Introducing Photonics

Brian Culshaw

Cambridge University Press 2020

Further thoughts and comments on the problems

The aim of this discussion is to shed a little more, hopefully helpful, light on the problems which appear at the end of each chapter and to expand a little on the initial hints presented in Appendix 6. And, of course, feel free to make contact through this site or via e mail…

Chapter 1

1.2

There are a number of other aspects to this which may be interesting for readers to ponder. The topic of microwave photonics is most frequently used in the 'microwave frequencies modulating light' context. However, as mentioned, individual microwave photons have been detected. So perhaps another aspect of this is – if an optical signal is modulated at microwave frequencies – what might be the implications of the modulation process on the photon energy distribution. How high does the modulation frequency need to go for the resulting photon stream to be usefully thought of as a collection of photons of slightly different (by the microwave photon energy) energy photons? The point here is that the normal way of viewing modulated carrier is as the carrier plus sidebands in frequency space. Could there be occasions where it is more appropriate to consider the 'carrier plus sidebands' as a collection of slightly different energy photons? How might this become relevant in final demodulated signal to noise ratio? (maybe have a look at shot noise implications)? This is but one example of where we use the approach (here photons or waveforms?) which is appropriate to a particular situation.

1.4

Is an X ray best viewed as proton, ray or wave? For medical imaging, we've always been very comfortable with X rays as waves travelling along straight lines and producing, in effect, shadows. And then there's X ray diffraction and examining the structure of, for example crystalline solids with Angstrom or better resolution. However, to make sense of these we need to know the X ray wavelength, and this is conveniently related to the X Ray photon energy (assumed equal to the voltage applied between the anode and the cathode of the X ray generation tube). So yet again, we have a situation where the radiation model to be considered varies depending on the particular issue we are addressing

Chapter 2

2.1

Perceived colour is an amazing intriguing arena. The colour we see depends on basically the colour distribution in the illuminating light and which parts of this illuminating light are scattered back to be received by your eyes and which are absorbed in the object you're viewing. This can also have a viewing angle component as well – some structures might, for example, reflect blue in one direction and yellow in another. And sometimes objects look different when views in reflected or transmitted light

We then have to ask – what determines which wavelengths are absorbed and which reflected, and of course, which wavelengths go straight through? Much is determined by atomic / molecular absorption properties which can be thought of through either energy levels or the Clausius Mossotti equation, as discussed in this chapter. However, this is by no means the whole story. What is returned to your eye depends on surface reflectivity combined with the shape of the illuminating beam (the sun can be thought of as pretty much parallel light. So reflection only works on the Snell's Law, Fresnel reflection etc and similar considerations.) but we do see colour over a wide range of angles – so light scattering is important and in general this is not a reflection process; it is related to the properties of the surface of the material and how these vary randomly over sub wavelength dimensions so that the is no 'reflection interference' type of impact.

And then there's regular – at sub wavelength level - structure within the material that you're viewing – the most common examples being butterfly wings. The colour here is not determined by the material itself, but by its regular structure over sub wavelength tolerance levels.

So, yet again, the approach to understanding and designing around a particular situation depends on the situation itself.

2.2
This needs some thoughtful exploration around how the various light sources function, how they can be applied and what our reactions are to these variations. There's much to observe here and also to think about in terms of the multiple uses to which we put light sources in the present time (and there will be more as time goes on…) and perhaps in terms of reflecting on the none too distant past a century or so ago at which time the electric light bulb was well invented but was far from being an everyday item!

2.3
Well, this is just another viewpoint on problem 2 – all this concerns thinking about the world around us, how it works and all those things to do with light that we take for granted – but how do these work in practice, and what lessons can we take?

2.4
This problem is covered in Appendix 6.

2.5
The gold nanospheres query here have been mentioned in the text. Yes, the colours are the same, and the reasons lie in the fact that the nanosphere modifies the actual energy levels in the material through its being small enough to modify the solutions to the Schrödinger wave equation for gold particles. Again – we've shifted into viewing the phenomenon to the perspective of physical phenomena which happen to best fit the situation!

