Photomultiplier Tubes

Introduction

Photomultiplier tubes (PMT's) are fairly simple devices, although sometimes a little tricky to use since the output electrical signal is dependent on both the light level but also (critically) on the voltage used to power the PMT. Their operation depends on two physics topics

- The photoelectric effect, in which incoming photons strike a metal surface, and eject an electron from the surface.
- Emission of electrons from a metal following the metal being struck by a previous electron. Providing the incoming electron has sufficient kinetic energy, more than one electron is emitted, the actual number being a strong function of the incident electron's kinetic energy.
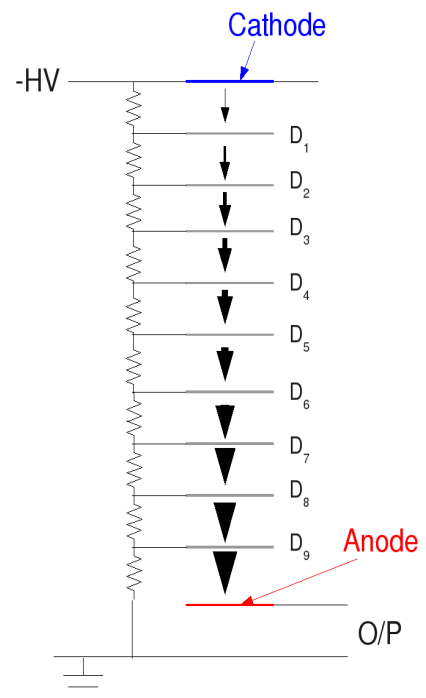
Photomultiplier operation

See the Hamamatsu Photomultiplier Handbook for a more detailed explanation of how a photomultiplier works.

A photomultiplier consists of five main parts
- a cathode which is coated with a material of low work function (see photoelectric effect)
- An anode to collect electrons
- A series of dynodes between the cathode and the anode
- An external power supply which produces an electric field between the cathode an first dynode, between the dynodes, and between the last dynode and the anode.
- An external current meter to measure the number of electrons collected at the anode, and a recorder to collect the information.

When photons of sufficient light falls on the cathode an electron is released (photoelectric effect). Each of these electrons is then accelerated towards the first dynode by the applied electric field. The energy that it receives is sufficient to liberate more than one electron when it reaches the dynode. Each of these secondary electrons is

Schematic of a Photomultiplier

accelerated towards the second dynode where they release more electrons. This is repeated as the swarm of electrons which is produced makes its way from dynode to dynode until the all arrive at the anode where they are collected and fed to the external current meter. Note: the detected signal is a stream of electrons, and hence a current, not a voltage.

The detected current depends on two factors; the number of electrons ejected from the cathode (which in turn depends on the number of incoming photons and on their energy), and the gain of the photomultiplier. The gain is defined as the number of electrons collected at the anode relative to the number of electrons ejected from the cathode. Let

- g be the number of electrons liberated at any one dynode when a high energy electron strikes it (the dynode gain.)
- N be the number of dynodes of the photomultiplier

Then the overall gain of the PMT is equal to g^N. For example if g=3 and N=9 (a nine stage tube) then the gain is 3^9 = about 20,000. However this gain is strongly dependent on the applied voltage. If the voltage is raised then the accelerating fields are stronger. That means the the energy of an electron as it strikes a dynode is also increased, and more electrons are released from the dynode. Suppose that the applied voltage is raised enough to make g=5. Then the PMT gain becomes 5^9 = 2 million. Adjusting the gain of the photomultiplier tube by changing the voltage is commonly used to adjust the strength of the signal to accommodate both strong and weak light sources.

To calculate the current we also need to know the number of electrons ejected from the cathode, which depends on two factors

- the number of incoming photons per second
- the quantum efficiency, defined as the average number of electrons released from the cathode by the photoelectric effect per photon.

Let us suppose that 1 pW[1] of light of wavelength 500 nm falls on a cathode whose quantum efficiency is 15%[2]. Then
- energy of one photon = $h\nu = hc/\lambda = 6.626 \times 10^{-34} * 3 \times 10^8 / 5 \times 10^{-7} = 3.98 \times 10^{-19}$ J
- number of photons ejected from the cathode = 10^{-12} W / 3.98×10^{-19} J = 2.5×10^6 per second
- if the gain of the photomultiplier tube is 2×10^6, then number of electrons arriving at the anode is $2.5 \times 10^6 * 2 \times 10^6 = 5 \times 10^{12}$ per second
- This is a current of 5×10^{12} electrons per second * 1.6×10^{-19} C/electron = 8×10^{-7} A

This is a typical figure. Generally speaking is is advisable not to let the current from a photomultiplier rise above 1 µA to avoid potential problems with non linearity and possible damage to the photomultiplier. On the other hand with a stable light source it is not difficult to measure currents down to 10 nA, perhaps even a little less.

