4B11 Crib 2026

1 a) [30%] Let us assume we have an arbitrary aperture function, *A(x,y)*, with coordinates [*x,y*]. The light passing through this aperture will be diffracted at its edges and the exact form of this pattern can be calculated using Huygen's wavelets.

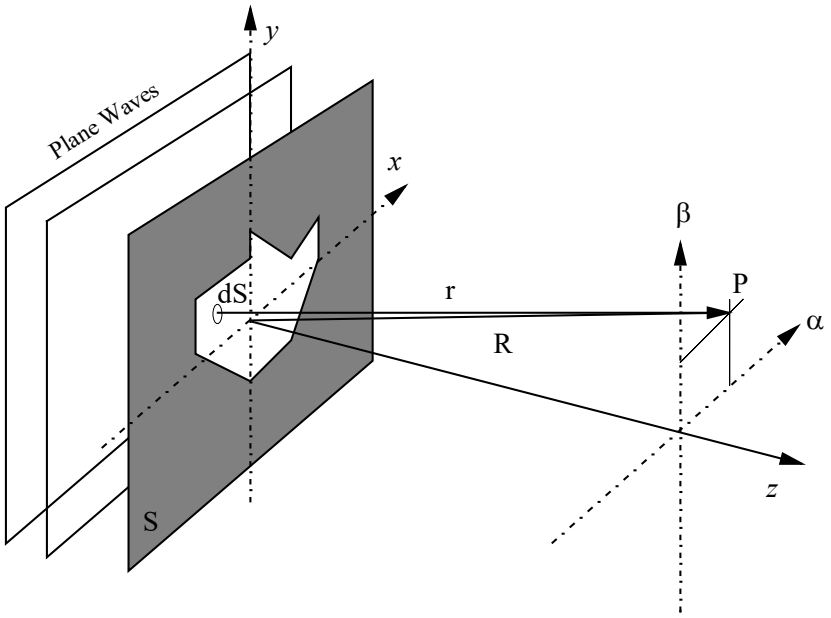

If we consider an infinitely small differential of the aperture, *dS*, we can model this as a point source of light emitting spherical 'Huygens' wavelets with an amplitude of *A(x,y)dS*. The wavelet acts as a radiating point source, so we can calculate its field at the point *P*, a distance *r* from *dS*. The source *dS* can be considered to radiate a spherical wave front frequency ω.

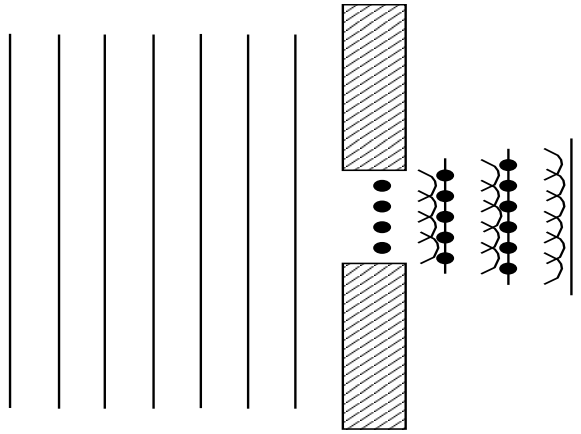

In order to understand and analyse the propagating wavelets, a series of approximations and assumptions must be made. If we consider only the part of the wavelets which are propagating in the forward (+*z*) direction and are contained in a cone of small angles away from the z axis, then we can evaluate the change in field *dE* at the point *P*, due to *dS*. As the wavelet *dS* acts as a point source, we can say that the power radiated is proportional to $1/r^2$ (spherical wavefront), hence the field *dE* will be proportional to *1/r*. We can see that for a real propagating wave of frequency ω and wave number *k*, ($k = 2\pi/\lambda$) we have the cosine component of a complex wave.