2.6
This is simply a matter of digging for the relevant free electron concentrations and substituting in the equation. Transparency above the plasma resonance is then to be expected – or is it? In the ionosphere – well we have satellite communications (and radio telescopes) as manifestations of this, but short wave radio bounces off! As for gold – well similar but at a much higher plasma frequency. In silicon – well the neatly organised crystal structure comes into play above the n type plasma resonance and the standard 'energy levels' (or Clausius Mossotti) approach is relevant. However, silicon transistors do operate over a very wide range of frequencies – many above, many below this

plasma resonance – so maybe here's an occasion when the concept is simply a distraction? Something to think about here and discuss among yourselves and with your teachers. **2.7**
The phase and group velocity discussion applies anywhere we look at wave motion. The group velocity is the speed at which energy is transferred along the propagation direction. The refractive index is however the ratio of the phase velocity to the velocity of light in vacuum for a material of infinite extent (though what may be the criteria for 'infinite'? The phase velocity is the speed at which a wavefront happens to travel along the direction in which you're observing this. You can visualise this by pebbles into ponds or numerous other simple experiments – and a phase velocity exceeding group velocity is exemplified in the picture below:

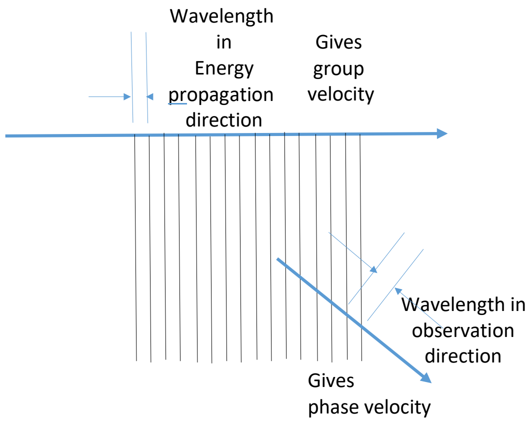

Figure A4.2 also indicates this principle.

Chapter 3

3.1

The key to what radiation comes to the earth can be found from an internet search – which may well yield something like this:

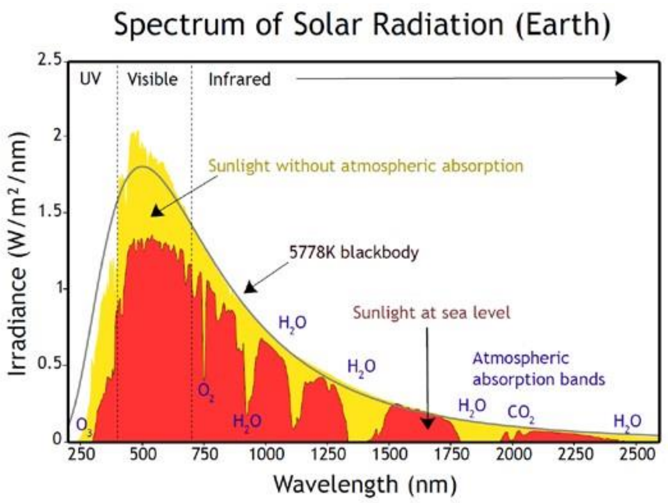

…and the absorption bands in the visible are clear on the above. As for the kW per square meter – well that is a noontime figure on a sunny summer day but now think about the size of the magnifying glass – hence its collection area and its focal length. Then assume a reasonable focal spot size (what's the criteria for 'reasonable'?). Then consider the implications of the figure you obtain.

3.2

Really little useful to add over Appendix 6 at this present time – this is one for exploration and discussion among yourselves – and there's much more to learn from that process than from reading a solution on line…

There may be a few things worth considering – you'll need to assume a diameter for the fibre core – 8 microns is probably a good starting point, but there are many variations in practice. Then you'll need to consider transverse power distributions (go for uniform for a start – but consider among yourselves the implications for the actual situation). There's also the observation that the Kerr effect depends on optical power density.

And on the impact of the actual situation – well the power levels (and hence the Kerr contribution) will vary logarithmically along the fibre a function of attenuation. So the Kerr induced phase difference follows exactly the same dependence as the power in the fibre. So if the average power is 1mW and you've worked that one out already, will that change in the presence of the attenuation?