Operating the monochromator with the PMT

There are three adjustments you need to make before taking a spectrum

1. Using the computer program for Monochromator Control, choose 'Diverter', and then select 'Side'.
2. Set the entrance slit width. (The exit slit is set to 20 µ, and should not be altered.)
3. Set the PMT voltage.

The last two of these are used in conjunction to get an appropriate maximum current from the PMT, say between 100 nA and 1 µA. There is no correct setting, you can use wide slits and a low voltage, or narrow slits with a higher voltage. In practice the slit widths are made a small as possible to get the best resolution, and the voltage adjusted to obtain the correct current, providing

---

1 That is not much light!
2 the stated quantum efficiency of an RCA R928 photomultiplier tube

- The PMT voltage does not drop too low (say below 50 V per dynode). Otherwise the PMT may not be turned on correctly.
- The PMT voltage does not exceed the safe operating voltage for the tube (see manufacturers specifications for the voltage ratings.)

A Practical Example – the 6383 and 6402 Å lines of neon

The first step is to get a signal. Adjustments can be made to then make the signal current lie in the appropriate range.

1. Using the computer's monochromator control, set the wavelength to be 640.32 nm (the approximate center of one of the neon lines. (The monochromator scale is slightly off from the actual wavelength scale.)
2. Set the entrance slit to 50 μ.
3. Set the PM voltage to 600 V.
4. Make adjustments to the PMT voltage to get a current of approximately 500 μA.

You are now ready to take the spectrum. There are however some additional parameters to consider.

1. Start wavelength. This should be just a little less than the lower wavelength. 638 nm is a reasonable choice.
2. Stop wavelength. This should be just a little larger than the higher wavelength. 640.5 nm is a reasonable choice.
3. Scan rate. A slow rate is preferable to get the best spectra. (Going too fast can make you miss lines.) For this measurement a scan rate of 0.1 nm/min gives decent results. However, note that if the scan rate is slow then the time to take the spectrum is high. For this example the total scan is 2.5 nm, so at 0.1 nm/min the full spectrum will takes 25 minutes. However if you wanted to record a spectrum over the full visible spectrum (400 nm to 700 nm) then at 0.1 nm/min that would take 3000 minutes (=50 hours!) A faster scan rate might be more appropriate.
4. Number of scans. Set equal to 1.
5. Scan delay. Set equal to 0.

The data is collected by a separate program, interfaced to the computer using a LabJack connecting board. The program has its own settings

1. Channel number. Set to 0
2. Data sample frequency. Sets the rate at which the computer samples the output of the electrometer. Five samples per second is a good initial choice, although if you are expecting the spectrum to be long a slower sample rate is better. (Otherwise you exceed the amount of allocated memory.)
3. Horizontal width. The display scrolls the spectrum across the screen, right to left. This setting controls how many data points are on the screen at one time. A low number shows only a small portion of the spectrum, but at high resolution. A high number can show the whole spectrum. A setting of 1000 is a reasonable starting point. (Note: This setting only affects the display, it does not affect your data. You can change it at any time, as often as you like.)
4. Vertical scale. Allows you to control the vertical range of the display, including the option for dual polarity. The signal from the electrometer is unipolar, and a range of 0 to 5 V is reasonable. (Note: This setting only affects the display, it does not affect your

data. You can change it at any time, as often as you like.)

To take data, first sett all the parameters above, then press start scan on the monochromator control, and the 'Acquire' on the data taking program. (Note: the monochromator control will automatically stop when you reach the end wavelength, the data taking program will not. You will need to stop it yourself.)