$$dE = \frac{A(x,y)dS}{r}\cos(\omega t - kr)$$

The full complex field radiating from the aperture can be written in terms of exponentials as the cosine is just the real part of the complex exponential.

$$dE = \frac{A(x,y)}{r} e^{jwt} e^{-jkr} dS$$

b) [25%] Now, we need to change coordinates to the plane containing the point *P*, which are defined as [α,β].

$$R^2 = \alpha^2 + \beta^2 + z^2$$
$$r^2 = (\alpha - x)^2 + (\beta - y)^2 + z^2$$

Which can be rearranged to give.

$$r = R\sqrt{1 - \frac{2\alpha x + 2\beta y}{R^2} + \frac{x^2 + y^2}{R^2}}$$

The final full expression in terms of *x* and *y* (*dS* = *dxdy*) for *dE* will now be.

$$dE = \frac{A(x,y) e^{j\omega t} e^{-jkR\sqrt{1 - \frac{2\alpha x + 2\beta y}{R^2} + \frac{x^2 + y^2}{R^2}}}}{R\sqrt{1 - \frac{2\alpha x + 2\beta y}{R^2} + \frac{x^2 + y^2}{R^2}}} dxdy$$

Such an expression can only be solved directly for a few specific aperture functions. To account for an arbitrary aperture, we must approximate, simplify and restrict the regions in which we evaluate the diffracted pattern.

If the point *P* is reasonably coaxial (close to the *z* axis, relative to the distance *R*) and the aperture *A(x,y)* is small compared to the distance *R*, then the lower section of the equation for *dE* can be assumed to be almost constant and that for all intents and purposes, *r* = *R*.

$$dE = \frac{A(x,y)}{R} e^{j\omega t} e^{-jkR\sqrt{1 - \frac{2\alpha x + 2\beta y}{R^2} + \frac{x^2 + y^2}{R^2}}} dxdy$$

The similar expression in the exponential term in the top line of the original equation is not so simple. It can not be considered constant as small variations are amplified through the exponential. The regions of the approximation are defined such that in the case where the approximation is bearably accurate, we are in the Fresnel region (often referred to as the angular spectrum). The exact boundary of the Fresnel region will depend on the acceptable accuracy and there are several methods that can be defined to create the Fresnel region of propagation..

c) [20%] To simplify the expression for the Fresnel region we must consider only the far field or Fraunhofer region such that.

$$R^2 >> x^2 + y^2$$

In this case, the final term in the exponential ($(x^2 + y^2)/R^2$) can be considered negligible. To further simplify, we use the binomial expansion,

$$\sqrt{(1-d)} = 1 \quad - \quad \frac{d}{2} \quad + \quad \frac{d^2}{8} \quad \ldots$$

and keep the first two terms only to further simplify the exponential expression.

$$e^{jkR\left(1-\frac{\alpha x+\beta y}{R^2}\right)}$$

Hence the simplified version of the field *dE*, can be expressed as:

$$dE = \frac{A(x,y)}{R} e^{j(\omega t-kr)} e^{jk\left(\frac{\alpha x+\beta y}{R}\right)} dxdy$$

The regions of the approximation are defined such that in the far field or Fraunhofer region, the approximations are accurate, hence the field distribution *E(x,y)* only changes in size with increasing *z*, rather than changes in structure. The total effect of the dS wavelets can be integrated across *dE* to get an expression for the far field or Fraunhofer diffraction pattern.

$$E(\alpha,\beta) = \frac{1}{R} e^{j(\omega t-kR)} \iint_{Aperture} A(x,y) e^{jk(\alpha x+\beta y)/R} dxdy$$

The initial exponential term $e^{j(\omega t - kR)}$ refers the wave to an origin at $t = 0$, but we are only interested in the scaling of relative points at *P* with respect to each other, so it is safe to normalise this term to 1.

$$\frac{1}{R} e^{j(\omega t-kR)} = 1$$

Thus, our final expression for the far field diffraction pattern becomes the Fourier transform:

$$E(\alpha,\beta) = \iint_{A} A(x,y) e^{jk(\alpha x+\beta y)/R} dxdy$$

d) [25%] If a positive focal length lens is included directly after the aperture, the far field pattern appears in the focal plane of the lens.

