

Thermopile Basics section authored by Dr. J. Buck and B. Mooney.

Pyroelectric Basics

Pyroelectric detectors are designed to measure the energy of short optical pulses that have a ***maximum*** width of 5 µs–400 µs, depending on the detector design. These detectors are made of a ferroelectric crystal which has a permanent dipole moment. When subjected to an optical pulse, the crystal is heated and causes the dipole moment to change. The changing of this dipole moment causes a current to flow, which is converted to a voltage in the detector head that can be measured by the optical power meter or oscilloscope.

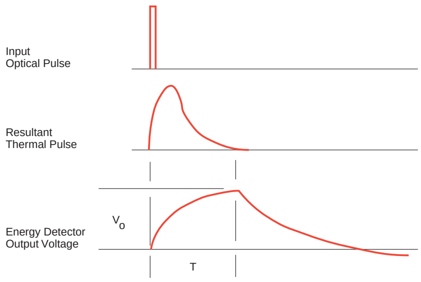

Figure 1—Typical signal behavior of a Pyroelectric detector is shown above.

As shown in Figure 1, the resultant thermal pulse is broadened relative to the short optical pulse. During this thermal pulse the current flows through the ferroelectric crystal, creating a voltage that increases in amplitude. The optical power meter has circuitry that measures the difference in voltage between when the output voltage just starts to increase and when the output voltage reaches its peak amplitude. This voltage difference is then numerically multiplied by the detectors responsivity, which is in units of Joule/Volt, resulting in the energy of the pulse in units of Joules.

Most of Newport's pyroelectric detectors require termination into a 1 MΩ input impedance, however, some have built-in amplifiers that require a 50 Ω input impedance. Our optical meters that use a calibration module provide the correct impedance automatically, but care must be taken to select the correct impedance when using an oscilloscope.

When using pyroelectrics, care must be taken not to exceed the maximum pulse width or the maximum repetition rate. If either of these specifications is exceeded, your measurement accuracy will degrade due to the electrical bandwidth limitation of the detector.

Photodiode Spectral Calibration

Newport Corporation performs its photodiode spectral calibration at their main headquarters in Irvine, CA. The calibrations are done using a double monochromator in order to minimize stray optical noise, especially in the ultraviolet. Three gratings and two light sources are used by the monochromator to maximize the signal to noise performance over the 190 nm–1800 nm wavelength ranges. A deuterium lamp is used in the ultraviolet range up to 310 nm and a tungsten lamp is used thereafter in the visible and near infrared.

Newport uses two standard detectors that are sent to NIST for calibration on an annual basis. One of the standard detectors is used for the wavelengths between 190–1100 nm and the other for 780–1800 nm. The absolute responsivity accuracy of NIST's standard detectors is based on a cryogenic radiometer which has a relative expanded uncertainty (k=2) to absolute SI units of 0.2%.

Prior to calibrating a manufacturing lot of detectors, the optical flux from the monochromator is measured in 10 nm steps using the NIST traceable standard detector throughout the wavelength range in which the detector-under-test (DUT) is to be calibrated. Since we know from NIST the responsivity of the standard detector, we calculate the optical flux of the monochromator using the following relationship:

$\mathrm{Flux}_{\mathrm{mono}}$ ([W] = i_{measured} [A] / $\mathrm{Responsivity}_{\mathrm{Std\ Det}}$ ([A/W]),

where i_{measured} is the measured current of the standard detector.

Knowing the flux coming from the monochromator, we measure the photocurrent of the DUT in 10 nm steps and divide this current by the monochromator flux to get the spectral responsivity of the detector in units of A/W.

Because the responsivity of a photodiode is temperature sensitive, especially near the ends of its usable wavelength range, we maintain the temperature of the standard detector and DUTs at the temperature where NIST calibrated the standard detector. This temperature control is critical for an accurate calibration. For example, silicon's temperature dependency induces a responsivity change of approximately 10% at 1100 nm with a 5°C change near room temperature.

