How long is a photon?

A 'HowDoesItDoThat?' puzzle for physicists...

"How long is a photon?" The initial reaction to this question is probably that it is rather a strange one to ask. However I've come to think that it is useful because it can (pun alert!) shed some light on more than one of the curious discoveries in physics. It first came to my mind as a result of my work in the past using two-beam interferometers, but it may perhaps now also help us to understand the results of experimental tests of the Bell Inequality! The path to that can start from a 'classical' approach to interferometry.

The Michelson Interferometer is a standard instrument described in many textbooks. It employs a beamsplitter to divide an incoming beam of light into a two output beams – ideally with each carrying half of the input power. These are then sent along two distinct paths before being recombined. By varying the difference between the lengths of the two paths we can alter the relative phases of the beams being recombined. Most physicists will have heard about this type of instrument and how it can be used. But behind the standard textbook descriptions lurks a surprise or two for many people who think they understand how it works!

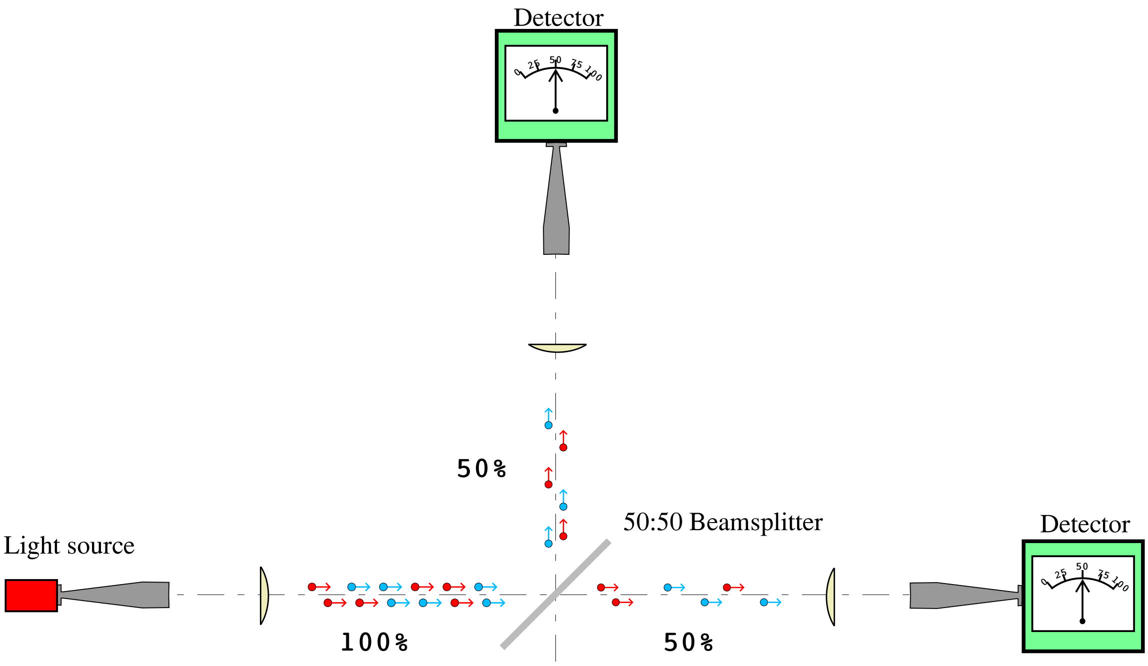

Figure 1 – The Beamsplitter

To explain this we can start with the arrangement shown above in Figure 1. The key item in a Michelson Interferometer is a 'Beamsplitter'. This is usually a flat sheet which acts so as to reflect half any light shone onto it, and to transmit the other half. Ideally it does this for all frequencies or colours of light that we are interested in, and doesn't absorb any of the light at all. In effect it acts as a mirror for some of the incident light whilst allowing the rest to pass though.

Figure 1 represents the light from a chosen source as a series of photons. For the sake of illustration these are of two colours (i.e. frequencies or energies), red and blue. The reflected and transmitted 50% of each colour are directed to a pair of light detectors which each absorb the photons that reach them and read out the intensity of the light they get – i.e. count the number of photons per second in this case because we are considering the behaviour in terms of photons. As we'd expect, each detector gets half the input light.