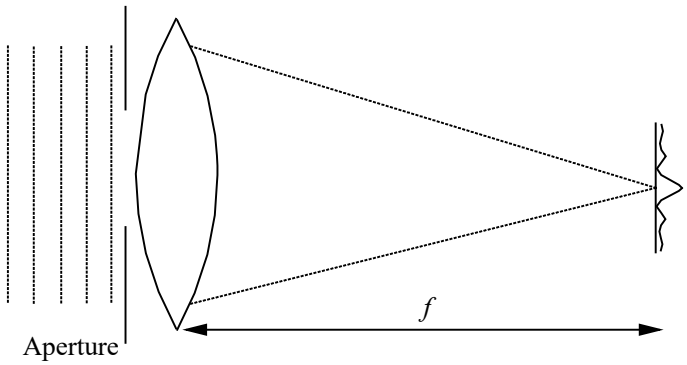

If we consider the aperture *A(x,y)* placed just before a positive lens of focal length *f*, then we can calculate the field just after the lens as described by Goodman, by the paraxial approximation.

$$A(x,y)' = A(x,y) e^{-j\frac{k}{2f}\left(x^2+y^2\right)}$$

The application of Snell's law at the spherical lens/air boundaries of the lens shows that the lens converts plane waves incident upon it into spherical waves convergent on the focal plane. For this reason, the diffraction to the far field pattern now occurs at in the focal plane of the lens, as the Fresnel approximation only holds true in this plane. The effect of this is to shift the far field or Fraunhofer region away from the distance *R* to the single position specified by the focal length *f*. The final result for the diffracted aperture *A(x,y)* through the lens is.

$$E(\alpha,\beta) = e^{\frac{jk}{2f}(\alpha^2+\beta^2)} \iint_A A(x,y) e^{\frac{jk}{f}(\alpha x+\beta y)} dxdy$$

The lens introduces the quadratic phase distortion term in front of the transform. This can lead to a smearing of the Fourier transform and must be corrected for in compressed optical systems..

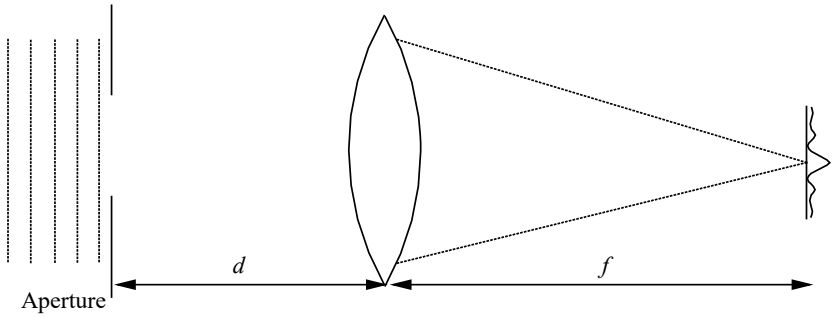

If the aperture is placed a distance d behind the lens, then there will be a corresponding change in the phase distortion term of the Fourier transform.

$$E(\alpha,\beta) = e^{\frac{jk}{2f}\left(1-\frac{d}{f}\right)(\alpha^2+\beta^2)} \iint_A A(x,y) e^{\frac{jk}{f}(\alpha x+\beta y)} dxdy$$

If the distance is set so that $d = f$, then the phase distortion is unity. This is a very important feature used in the design of optical systems and is the principle behind the *4f* system. This forms the basis of a low distortion optical system.