Spectral Resolution

The two figures below show the same spectrum, that of the two Sodium D lines (589.0 nm and 589.6 nm). These two lines are 0.5974 nm apart, far enough for the equipment to easily resolve them. However note the difference in the two spectra.
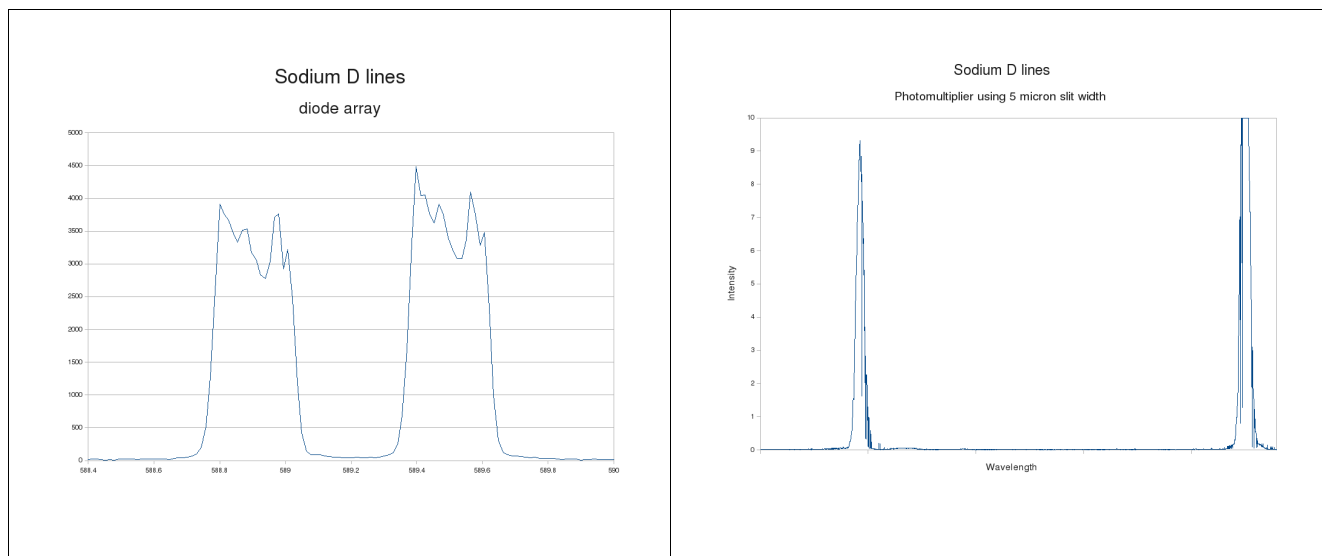
- The spectrum taken with the diode array shows two separate lines, but each one is very broad. Knowing the separation of the two peaks, their positions in the spectrum, and their widths in the spectrum, we can estimate the spectral width of each one at about 0.24 nm (2.4 Å)[3]. We would then expect that the monochromator using the diode array would be incapable of resolving two lines which are closer together than this figure.

  However, note that this spectrum is already calibrated by the equipment.

- The spectrum taken with the photomultiplier also shows two separate lines, but each one is now quite narrow. An estimate of the spectral width of each one is about 0.0086 nm (0.086 Å). We would then expect that the monochromator using the photomultiplier would be capable of resolving two lines providing that they are separated by at least this figure.

  Two points to note.

  o The width of these lies is dependent on the width of the monochromator slits. The spectrum above was taken with a slit width of about 5 μ. Making the output width larger increases the width of the peaks roughly in proportion.
  o This spectrum is not calibrated by the equipment. If we need to measure the actual

---

3 The spectrum taken with the diode array shows evidence of saturation of the detector, in which case the width of the line is overestimated. However, note the width of the region between the two lines in the two spectra. It is much wider in the spectrum taken with the PMT, showing that spectrum does have the better resolution.

wavelengths then we would need a separate calibration spectrum,

Spectral Response

The spectral response of a PMT depends on two factors

- The material from which the window is made. Commonly it is just the glass envelope of the PMT, in which case uv light below about 300 nm (3000 Å) will not pass through, and hence not be detected.
- The material from which the cathode is made. Commonly cathodes are made for two alkali metals (bi-alkali) because these materials have a low work function, and even photons of relatively low energy (long wavelength) can still eject electrons.

In order to find the spectral response for a given PMT you need to consult the manufacturer's specifications. The diagrams below show a 1P28 PMT (a very common 9 stage tube) and it spectral response.
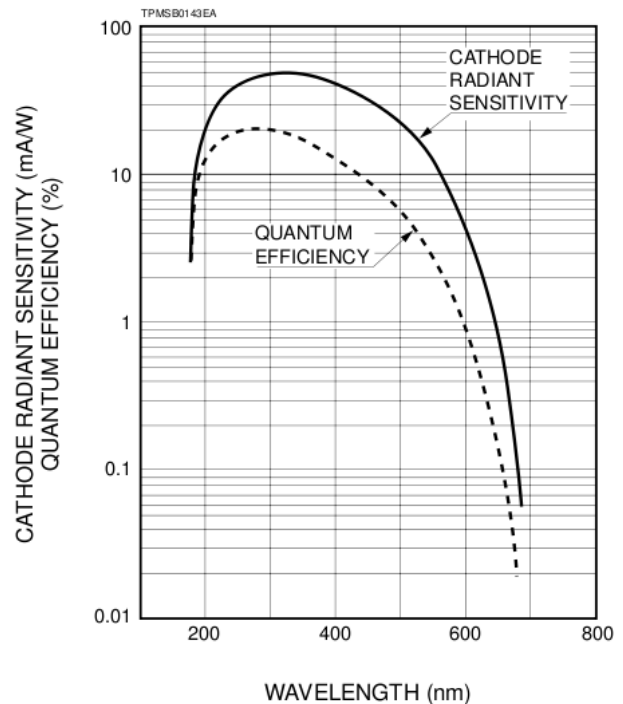
Note that the spectral response has too curves

- the sensitivity of the PMT to the incoming light measured as a power
- the quantum efficiency of the PMT measured relative to the number of incoming photons.

The difference is the energy carried by the photon, hc/λ.