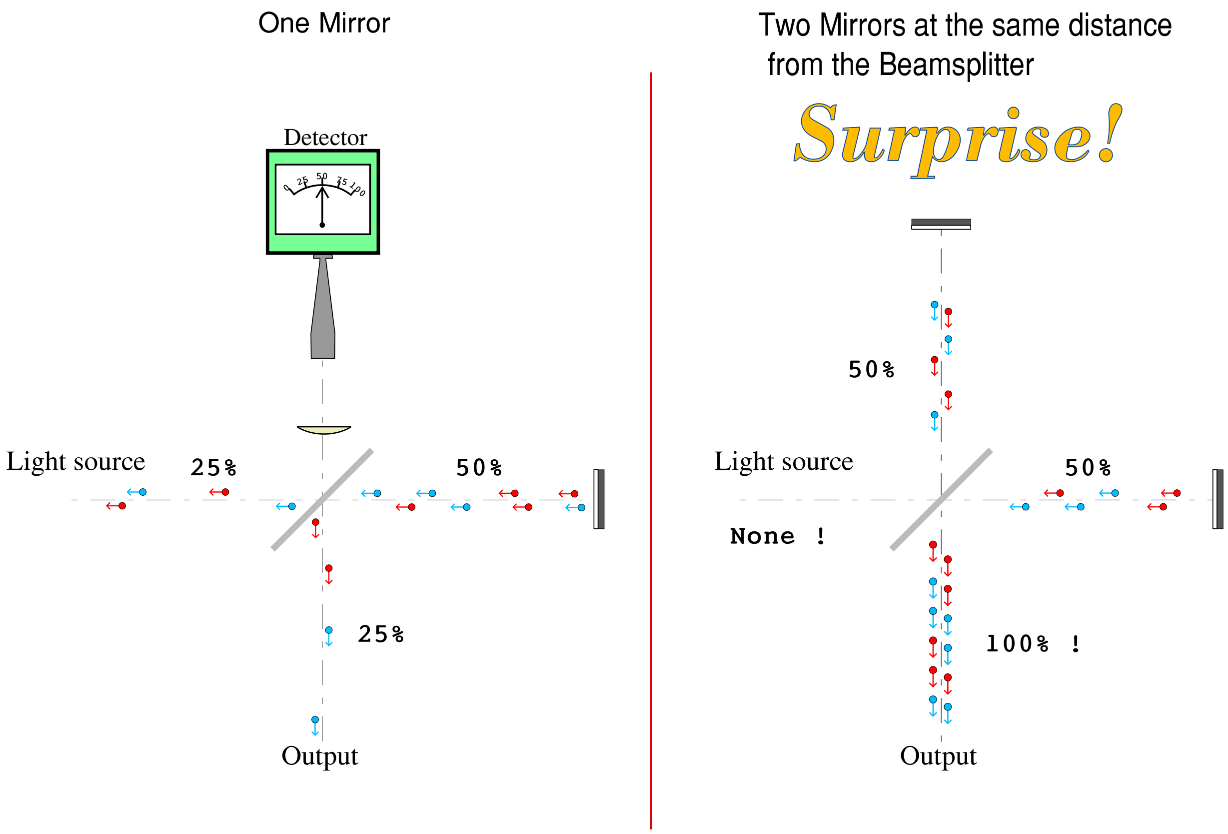

On reflection we get... Mirror Magic!

Figure 2 – Mirror, mirror...

Figure 2 shows what happens when we replace first one, then both of the detectors with plain mirrors. One mirror – as shown on the left – simply reflects back the 50% of the input photons which managed to reach it though the beamsplitter. These return to the beamsplitter and, again, get split equally. The result being that a quarter of the photons that were sent into the system are now directed back towards the light source, and a quarter are reflected in the direction of the instrument's output. However if we replace both of the original detectors with mirrors and carefully place them so that each mirror is at exactly at the same distance from the beamsplitter as the other, we get the result shown on the right of Figure 2. We find that *all* the input photons now get directed to the same place! In effect, all the photons we shone into the arrangement find their way to this output and *none* of them get sent back to the signal source!