After a manufacturing lot of detectors has been calibrated, we perform a single-wavelength check on ***every*** detector, with and without the OD3 attenuator, to ensure that nothing went wrong with the calibration. This check is done using a different working standard and test equipment in order to isolate any systematic problems.

Power Meter Basics

Although most people want to make measurements in units of dBm or Watts, an optical power meter is only capable of measuring either the current or the voltage generated by a photodetector.

When interfacing with a photodiode, the quantity that must be measured is current. There are numerous techniques in measuring this current, but only one will yield the detectivity, signal-to-noise, and accuracy that is expected from a semiconductor photodiode. A circuit known as a transimpedance amplifier is the circuit of choice when using a photodiode (Figure 2).

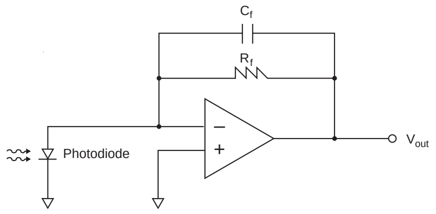

Figure 2—Transimpedance Amplifier

The advantage that the transimpedance amplifier has over almost any other amplifier configuration is that it does not bias the photodiode with a voltage as the current starts to flow from the photodiode. Typically one lead of the photodiode is tied to the ground and the other lead is kept at virtual ground by means of the minus input of the transimpedance amplifier. The resultant bias across the photodiode is then kept at virtually zero volts, a condition that helps minimize dark current and noise, and helps increase linearity and detectivity.

Effectively the transimpedance amplifier causes the photocurrent to flow through the feedback resistor which creates a voltage, V = iR, at the output of the amplifier. Since the meter knows the value of the precision feedback resistor, the current can be calculated with very good accuracy.

When interfacing with a thermopile or pyroelectric detector, voltage is the quantity that the optical meter must measure. There is, however, a considerable difference in how the measurement must be made between the two types of detectors. The optical meter's circuitry must be designed and configured to accommodate the two different types of voltage sources.

Thermopile detectors produce very slow bandwidth voltages (≈1 Hz) that can be measured in the sub-millivolt levels. One of the main concerns when trying to resolve such low voltages is to compensate for, or eliminate, thermoelectric voltages caused by dissimilar metals, which are generated in the connections and printed circuit board. It is somewhat ironic that the desirable physical effect that generates the voltage in a thermopile detector is similar to the undesirable effects that are present in the connections and printed circuit board. Precautions must be taken when choosing the electrical components to help minimize the unwanted thermoelectric voltages. Additionally, to resolve accurately small voltages, the optical meter must be able to zero any offset voltage due to temperature drift of the components and the thermopile.

Pyroelectric detectors, in contrast, produce relatively fast rise-time signals in the microsecond regime (see the figure in the Pyroelectric Basics section). The circuitry in the optical meter must sample-and-hold both the baseline voltage and the peak amplitude of the pulse. These two voltages are then put into a differential amplifier; and it is this voltage difference that determines the amount of energy in the optical pulse by way of the responsivity of the detector. Precautions must be taken to avoid accidental triggering of the sample-and-hold circuit since these circuits are sensitive to noise. Because the faster pyroelectric detectors have narrow upper peaks, it is crucial that the bandwidth of the circuit is fast enough to capture the level of the upper peak without degradation of amplitude accuracy.

Integrating Spheres

Newport's general-purpose integrating spheres can be used to make a variety of measurements. Optional sphere accessories are also available to enhance their utility.

Beam Power

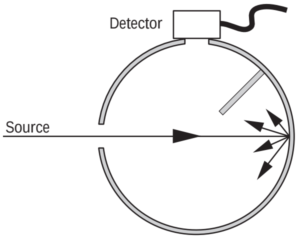

Figure 3—Beam Power

Measuring total collimated or uncollimated beam power (Figure 3), independent of polarization or beam alignment, is straightforward. The beam is admitted into the sphere and a detector, baffled from directly reflected radiation, measures the spatially integrated beam power. Integrating spheres are ideal for measuring the output power of divergent beams from laser diodes, lensed LEDs and lensed lamps.