3.3

Again, this is an explorative problem for group discourse. There's a lot of basic material implicit here: first of all, can we trigger electron flow using the optical electromagnetic field? The answer is 'yes' but the issue, as we've seen, is that the piece of wire is unbelievably lossy thanks to skin effect and the like, so it really is impractical. And the optical frequency electronic amplifier has yet to materialise (though – watch this space - may well apply)

Moving on to currently used detectors, the basic detection limits stem from the temperature coefficients of resistivity in the bolometer case; and the shot noise vs thermal noise in the band gap detector case. There's also considerations such as thermal time constants (which will determine bandwidth) feasible for the bolometer and the trade-off between the speed at which the temperature changes are achievable and the bolometer sensitivity. Also, there's the inevitable best wavelength effect in a 'band gap' detector – this occurs at a wavelength slightly shorter than the wavelength corresponding a photon energy equal to the band gap.

As for section d – how to optimise a broad optical bandwidth optical detector… Well, the implication in the text is that this is not designed for broad operational signal bandwidths modulated onto the source – so the electronic post detection bandwidth can be reasonably assumed to be not an issue. In which case you could envisage a stack of band gap detectors placed on top of each other with the band gaps increasing with depth (or maybe decreasing with depth from the input surface – now which one would it be?) and in principle one could design these into the wave guiding structure with the detectors along the guide and gradually going deeper with a means of adding the electron currents from all the detectors built into it. Again much to be clarified among yourselves here!

3.4

For part (a) - well it is safe to assume that the interface between the air and the LED is flat and has a

refractive index of? Well – go and check. It is also safe to assume that without any further design on the simple junction, the light is trying to get out uniformly over a 4π steradian solid angle. For all this light, only the light incident at the planar interface below the critical angle will escape. This give an absolute maximum on the escaping fraction – remember too that it is a good approximation to consider each atom as a source of light propagating at some random direction in the 4π solid angle. Even from this there will be a reflected and transmitted component, which can be derived from Fresnel reflection.

The lens issue can also be a quite complex one. Yes, it is feasible to minimise the reflections from one surface to the next based on Fresnel reflections – but remember you'll need to also minimise the reflections from the two interfaces – LED to lens and lens to air. So, in the end it is the minimising of this product which is needed, and several surface lensed versions of the LED have been explored over the years using exactly these guidelines.

There's another approach – deposit a layer of a carefully chosen combination of refractive index and thickness and use this to couple from the LED surface to the world outside. If you were to assume normal incidence – then how thick would this layer be – and what about its refractive index? Interestingly similar layers find their way into camera lenses and a host of other applications. The wiki article gives food for thought!

As for the lasing case – well the basic process through which this can be explored is described in the question and it is certainly well worth thinking this one through among yourselves. The basic need is – like any oscillator – that the gain available in the amplifying medium exceeds the losses from this amplifying medium. In the optical context – the implication here is carefully designed high reflectivity (but often limited bandwidth) mirrors and – as in all oscillators – the output wavelength is determined by a combination of the gain bandwidth of the medium and the resonant frequency (or often frequencies) of the resonant feedback circuit.

3.5

There's little to add here over the hints in Appendix 6. The story is essentially about changes in birefringence induced by compressing or expanding the sample in one direction and relating these changes into the characteristics of a filter comprising the two crossed polarisers and the birefringence and orientation thereof of the stresses in the plastic. Recall that the birefringence is effectively a difference in optical thickness between the two principal axes (stress direction and orthogonal thereto) and is, to a good approximation, pretty much constant in time over the spectral range we're looking at (that is twice as far - roughly – in phase in the blue as in the red.)

3.6

- On the first part – well this problem gets to be really interesting if you forget the 'zero field slope' approximation – but first it's useful to set up criteria for determining whether the zero field slope approximation is acceptable. These criteria you can assess among yourselves – always much more interesting and much more instructive that way! This also relates to the answer for the bias voltage – how far out may that be?