This result *isn't* what we might expect from what we see when using only one mirror because that implies using two would allow the beamsplitter to send 50% of the light back towards the light source, and 50% to the output. So may be quite a surprise!

The behaviour arises because we carefully arranged for the two mirrors to be equally distant from the beamsplitter. However if we now move one of the mirrors we can discover how the output light levels vary with the resulting difference in lengths of the paths light must take to travel to each mirror and back.

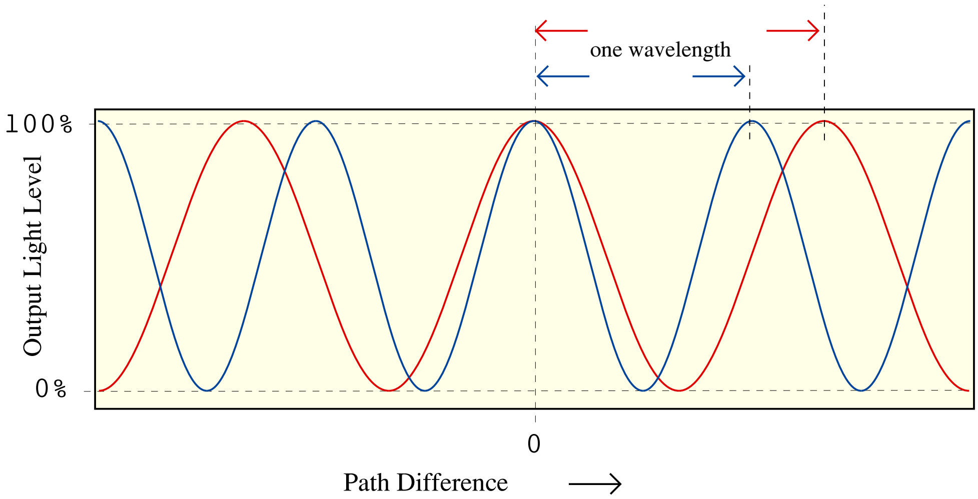

Figure 3 – How the output varies as we move a mirror.

The above shows the patterns we'd get for both the blue and the red light in our example. In an ideally constructed instrument we find that the power levels with a cosine wave shape. All choices of colour have a maximum output at zero path difference. But as we move the mirror we discover that each colour of light shows repeated peaks each time we have altered the round-trip distance by one wavelength for that particular colour. Blue light has a shorter wavelength than red light, so adjacent peaks in its interference pattern are closer together. In practice, this behaviour is what allows us to use the system as an instrument to measure the spectrum of any input light. The 'missing' light when the output is below 100% is, in the basic Michelson Interferometer, sent back to the light source.

Although I won't go into detail here it may be worth adding that the Michelson arrangement is just one example of a general class of instruments which are all forms of "two-beam interferometer".