Transmittance

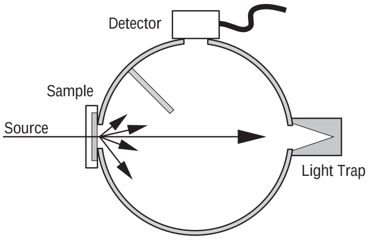

Figure 4—Diffuse Transmittance

Transmittance (Figure 4) can be measured by using the integrating sphere to collect transmitted radiation from a sample held in one of the ports. The sample is irradiated, then compared with a direct source measurement made outside the sphere. A baffle is used to shield the detector from non-integrated transmission, and a light trap can be used to remove the unscattered component. Measurements of total integrated scatter, fluorescence, bulk scatter and forward and back scatter can also be made.

Reflectance

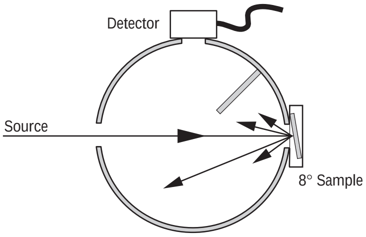

Figure 5—Specular + Diffuse Reflectance

To measure reflectance, a sample is held in one of the ports and irradiated by an incident beam. Total reflected radiation is spatially integrated by the sphere and measured by a baffled detector. The specular component of the reflective radiation can be eliminated by using the normal-incidence sample holder, which reflects the specular beam back out of the input port. An 8°-incidence sample holder allows measurement of the "specular plus diffuse" reflectance (Figure 5). The reflectance of a sample relative to a known standard can be calculated by measuring both and taking their ratio. The sample and standard should have a similar reflectance to avoid errors caused by sample reflectivity. A dual-beam system can be used to eliminate this potential source of measurement error.

Fiber Optic Power Output

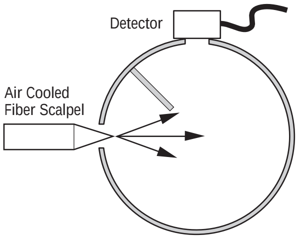

Figure 6—Fiber Scalpel Power

An integrating sphere also is ideal for measuring the output of optical fibers. In particular, this approach avoids the sensitivity of thermopiles to air currents and provides reliable NIST-traceable calibration of high-power, air-cooled fiber scalpels for surgical or ophthalmic applications (Figure 6).

Laser Diode Power

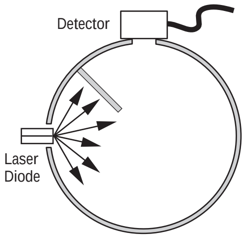

Figure 7—Laser Diode Power

An integrating sphere and calibrated detector setup is suitable for accurate, absolute value light power measurement of laser diodes. Your measurements will be insensitive to errors caused by detector positioning or problems associated with overfilling, or saturation, of the active area of the detector. A baffle, positioned between the input port and the detector port prevents the detector from directly viewing the emitting aperture of the laser or the direct area of illumination. In an integrating sphere the detected flux is always a small fraction of the incident flux. This attenuation, caused by light reflecting many times before reaching the detector, makes the integrating sphere an ideal tool for measurement of output light power of high-power lasers (Figure 7).

High-Speed Detector Terminology

Several terms are used to describe the performance of high-speed detectors, and are defined as follows:

Conversion Gain, CG: The sensitivity of a detector or amplified detector (usually into 50 ohms) converted to Volts/Watt via Ohm's Law. CG = Responsivity x 50 ohms.