- On the transit time issue – well the necessary insight here is that more (typically many more) carriers are generated at the light input end than at the distant end of the depletion region. However, whilst the carriers are traversing the depletion region, they are inducing current in the photodetector load circuits. In other words, each absorbed photon creates charges which contribute to the external current for the total transit time, here for the 10 microns – 100psec. Think this through – this is equivalent to saying that a very, very short impulse (say

1 psec) will produce a detected current which last for 100psec. Now this corresponds to a bandwidth of???? But when you ponder the implications of the very high bandwidth optical fibre links which facilitate the internet, in the 100'sGB/sec per channel region, then there's an appreciation of the elegance in designing a suitable photodetector! Not to mention all the other contributory components!

- On the final section – well the absorption coefficients appear in figure 3.12. These are in slightly alien units, and assume an already 'optimised design' – they represent the best you'll get! And the absorption spectra of silicon over the near UV to Near IR range vary hugely (over around 10 orders of magnitude – look them up!) Silicon is transparent in the near IR (hence silicon photonics!). So it is interesting to explore among yourselves the ideas of absorption coefficients, quantum efficiency and responsivity and how all these factors might relate into photodetector design!

Chapter 4

4.1

FIRST and FOREMOST – there's a misprint in question 1(a). The reference to figure 3.1 should be to figure 4.1. It's also a usable approximation to assume that glass has a refractive index of 1.5

(though you can look up values for a range of glasses if you so wish!)

4.2

Again, little to add! This is straightforward Snell's law on the input on the hypotenuse side of the prism coupled to total internal reflection needs on the other two sides. Again assuming a glass with an index of 1.5 will give the necessary insight here.

There are however some subtleties here. Totals internal reflection is the ideal, but there is a minimum angle of incidence for which this occurs, and then there are (polarisation dependent) reflections which will contribute partial reflection at lower angles of incidence on the reflecting air lass interfaces. Calculating the angles over which total internal reflection occurs on both the input and output reflecting faces will give a useful indication of the range of useful incident angles. There are also specular reflections at the input surfaces for all input conditions. And suppose that somehow or other we could make a prism from liquid water – index 1.33. Any comments on the potential of this as a retroreflector?

4.3

And – if you take careful notice – the red sky at night can be much more dramatic when there's a few clouds around, and relatively 'ordinary' without the clouds. There's clues in here too!

4.4

As mentioned in the hints appendix – the perfect parabolic mirror can be viewed in ray optics as the ideal means to take incoming rays along the axis and focus them into a single point.

The second part is more subtle. On the face of it here we have an aperture of 5 metres – 10 million wavelengths at 500nm. So the resolution should be around 10^{-7} radians. (figure 4.1) The next question though – over what range of input angles might this resolution be maintained. In other words, how far from the principal axis in angular terms can we go before this perfect focus starts to fade? This is one to discuss among yourselves – what would be good criteria? How would you estimate this? And then there's the resolution of the observer's eyes as the limit on the effective number of useful pixels that the resolution can usefully fill? Exploring this gives interesting insights – including on the ever-increasing number of pixels on digital cameras and why they might be there.

Chapter 5 has a few comments (figure 5.12) on the turbulence issues, but might there be other approaches too? Think through the implications of what you've explored thus far in this. Could there perhaps be ways of sectioning the image and applying some form of computer based image processing for example (a very fancy version of Photoshop). Maybe take small sections in the image and use adaptive interfaces? Maybe images averaged over time (with the tracking system operating too of course!)?

4.5

The principal point here is that there is much to be learnt from the diffraction pattern of an object. Of course, it only happens if the object is laser illuminated – as you demonstrated in part (a). The interesting aspect here is that specific features of the object can be highlighted by selecting features in the diffraction pattern – image processing! Take the patterns produced by the laser pointer – and put a lens in there too to get an image… How might this image look if for example you only let through the central vertical bit of the pattern to form the image by using a slit? There's much to be got from simply playing around in the spatial frequency space and relating this into image processing – much microscopy for example has relied on these ideas for decades! This was all done though in the analogue domain – what about images from a digital camera?