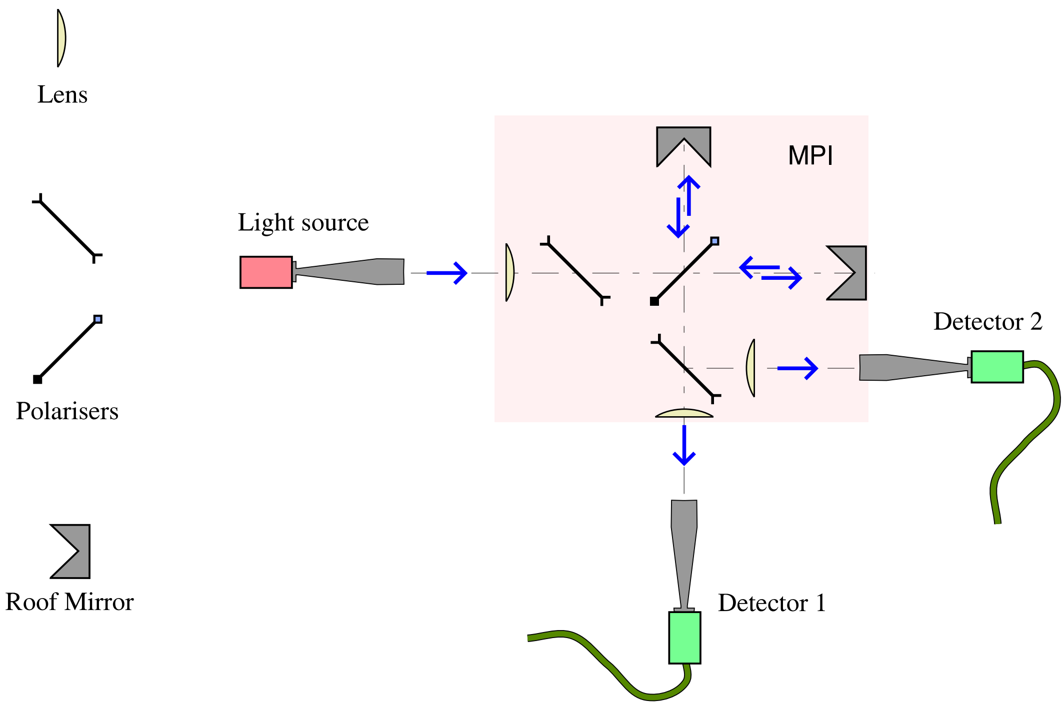

Figure 4 – The Martin-Puplett Interferometer.

The above diagram illustrates one alternative arrangement which used polarisers as the beamsplitter, etc. The details are more complicated so I won't explain how it works, but one advantage is that it avoids shining any input light back onto the source. Instead it provides *two* outputs and alters the relative levels which emerge from them as we alter the path difference. This is useful because some sources of light can be affected by having some output shone back onto them, and that can alter what we are trying to measure. It also improves the accuracy of measurements because all the light power can contribute to the results.

Useful as this is. if we think of the beamsplitter action as simply being due to each individual photon hitting it having a random 50:50 chance of being reflected or transmitted the overall behaviour makes no sense at all. Clearly something is wrong here with regarding photons as individual particles like energy carrying tiny billiard (or, if you're American, pool) balls...

At this point the standard response of a card-carrying physicist would be to step forwards, wave a magic wand, and pronounce the mystic incantation: *"Waves!"* The point being that although the above behaviour can look really weird when treating the light as a series of Quantum-Mechanical particles it all makes sense if we, instead, think of light in terms of the more traditional Electromagnetic Wave description. By treating the light as series of sinewave components we can understand the effect of using two paths of different lengths as the consequence of adding two contributions that have different phases and give an interference effect. However neat as this explanation may be, it still leaves some potential puzzles which can be exposed by considering two questions:

- What happens if the light level is so low that only one photon's worth of light power will be in the instrument at a time?
- What happens if we also make the difference in path lengths extraordinarily long?

For the sake of simplicity and brevity I'll ignore issues like beam diffraction and just consider the main implications, starting with the question regarding ultra-long path difference.