Dark Current: The DC current that flows through a detector when there is no light present. Usually measured in the nanoamp range.

dB: Logarithmic unit of relative measure [e.g. 3 dB = ratio of 2:1].

dBm: Logarithmic unit of absolute measure for power [0 dBm = 1 mW].

NEP: The amount of optical input power that produces the same output level as the inherent noise level of the detector/receiver, i.e. a signal-to-noise ratio of one. Usually given in picowatts per root bandwidth. Total noise level is calculated by multiplying the NEP by the square root of the full bandwidth.

Optical Return Loss, ORL: The amount of light reflected (lost) back out of the detector towards the light source. Measured in dB relative to the input power level. For commercial single-mode systems, typical ORL values for a detector must be less than -27 dB. For multimode systems, -14 dB is usually the maximum tolerable value.

Power Bandwidth, -3 dB: The frequency at which the electrical output power of the detector falls to 50% of its value at DC. Same as "electrical" bandwidth. Typically used for specifying analog microwave detector bandwidths.

Pulse Width: The full duration at half the maximum value (FDHM) of the output current pulse when the detector is illuminated by a negligibly short optical pulse.

Responsivity, R: The sensitivity of a detector element to light given in amps/watt, independent of load resistance.

Rise Time: The 10–90% rise time of the output voltage step when the detector is illuminated by a negligibly short optical step function. This is difficult to do in practice, so the measurement is simulated mathematically by integrating the pulse width (see above).

Sensitivity: The optical input power (in dBm) required to achieve a particular Bit Error Rate, BER (or signal to noise ratio) at the output of the detector/receiver. Usually specified for BERs of 10^{-9} (or a S/N of 6). BERs of 10^{-12} require a S/N=7.

Voltage Bandwidth, -3 dB: The frequency at which the output current or voltage of the detector falls to 50% of its value at DC. Same as optical bandwidth. Same value as the -6dB power bandwidth.

Impulse Response or Rise Time?

Detectors temporal performance is often specified by either impulse response or rise time. Which one of these parameters is appropriate for your application?

Impulse response is best used when you are actually measuring pulses, i.e. signals that turn on and then return to zero. The impulse response of a detector tells you the shortest pulse you could ever expect to see output from the detector. For good resolution, you need to select a detector whose FDHM is at least three times shorter than the pulse you expect to measure.

Rise time is the parameter of choice when you are measuring either rising or falling edges. This type of measurement is especially common in digital communications systems where bit streams are comprised of an endless series of rising and falling edges. The rise time of a detector should be at least three times shorter than the risetime you expect to measure.

Clearly, impulse response and rise time are related quantities. Mathematically, the rise time of a detector can be obtained by integrating its pulse response. Clean pulses without tails or ringing approximate a Gaussian shape. Such pulses have rise times (10–90%) that are only 10% longer than the FDHM. In this case, the difference between the two values is negligible.

However, when pulse shapes deviate from the ideal, the difference between impulse response and rise time can indeed become significant. Pulses with positive tails produce longer rise times (and have less bandwidth), while pulses with negative ringing produce shorter rise times (and have enhanced bandwidth). See application note for more discussion.

Time-Domain or Frequency-Domain?

There are many common parameters one considers when selecting a detector for a particular application. These include pulse width, bandwidth, responsivity, spectral sensitivity, noise level, linearity, power handling, bias voltage, power consumption, etc., to name several. However, with the explosion in optical communications, detector applications have evolved into two major groups that have significantly different requirements for the *shape* of either the temporal or frequency response. The particular response shape requirement is usually determined by whether the user has a time-domain, or a frequency-domain, application. Knowing this, what does one actually look for when making a selection?