4.6

The sinusoidal phase object has a constant unity amplitude transmission, but does produce a diffraction pattern – the Fourier transform (diffraction pattern in effect) does have harmonic components. The key here lies in the relative phase of all these components, notably with respect to the DC component. But – you may say – the dc component doesn't have a phase – but in the optical imaging context it for sure does (the light carrying this component has a phase relationship with that carrying the other spatial harmonic components (the diffracted beams). And changing this relative phase can have profound effects – the phase contrast microscope being a well-established example. As indeed can removing this component entirely – dark ground imaging, also well established.

4.7

The angles as is evident in figure 4.12, do need to be equal – it is a reflection process. And all the reflections need to interfere (optically) constructively.

As for the Doppler shift – this, as we've said, needs careful thought, but here are some clues. Looking at the optical phase of the same point in the diffracted beam as the ultrasonic wave moves along is the starting point and considering how this changes as the ultrasonic beam passes through one cycle. Looking at figure 4.12, imagine that the horizontal lines represent the peaks in the ultrasonic wave pressure. Then let these move along half a cycle and, recalling that the deflected optical beam hasn't moved, look at the pressure wave induced phase changes in that overall beam. The parts of the optical beam which went through the peaks are now in the troughs, so the overall relative optical phase has shifted through half a cycle. This is the key – but it still needs to be carefully thought through!

4.8

There's little to add here over and above the previous hints. In essence the dielectric guide has less than unit reflection coefficient at the core to cladding interface – so the electric field is finite at the interface. However, there is some reflection and so the 'interfering beams' view of the field in the core of the guide is still useful – just that the interface reflects less than 100%. You may even be able to get some insights on this from Fresnel reflections? To get to the heart of this needs Maxwell's equations, but what's here is good to get the feel of how it works…

Chapter 5

5.1

There's much to think about around this problem - for example, suppose the photodetector area could somehow be made to match the core size of the fibre (about 8 microns diameter typically) – would this help in any way and how might this impact on the device design?

There's also a bandwidth issue in the 'transit time for carriers' effect discussed earlier – if you do have a need for a 1GHz bandwidth – how far can the carriers go before there's an issue here? This also in effect defines the maximum width of the stray capacitor in the photodiode, implying a minimum capacitance per unit area.

Yet another aspect is the necessary signal to noise ratio in the practical receiver. Bit error rates come into this – in other words what is the probability of a transmitted '1' being interpreted as a '0' or vice versa. This is typically set around 10^{-9}. How would you go about calculating this? In essence, if there are on average N photons per bit arriving at the receiver, what is the probability that there may be 0.5N or less actually arriving? Interesting to explore!

The details of this overall design are what in the end will make the receiver work reliably and predictably over a long period. Possibly the principal aim of this discourse is to highlight that whilst the basic principles are – usually - fairly straightforward, finding the actual devices, stitching them together and putting them into an appropriate enclose requires significant teamwork and lots of expertise.

5.2

The starting point is the spectral resolution, here set at 0.5nm over a range of 1micron to 200nm – that is extending over 800nm. So the minimum number of resolvable points here is 1600. Probably going for 2000 would give a little leeway – or – a question to ponder – are we better settling for the 1600, or alternatively going for the 3200? The 2000 points will have, albeit small, blank spots between the adjacent detectors. Inevitably some of these may well coincide with the 'centre 'of one of the 0.5nm resolution points. In other words, should the number of photodetectors in the array be equal to the number of resolution points, or multiples thereof?

The required per detector SNR is also up for debate. The sample (see figure 5.2) absorbs some of the light and the relative absorbance is the parameter of interest. This value is invaluable as a guide to sample composition. The required SNR here will be at least 100, and much more for some applications. Then the source power distribution can be arrived at – how may photons for SNR of 100 on each one? That does of course depend on bandwidth – so let's start with 1Hz.

And now you have the essential features of the design problem. As in problem 6.1, this gives some indication of the range of possible specifications for practical spectrometers (and by implication, the range of prices the suppliers may charge – look them up!)

5.3

Really no further hints spring to mind here over the comments in Appendix 6….

