BY HOW MUCH DOES THE PERFORMANCE OF AN OPTICAL FILTER CHANGE WITH ANGLE OF INCIDENCE?

First of all, it's important to make clear that this analysis applies specifically to filters that operate by optical interference effects, which in practice most “precision” optical filters do. For filters that just work by incorporation of wavelength-dependent light-absorbing molecules, such as in Schott glass, the performance is going to show little if any significant angular dependence, with light going through at more extreme angles just seeing an apparently greater filter thickness and hence correspondingly more absorbance. But the perhaps unexpected finding is that the response of an interference filter shows, at least for small angles, relatively little angular dependence too, which is not what one might have expected. So what is actually going on here?

The reason why there is going to be some angular dependence for interference filters is because they consist of multiple layers of optical materials of differing refractive index and carefully designed thickness, and which have been deposited onto a glass substrate by a sputtering process Internal reflections between these refractive index boundaries are going to be either constructive or destructive according to the wavelength of the light with respect to the thicknesses of all these layers (the number of which may be well into double figures), and this is inevitably going to depend on the angle at which the light goes through them. The angular dependence issue is of potential importance in two basic types of situation. In the first, the filter is in a beam pathway that is “normal”, i.e. at 90 degrees, to it, and which we will define as an angle of incidence θ of zero degrees, but the beam encompasses a range of angles around this. In the second, the filter is at some other angle to the incoming beam, specifically in the case of a dichroic mirror, where the angular range is usually centred around 45 degrees. As far as the filter performance itself is concerned, it doesn't matter whether it is optically in a collimated (“infinity”) space or in one which will generate an image, although angling a filter in an “imaging” space will introduce astigmatism. However, that effect is due to the refractive index of the filter, hence is true for any other angled optical component, and is therefore a separate matter from the one being considered here. This particular note, although also relevant to the 45 degree case, is primarily intended to cover the case of a significant angular range around normal incidence. So the question is, in practice how big can θ be, before the performance of the filter is significantly affected?

To understand what happens when the light is at some angle θ with respect to this normal condition, we need to resolve it into two vector components, as shown in the left part of the Figure. One of these components, which we've labelled x, actually IS normal, so that it goes through at 90 degrees to the filter surface, and another, which we've labelled y that is parallel to the surface, so that it isn't actually "seen" by the filter at all. What may not be immediately apparent is that since the path length of the normal component is somewhat shorter than the actual "angled" path length, its wavelength as seen by the filter is also effectively shorter. If we normalise the angled path length to 1, then from Pythagoras we obtain the length of the "straight through" component x as sqrt($1-y^2$), where y is the length of the surface one. Or if we know the angle θ with respect to the straight through component, then the length of that component, and hence the effective wavelength of the light, just varies with cosθ. So to summarise, as the angle of incidence increases, the response of the filter shifts to shorter wavelengths, although relatively not by very much for small angles.