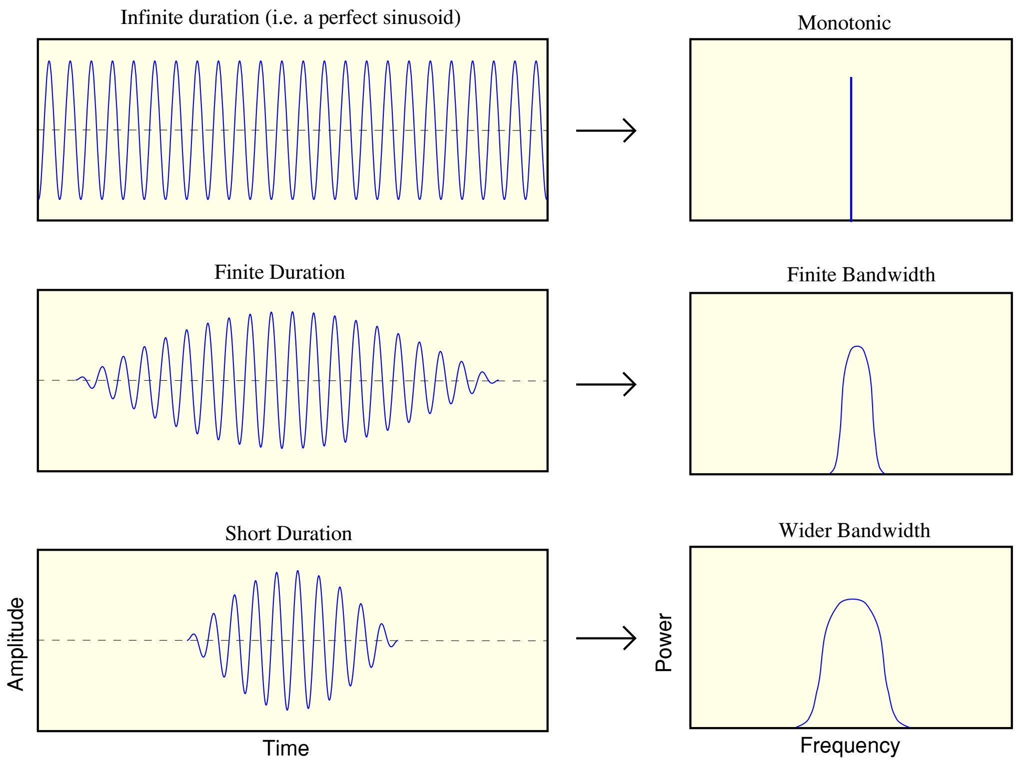

Figure 5 – Longer is better

We have already seen (Figure 3) that the output pattern we observe when we vary the path difference of an interferometer depends on what wavelength(s) – and hence colour(s) or frequency(ies) – are in the light being shone though it. This turns out to have some implications for the choice of our light source when considering the above questions. So a diversion to, erm, illuminate this makes sense...

The way we define a sinusoid that has a precise frequency can be taken to imply a wave with an infinite duration. The formal mathematical proof of this point tends to hinge on using the Fourier Transform. Alas 'hard sums' like those tend to be as welcome as Vogon poetry. Fortunately, Figure 5 can summarise the main result which is relevant here. A pure and perfect sinusoid has an infinite extent or duration, by definition. This goes with being perfectly monochromatic. However real-world signals, waves, and perhaps (?) photons may not behave like this. They may have some form of beginning and end. And even at the speed of light, that also implies a finite length. The overall result is that the shorter the duration or length of an oscillation, the wider or more blurred its frequency will appear to become.

In practice the above agrees with what we find when we use an interferometer to measure the spectrum coming from a light source. As we vary the path length we see that the pattern fades away as we move away from having a zero path difference. The wider the range of frequencies in the detected light, the narrower the range of mirror movement within which we can observe a clear pattern. In terms of 'waves' therefore, if we want to get an interesting result when the path difference is very long this tells us that we probably will need a to choose a highly monochromatic source. i.e. one that essentially only produces as close as necessary to a pure sinewave.

As an aside, it is possible to 'cheat' here: An alternative is to use light detectors that only respond to a suitably ultra-narrow range of frequencies. Any other light then 'doesn't exist' so far as the detectors are concerned. So we may then get the same results using a wideband source as if we'd used a narrowband source! Weird, eh?

Either way, if we want to find out what happens when we have an extraordinary large path difference we can assume we are using/observing a highly monochromatic light source. Which in turn implies that any photons that contribute to our results must all have the same specific colour / frequency / wavelength.

The obvious change we can expect when we use a large path difference is that the beam sent via the extended route should take considerably more time to return to the beamsplitter than the beam sent via the short route. Given a classical ideal monochromatic source, and perfect optics, this won't matter and we should find that when we vary the path difference we will still trace out the familiar perfectly sinusoidal variations. Now consider carefully adjusting the path difference so that with our ideal setup *all* of the beam power emerges from one output, and *nothing* emerges from the other or or goes back to the source, despite the very large path difference.

Returning to the first of our two questions: What happens now when we reduce the input power to such a low level that only one photon at a time is allowed to be passing though the system?

Laboratory experiments carried out in the past on more conveniently sized systems tend to confirm that in this kind of situation every individual photon will emerge only from the one output which we get light from at higher input levels, and no photons ever emerge from the other. But now consider another question: *when* do these photons emerge? Is the transit time from source to output detection set by the duration required to traverse the shorter path, or the longer one? Which transit time should we expect, and why?