This question was not as well answered as I expected, especially as it was mostly book work and a fairly standard question. A lot of answers had the right derivation, but not in the right sections. Part a) should end with the element of energy dE not for Fourier transform. in Part b) a lot of answers confused the two approximations that define the Fresnel region and applied the binomial expansion without definition. The final part of c) should have need the Fourier transform but quite a few did not apply the approximations in the right way. d) was well answered with most knowing the 4f system and quite a few discussing phase distortion. Quite a few tried to use the shape of the aperture to define things which is not correct.

Q2 a) [20%] The assumption of matched sampling and square pixel size generates the replay field by adding a sinc envelope and outer orders. We can now look at the scaling (dimensions) of the replay field. The spatial frequency coordinates (u,v) are related to the original absolute coordinates in the far field plane (α,β) by the relation.

$$u = \frac{k\alpha}{2\pi f} \qquad v = \frac{k\beta}{2\pi f}$$

We can use the above relationships to directly calculate the positions of peaks in the hologram's replay field. A two dimensional hologram comprises of *NxN* square apertures (pixels) with a pixel pitch Δ. The two dimensional envelope due to the fundamental pixel which covers the far field diffraction pattern (or FT) of the hologram is a 2-D sinc function.

$$A\Delta^2 \mathrm{sinc}(\pi\Delta u)\mathrm{sinc}(\pi\Delta v)$$

Hence it forms the envelope function for the replay field of the hologram. The useful information of the replay field is contained in the central first lobe of the sinc function, so we can calculate the width of the replay field as where the first zero of the sinc function occurs ($\pi\Delta u = \pi, \pi\Delta v = \pi$). We want the coordinates in terms of [α,β], so we use the above transformation to get.

$$\alpha_M = \frac{f\lambda}{\Delta} \quad \beta_M = \frac{f\lambda}{\Delta}$$

As the sinc function is centred on the 0th order, $2\alpha_M$ and $2\beta_M$ will tell us the width of the central lobe of the sinc envelope, due to the pixel pitch Δ. At the simplest level, the replay field will have an estimated 'spatial frequency pixel' of pitch.

$$\alpha_0 = \frac{2f\lambda}{N\Delta} = \frac{2 \times 100 \times 10^{-3} \times 880 \times 10^{-9}}{10 \times 10^{-6} \times 256} = 6.88 \times 10^{-5} m$$

The spots in the replay field are separated by 4 zero spaces, hence the spacing will be

$$4 \times 6.88 \times 10^{-5} = 2.75 \times 10^{-4} m$$

b) [30%] The height of the spots in the ideal version of the replay field is given purely by the Sinc envelope. The inner most spot to the central optical axis will be at position $[u_1,v_1]$ = [2,2] in spatial frequency coordinates and the outer most spot at position $[u_{10},v_{10}]$ = [18,18]. The ratio can be calculated via the 2D sinc function above, which given that $u_1 - v_1$ and $u_{10} = v_{10}$, simplifies to

$$\frac{sin^2(\pi\Delta u_1)(u_{10})^2}{sin^2(\pi\Delta u_{10})(u_1)^2}$$

We need to convert the coordinates of the replay field into [u,v] spatial frequencies.

$$u_0 = \frac{\alpha_0}{\lambda f} = \frac{6.88 \times 10^{-5}}{880 \times 10^{-9} \times 100 \times 10^{-3}} = 782\ m^{-1}$$

Using these values, gives a ratio of 1.067 which means that the spot at position u_1 is 6.7% higher than the spot at u_{10}. This assumes there are no imperfections I the system and the far field is purely the Fourier transform of the pixel array.

The ratio can never be one for the reason seen here, which is the effect of the sinc envelope. It can also never be one as the direct binary search algorithm running with binary phase cannot guarantee all the spots to be the same height.

c) [30%] Direct binary search (DBS) is an algorithm where we take a hologram of random pixel values and then calculate its replay field by optimisation. We flip the binary value of a randomly positioned pixel and calculate the new replay field. We then subtract the two replay fields from the target replay field, sum up the differences to form a cost function for the hologram before and after the pixel change. If the cost function after the pixel has been flipped is less than the cost function before the pixel was flipped, then the pixel change is considered to be advantageous and is accepted. The new cost function is then used in comparison to another randomly chosen flipped pixel. The process repeats until no further pixels can be flipped to give an improvement in the cost function.