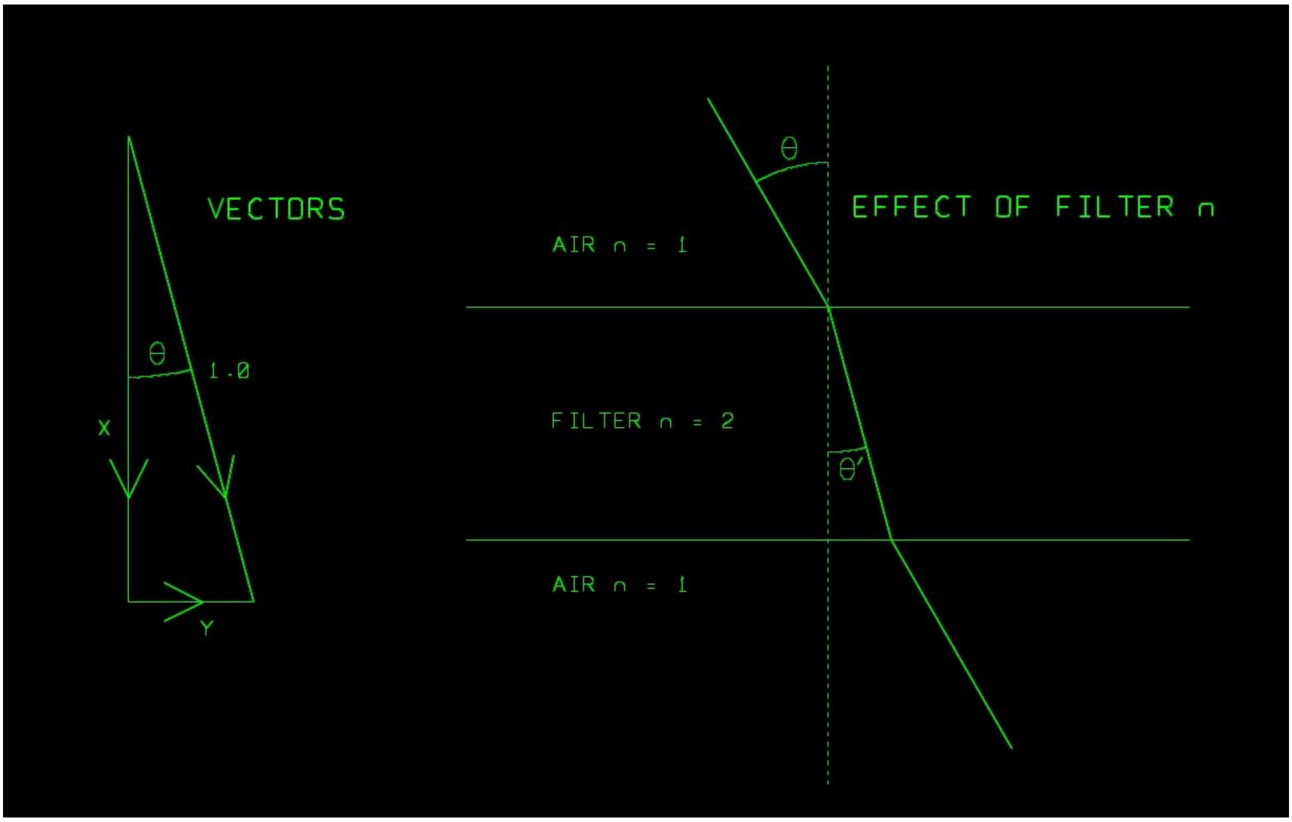

But what happens at larger angles? We've already mentioned the case of a dichroic mirror, designed such that some wavelengths (usually the shorter ones) are reflected, and others are transmitted. Depending on their application, they may be designed for use either at normal or 45 degrees incidence, and a typical rule of thumb, that we learned years ago from a Comar catalogue, is that if you use such a component at 45 degree incidence when it was designed for use at normal incidence, its spectral response for mid optical wavelengths (say 500nm) typically shifts towards shorter wavelengths by around 35-50nm, hence to 450-465nm. But wait a moment! We might think that since the cosine of 45 degrees is around 0.7, the shift should be by this same factor, which would be to 350nm, so clearly something else is going on here.

The answer is that we've used the wrong θ! There are in fact TWO angles that we need to consider, which we'll term θ and θ'. The reason is that we have to take into account the relative refractive index n of the filter material with respect to air, and as this link to the Semrock website shows, the effective value of n (bearing in mind that it is in practice going to be a mixture of different values) is relatively high, at typically around 2. So, as the right part of the Figure shows, the beam angle going through the filter, which we are calling θ', is significantly less than the incident angle θ at the filter surface. From Snell's law of refraction from an air interface, where the refractive index on the air side approximates very closely to 1, we obtain $\sin\theta' = \sin\theta/n$. Since for relatively small angles, $\sin\theta$ is close to linear with θ, then in such cases we can in practice say $\theta' = \theta/n$. Hence if we further approximate n to 2, then we can say that the angle going through the filter (with respect to normal incidence) is only about half that at the filter surface.

In practice this makes a BIG difference! Although using this approximation at 45 degrees means that we are pushing the linearity of the sine function, if we nevertheless approximate the angle through the filter (or in this case dichroic mirror) as just 22.5 degrees, then the cosine of this is about 0.92, from which we would calculate an effective wavelength shift from 500nm to 462nm in this case – exactly in line with Comar's rule of thumb. But perhaps more importantly, the improvement (in terms of the reduction in shift) from the effective n=1 assumption is significantly more than 2, and hence is more beneficial than we might have thought.