Figure 8a shows the 15 ps pulse response of a detector designed for time-domain applications. Figure 8b shows its corresponding frequency response curve. Note that in the time-domain the pulse is very "clean" showing very little tail or ringing. This type of Gaussian pulse response is ideal for applications where the temporal behavior of a waveform is under study, or where the temporal behavior of an optical signal must be converted to an electrical replica as accurately as possible. The most common applications are in signal diagnostics and receivers for digital communications, where temporal distortion can create bit-errors. Note that for this type of detector, the frequency response smoothly rolls off to a 3 dB point near 21 GHz that yields a Gaussian time-bandwidth product of 310 GHz-ps (power bandwidth).

Figure 9b shows the frequency response of another detector designed for frequency-domain applications. Figure 9a shows its corresponding temporal waveform. Note that in this case, the frequency response has been designed to be flat within 1 dB from DC to 20GHz. Beyond this, the response decays rapidly. This type of detector is ideal for many analog, microwave applications where a narrowband signal can be detected anywhere within the operating bandwidth with essentially the same sensitivity as at DC. Common applications include microwave communication links and radar arrays.

Note that the squared-off shape of the frequency response in Figure 9b results in significant ringing in its corresponding time-domain response. As a result, this type of detector would be a bad choice for time-domain applications. Similarly, frequency-domain users would be disappointed with a time-domain detector whose responsivity naturally drops by 3 dB at high frequencies, as shown in Figure 8b.

The nominal pulse width for these detectors might be specified as 15 ps for both products. However, it is the *shape* of either the temporal response or the frequency response that determines its usefulness for a particular application. Note that the pulse width is really only an accurate measure of comparison for time-domain detectors when there are no aritifacts on the waveform and the pulse shapes are the same.

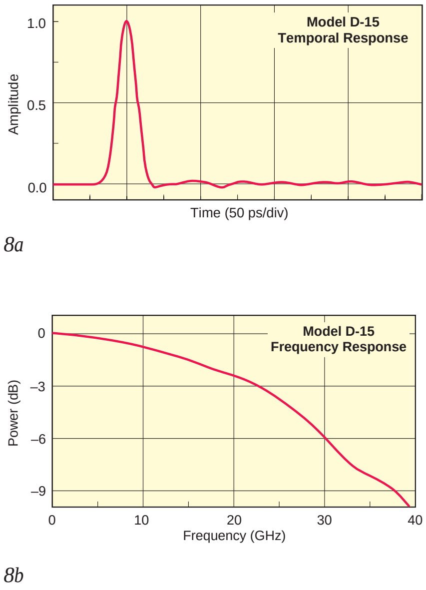

8a

8b

Figure 8—Temporal response (8a) and frequency response (8b) of a time-domain photodetector (Newport D-15) having a nominal pulse-width of 15 ps (full duration at half maximum).

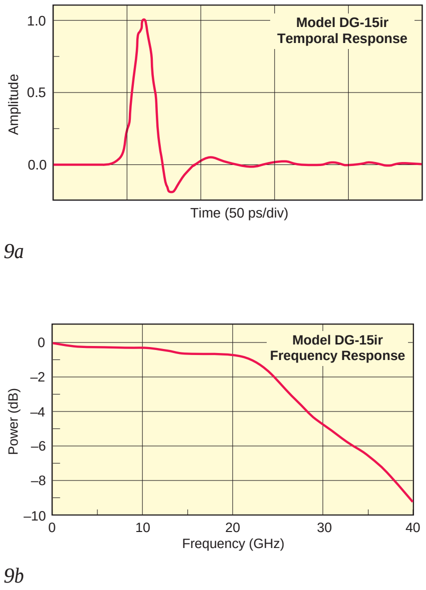

9a

9b

Figure 9—Temporal response (9a) and frequency response (9b) of a frequency-domain photodetector (Newport DG-15ir) having a nominal pulse width of 16.5 ps (full duration at half maximum).

Bandwith Terminology

There are many terms used to describe the bandwidth of a photodetector, but the two most common, "optical bandwidth" and "electrical bandwidth" also tend to be misleading, and often lead to some confusion when making detector comparisons. Let's attempt to clarify the nomenclature by describing a technique for measuring bandwidth.