1) Define an ideal target replay field, T (desired pattern)
2) Start with a random array of binary phase pixels.
3) Calculate its replay field (FT), H0.
4) Take the difference between T and H0 and then sum up to make the first cost, C0.
4) Flip a pixel state in a random position.

6) Calculate the new replay field, H1.
7) Take the difference between T and H1 then sum up to make the second cost, C1.
8) If C0 < C1 then reject the pixel flip and flip it back.
9) If C0 > C1 then accept the pixel flip and update the cost C0 with the new cost C1.
10) Repeat steps 4 to 9 until |C0 - C1| reaches a minimum value.

The DBS algorithm does not require the full replay field to calculate the hologram. It uses a target function that represents all the spots in the hologram replay and often a few zero points as well to suppress the noise. It only calculates the replay values at the target positions, which greatly speeds up the algorithm. DBS works well with binary phase as it only has to calculate the change in phase in the replay due to a flip in each pixel value.

This is not a fully optimum means of generating a hologram, but it gives a very good approximation. One of the limitations is that the algorithm will always find a local minima rather than a global one. This means that the replay field will have excess noise and that uniformity is very hard to guarantee. One improved algorithm is simulated annealing, which uses DBS, but also includes a probabilistic evaluation of the cost function, which changes as the number of iterations increases. The idea is to allow the hologram to 'float' during the initial iteration, with good and bad pixel flips being accepted. This lets the optimisation float into a more global minima rather than getting stuck in a local minima as is the case with DBS.

d) [20%] One possible technique is sinc envelope compensation. When we calculate the target function, the uniformity of the spots in the array are inversely adjusted to compensate for the effect of the sinc envelope.

A second technique is to change the error function used in the algorithm. Rather than just taking the difference between the replay field before and after a pixel flip (which just makes the spots as bright as possible), a different function could be used. A function the emphasises the relative heights of the peaks, or their uniformity rather than just pure difference would, work well.

This question was more challenging with very few getting it right. The key to this question was to define your coordinate system clearly which only a few did. Part a) was mostly ok, but a lot were out by a factor of 2 without defining the sampling in the replay field. Part b) was not answered correctly by anybody which was surprising as it was a numerical problem. A few got that it was a ratio of sinc functions, but most did not convert the real coordinates into spatial frequencies properly. Part c} was pure book work and was answered well in general. Part d) was more speculative and there were some excellent ideas on how to correct for the sinc variations. The best one was to repeat the annealing until the local minima gave you a flat replay field.

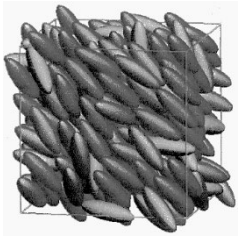

Q3 a) [30%] The nematic phase is the least ordered mesophase before the isotropic. The molecules have only long range order and no longitudinal order. This means that the molecules retain a low viscosity, like a liquid, and are prone to flow. The anisotropy in the shape of the calimatic molecules leads to a dielectric anisotropy which is represented by a pair of dielectric constants, ε_{perp} and ε_{para}. The total dielectric anisotropy $\Delta\varepsilon$ is defined as the difference between the two. It is important to note that can be both positive as well as negative (vertically aligned nematic (VAN) devices). Molecules with a negative $\Delta\varepsilon$ behave in the opposite way electrically to those with a positive $\Delta\varepsilon$. The existence of this dielectric anisotropy means that we can move the molecules around by applying an electric field across them, combined with the flow properties means that a nematic molecule can be oriented in any direction with the use of an electric field which leads to their ability to perform greyscale modulation of the light.