We can imagine building and using an almost arbitrarily well tuned laser beam that is extremely near to being monochromatic. Then passing this though an attenuator to reduce the beam power to a level such that only one photon at a time is likely to be passing though our interferometer at any time. When we have arranged this, do we expect, in principle at least, to be able to adjust the interferometer so that the photons only ever emerge from one output – even if one path reaches to, say, the Moon and back? (I'm behaving like a theoretician at this point and handwaving away or ignoring 'trivial' details like atmospheric scatterings, losses, etc.) If we can do this we should be able to get an experimental result that answers the question about transit time. But I don't know of any experiment which has been done to tell us this result, presumably because it isn't an easy experiment to do in practice, despite handwaving away the practical requirements as being trivial!

I suspect the general view, though, is that for the interferometer to still be giving the expected output each photon has to have simultaneously "travelled down *both* mirror paths" and this then ensures that it will emerge from the expected output. Yet one path might, say, take a second longer to traverse than the other. Which seems to imply at least one of the following:

• That the photon must be at least one light-second long to be sufficiently monochromatic.

• That *either* the 'early path' part must get back to the recombination point and then wait for its twin to arrive, *or*:

• The 'longer path' part must somehow indicate what it should do as they are split.

How do we tell via an experiment? Indeed, does the answer vary from one photon to another? i.e. do some photons emerge after a 'short path' delay whilst others emerge after a 'long path' delay?

At this point I should confess that I started thinking about all of the above after seeing a report on a recent 'Bell Inequality' experiment that used two different astronomical sources to choose the measurement settings for the paired (entangled) photons they used. For some reason, the thought that popped into my mind at the time was, *"Did the paired photons arrive at the two detectors at exactly the same time?"* Which in turn also dragged Relativity into the argument because this question then hinges on what do we mean by the *"same time"*! So the conundrum now involves classical physics 'waves', Quantum Mechanics, *and* Relativity!

One feature of photons is, of course, that by their very nature we assume they must travel at the speed of light. But Relativity actually then tells us that this means any journey they make from a source to an end-point has – from the *photon's* point of view– *no* duration or interval at all! Its point of origin and point of arrival are co-incident! So far as a photon knows, it takes no time to travel no distance.

The Bell experiment employed a technique where one initial photon was 'split' into two 'entangled' photons that are essentially Tweedledee and Tweedledum twins. These then went - from our point of view – in two opposing directions to reach two measurement points. Given that they were twins, they presumably had identical frequencies and energies, and had a coherent relationship. They then duly gave the results predicted to validate Bell's expectations. When this sort of behaviour was first suggested it worried Einstein and others who called such behaviour *"Spooky action at a distance"*. Yet from the above it seems like *neither* of the photons, whatever their 'length' in other contexts, actually travel any real 'interval' at all. Hence there was no 'distance' between the two detection points – so no 'Spooky' aspect at all, according to Einstein himself.

The result is that we can perhaps now regard the two points of measurement, and the point of their joint origin, as all actually being only *one* point so far as the pair of entangled photons were concerned. Which for me raises the questions of how *different* in lengths – as perceived by us –

can the two paths from the generation point and the two detection points be whilst still obtaining this result? Is it infinite? And in our own terms, at what times do we find the photons arrive? One potential advantage of the methods used to test Bell's predictions is that it shows we can generate and use simultaneously created pairs of photons. So we could perhaps then use one of them to indicate the 'starting gun' and the other to indicate the 'finish of the race' when we send a single photon though an interferometer with a very large path difference. And thus make the measurements required to resolve the above. In the process perhaps again showing that so – far as the photon is concerned – it treats both mirror paths as being the 'same' even though they look very different to us!

If we can answer these questions by experiment, it may also lead us to finding the answer to an even deeper one: How long is a piece of string? ...Or have the string theorists already answered that one?

2700 Words

Jim Lesurf
http://jcgl.orpheusweb.co.uk/history/ups_and_downs.html
http://www.audiomisc.co.uk/index.html
https://www.st-andrews.ac.uk/~www_pa/Scots_Guide/intro/electron.htm
(N.B. The email address on the last of those sites no longer works since I retired!)

19th Nov 2020