A photodetector is a converter of optical power (mW) to electrical current (mA). This is why responsivity is specified in amps/Watt. High-speed detectors are simply designed to perform the optical to electrical conversion extremely quickly so when a short pulse of light arrives, the detector produces an exact replica of the input as a current pulse at the output. The shortest current pulse that can be produced at the output determines the speed of the detector.

The speed of a short-pulse detector can be determined by applying an extremely short optical pulse to the input, and then measuring the duration of the current pulse produced at the output. The output pulse is directed through a load resistor (usually 50 Ω) in order to generate a *voltage pulse* that can be displayed and measured on an oscilloscope. It is here that the pulse duration is determined.

The frequency response can be determined from the voltage pulse by mathematically transforming it to yield a voltage spectrum that shows how the response rolls off at higher frequencies (see Figure 10). The *bandwidth* of the detector is then defined as the frequency at which the response drops to 50% of its value at DC. On a log scale, this is the -3 dB point of the voltage spectrum, and it is referred to as the *voltage bandwidth*. It is this same measure of bandwidth that is referred to as *optical bandwidth* by other manufacturers.

"Voltage" Bandwidth = "Optical" Bandwidth

In analog, microwave applications the frequency response of a photodetector is often measured by using a microwave *power* meter, which gives a reading proportional to the *square* of the output voltage and therefore results in a *power* spectrum (see Figure 10). In this case, the bandwidth is defined as the point where the output *power* drops by 50% relative to its DC value. Once again, on a log scale, this is the -3 dB point of the *power* spectrum, and it is referred to as the *power bandwidth.* Historically, this has also been called the *electrical bandwidth*, in spite of the confusing fact that both voltage and power are electrical terms.

"Power" Bandwidth = "Electrical" Bandwidth

Newport's high-speed detectors are specified by both voltage and power bandwidth to avoid any confusion. The relationship between voltage and power spectra can be seen in Figure 10. The power spectrum simply goes as the square of the voltage spectrum, because power is proportional to the square of the voltage. On the log scale, this squared relationship appears as a factor of two difference in decibels (dB). Therefore you see that when the voltage spectrum has dropped to its -3 dB point, the corresponding power spectrum has dropped to its -6 dB point at exactly the same frequency.

Mathematically:

-3 dB voltage bandwidth = -6 dB power bandwidth

As a result, when comparing detectors, be sure you are comparing apples with apples, or in the case of detectors, the same measures of bandwidth. For most detectors, the voltage bandwidth is always greater than the power bandwidth, although the exact relationship is highly dependent on the shape of the individual detector's response curves.

In general:

-3 dB voltage bandwidth > -3 dB power bandwidth

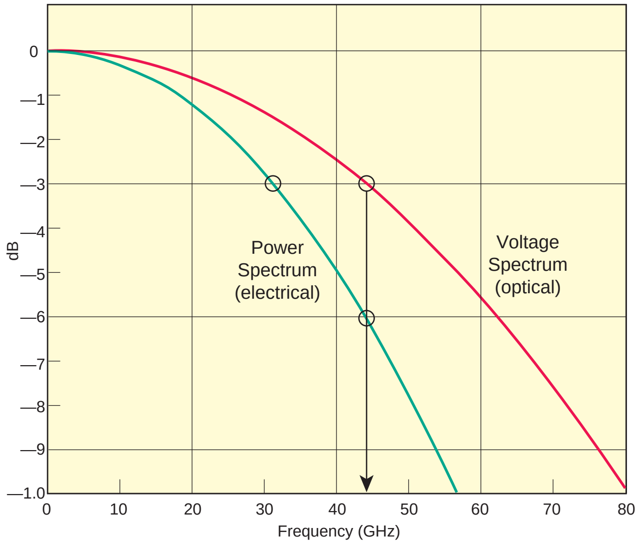

Figure 10—Frequency response of an ideal 10 ps detector