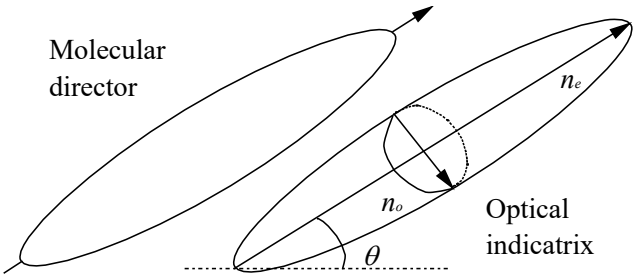

The calimatic molecular shape also leads to an optical anisotropy in nematic LCs, with the two axes of the molecule appearing as the refractive index. The refractive index along the long axis of the molecules is often referred to as the extraordinary n_e (or fast n_s) and the short axis the ordinary n_o (or slow n_s) axis. The difference between the two is the birefringence. $\Delta n = n_e - n_o$

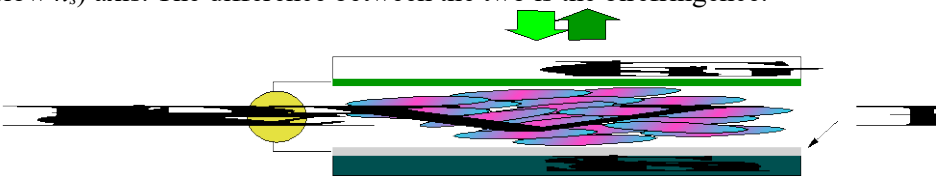

The optical anisotropy means that we can effectively rotate the axes of the indicatrix as the molecules move, creating a moveable wave plate or optical retarder. This along with polarising optics makes the basis of phase modulation. If we have an optical indicatrix oriented at an angle of □ to the plane of the cell (usually this corresponds to the plane with the glass walls and ITO electrodes), then we can calculate the refractive index seen by light passing perpendicular to the cell walls. We can then calculate the retardance Γ of the liquid crystal layer for a given cell thickness d and wavelength λ.

$$\Gamma = \frac{2\pi d\left(n(\theta) - n_o\right)}{\lambda}$$

We can now use this expression in a Jones matrix representation of the optical LC retarder to get the optical characteristics of the LC material. In the case of nematic LC materials, there is little restriction on the flow properties of the material, hence it is possible to continuously vary Γ through several rotations of π providing a thick enough cell. The main drawback of these materials is due to the fact that the flow of the material means that the speed of the modulation is often very slow (10's of msec to seconds). It is by definition an out of plane electro-optical effect.

If the nematic were homeotropically aligned, then it would be perpendicular to the cell surface, In order for the LC to be moved by an applied electric field, it would have to have a negative dielectric anisotropy (such as in a VAN cell).

b) [30%] One of the main limitations of nematic LCs is that they are inherently polarisation sensitive when in the planar geometry. This comes from the fact that the optical anisotropy is only in two dimensions and the planar orientation is an out of plane effect.

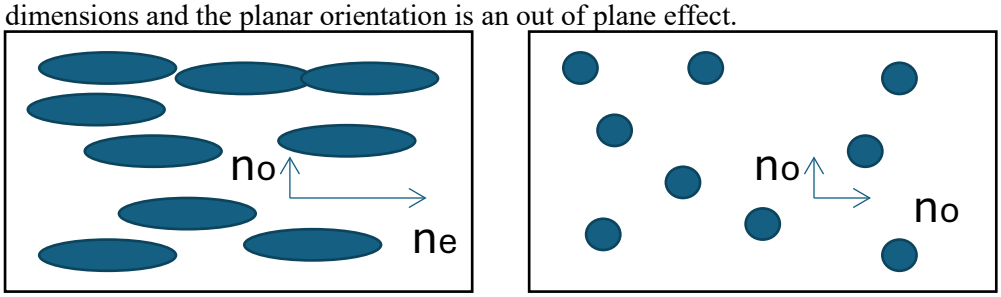

When seen from above in the 0V state, light polarised parallel to the long axis of the molecules sees ne and light perpendicular sees no. When the LC is switched into the homeotropic state by the applied

electric field, the light polarised parallel to the long axis of the molecules now sees no but the light perpendicular still sees no. Hence one polarisation is phase modulated, while the perpendicular polarisation is not.