5.4

Just a couple of quick observations – in OCT, the resolution improves as the bandwidth of the source increases. In essence for a 10 micron resolution – you'll need to be sure that the source coherence length is less than twice the specified resolution (the light goes to and returns from the sample - think about it…). As for range restrictions in the sample – well that is all about attenuation between entering the sample and later leaving it.

As for the eye and silicon – the eye is – by definition transparent in the visible – the silicon isn't (though all is well in the silicon photonics band). But – is the eye transparent in the near IR – one to check out!

5.5

Somewhat along the lines of problem 3 – no further useful hints spring to mind at present – except maybe to mention that for measuring very low rotation rates – ($5x10^{-4}$ degrees per hour) – relatively long integration times are OK. How long, and why would this low bandwidth be useful here?

5.6

Just a hint – though you've probably found this already – in addition to the texts on tweezers, there's a very nice Wiki account which should, with a careful read and a chat among yourselves, give more than the necessary background!

Chapter 6

6.1

Yet again, no further hints spring to mind here over the comments in Appendix 6….

6.2

The *Nature Photonics* article referred to in Appendix 6 presents a very good and pertinent overview of what is now an increasingly important topic –thanks to the semiconductor devices industry.an immense processing portfolio has emerged for silicon – and it is transparent at communications wavelengths (around 1.5 microns). What we have in figure 6.6 is the P type silicon as the waveguide with a junction into the N region at the top and P^{++} and N^{++} regions as the contacts. This could work in several ways essentially either introducing carriers into the P region or depleting the P region of carriers? What might the mechanisms be in either case – and which would you go for as the preferred option?

This last comment is also the key to any comparisons in terms of phase or intensity modulation depths from the modulation zone here as compared to the responses to he previous question. For this discussion – far simpler to ignore the fact that niobate doesn't transmit too well at all in the near IR!

6.3

For the first part, no further hints spring to mind here over the comments in Appendix 6…. For the second part – well, we've looked at the ideas of skin depths in metals at optical frequencies – for the TE case the incoming wave will attempt to set up an electric field in the plane of the metal overlay, and, expressing it simply, this wave will experience severe losses in the overlay, given that the skin depth (figure 3.3) is dropping to the nm level. This can be (and has been - very successfully) used as the basis of an 'on waveguide' polariser. There's always a 'but' though and here it is – how would surface plasmon resonances impact on all this? Something more to ponder… **6.4**

There's a straightforward key here to appreciating the operation of the device – the electrons flow as surface plasmons. In other words, there are no (or, at worst, a safe assumption to work on this basis) electron to lattice collisions involved in the current flow. So really in principle all that's needed is to change the thermally present electron flow with zero net velocity into a net flow at the velocity needed for the 180^{o} phase change. A little algebra with the necessary electric fields in mind will lead into a suitable drive voltage.

As for the sheet resistance concept – the idea of 'Ohms per square' and how it all works out is well described in a Wiki entry (https://en.wikipedia.org/wiki/Sheet_resistance) and this will lead onto the necessary insights for graphene.

6.5

Just a few more comments on this one - if you look at the structure to the right of the diagram – there are 36 holes in each of the outer sections over a total length in the region of 6 to 7 microns – in other words – yes this could be an optical cavity. But it also needs to be thin (remember all the plasma resonance things and the like) and very accurately machined to ensure that all the sections have the same resonance characteristics. There also needs to be mechanical as well as optical coupling, so there's a lot to consider for 0.5dB of squeezing! The cryostat is needed to access the minima for both optical and mechanical losses. You can, as mentioned in Appendix 6, check all the details in the original paper. As for the future… well look into the plethora of squeezing approaches (including LIGO – where the squeezing idea actually worked and was useful, even essential) and maybe check into more recent activity from the authors of the reference here, and others. Then come to your own conclusions!

6.6 to 6.8…

…are, as mentioned in the appendix, really short projects for discussion among yourselves – but they do point to the immense e potential for photonics as “the electronics of the 21st century…” The younger, less experienced and therefore minimally prejudiced, voice has much to say on this for sure!