In fact, it's easy to show that for small angles, the improvement approaches a square-law one. This comes directly from the relation that for any angle θ, $\sin^2\theta + \cos^2\theta = 1$, which is just the re-expression of the Pythagoras relation in this particular form, as the Figure also clearly shows. Hence we can say that $\cos\theta' = \mathrm{sqrt}(1-\sin^2\theta')$, and since for small angles we can make the approximation that $\sin\theta = \theta$, we can see that for small angles, the departure of $\cos\theta$ from unity increases with the square of θ. Hence if we halve θ in the filter as a result of having $n=2$, then we are actually four times better off!

We can put all this into a single relation (while still estimating some equivalent overall value for n) by replacing $\sin\theta'$ in the above equation by $\sin\theta/n$, and then expressing $\cos\theta'$ in terms of the effective wavelength for a given angle of incidence, compared with that for normal incidence. This gives (as others have done before us, so we are merely confirming here!)

$$\lambda_\theta = \lambda_0 \, \mathrm{sqrt}[1 - (\sin\theta/n)^2]$$

The main potential pitfall here is that in practice the effective value of n tends to be somewhat polarisation-dependent. As soon as we go away from normal incidence, then polarisation issues are potentially important. Instead of rays, we now need to think of wavefronts arriving at an angle θ with respect to normal. The electric vectors of these wavefronts can be resolved into two components, namely the p component, which is parallel to the wavefront, and the s component, which is perpendicular to it. This can be a particular issue for dichroic mirrors, where θ is likely to be around 45 degrees, so any difference in n between the two polarisations is going to be correspondingly more important. Such components therefore need to (or should!) be designed to minimise this difference, whereas no such equivalent care is likely to have been taken for filters designed for use at normal incidence, so perhaps rather more polarisation dependence may be observed in this case. However, it is always likely to be a secondary effect here.

Note also that the primary topic here is the use of filters at relatively small angles to normal incidence. Once we get up to 45 degrees, for the dichroic mirror case, the approximations become much less accurate, and the angular effects become correspondingly more significant, to the extent that they can now be used to “tune” the filtering over a potentially useful range.

However, it remains the case that the range is still not as great as might have been expected. In this case, by using the n=2 approximation if a more precise figure isn't available, we can use the above equation to explore whether the effects of angular variation within a nominal 45 degree beam are in practice going to be a problem or not. But at the risk of making sweeping generalisations, one can nevertheless say that from how these components tend to be used in an optical system, it tends to be the case that the wavelength-splitting characteristics of a dichroic mirror aren't as critical as the bandpass characteristics of the subsequent filtering. Or perhaps we can say that the system can often be designed so as to make this the case, allowing us to live with the problem instead!

But for filters used at nominally normal incidence, the "take home" message is that their angular dependence is a relatively forgiving one. In fact this is often put to good use when a filter is used in an infinity space in an optical system. The potential problem here is that the back reflection from the filter can be refocussed by what was the preceding collimating optics, to give rise to a ghost image superimposed on the original object. This issue is routinely solved by angling the filter by a few degrees (we use 5 degrees in our own filter cubes and related products), so that the reflection is off-axis and will miss the preceding optics, and hence cannot be refocussed. In this case the effect really is negligible, as for n=2 we can calculate that the wavelength shift at 500nm is just less than 0.5nm. However, the point we are making here is that in practice the acceptable angular range is likely to be much greater. A rule of thumb that we use at Cairn is that if θ is less than 15 degrees, giving a θ' of say less than 7.5 degrees, then usually we don't need to worry too much. For example, for n = 2 the wavelength shift for a design wavelength of 500nm would be only to 496nm, and here θ is likely to be an extreme rather than an average angle in the incident beam, so the average effect will be correspondingly less. Our appreciation of all these matters is leading to our devising some potentially interesting optical arrangements, especially in respect of illumination, but that is going to be a story for another day!