The effects discussed above is the planar effect, where there is just the simple angle θ with respect to the cell walls. Because the nematic LC molecules are free to rotate in any position, it is possible to make more complex geometries such as twisted structures. In these devices, the molecules follow a twisted or helical path often rotating through an angle of 90º as is the case with twisted nematic or TN displays. This reduces the polarisation dependence as both polarisation states will see a change in reflective index when the LC material is switched, however it is not fully polarisation insensitive as each polarisation direction will see differing degrees of phase modulation. There is no twisting direction that will solve this inequality between the two polarisation states.

An FLC material is a possible alternative and in its ideal form is inherently polarisation independent. This independence means that the cell will either be lossy due to imperfect parameters or have a large zero order. The other main limitation s that FLCs are inherently a binary phase modulation material.

c) [20%] LCOS SLMs are reflective as they use a silicon backplane to support the mirror which reflect light through the LC layer. The LC is sandwiched between glass cover and the silicon backplane.

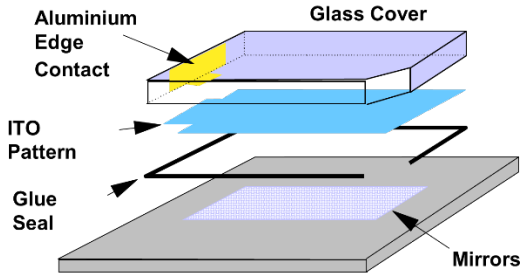

The use of the Si backplane means that the power of electronics can be used to interface with the SLM. This allow pixel arrays with very small pitches (<5um) and also allows them to be addressed at high speeds. Other features such as processing, demultiplexing and compression can be added to the electronics on the backplane.

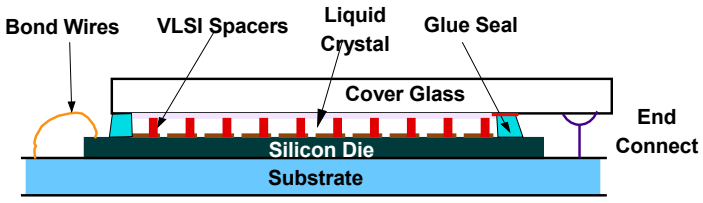

A quarter waveplate can be used to rotate the polarisation state of the light from linear to circular in transmission. If the quarter waveplate is used in a reflection, then the light will pass through it in both directions and both passes will combine to make if a halfwave plate, which can rotate a linear polarisation state by 90 deg. If the quarter waveplate is inserted next to the mirror in a LCOS device and it is oriented so that its axes are at 45deg to the planar alignment direction of the nematic, then it will rotate the horizontal polarisation to the vertical and vice versa. If the LC is [planar aligned I the vertical direction, then the vertically polarized component of the light entering the SLM will see the effect of the LC and the horizontally polarized light will see no change. The quarter waveplate then rotes the two polarisation states so that the vertical becomes the horizontal and vice versa. Hence as the light is reflected out of the SLM, the state that was horizontal entering is now vertical and will see the modulation of the LC. Over this reflection, both states get modulated and the SLM is now polarisation independent.

The main limit of this technique is the fact that making an LCOS device with an in built quarter waveplate is nearly impossible. The SLM would also have to be twice as thick to get the full effect of the retardance.

d) [20%] The penalty of using a nematic LC is that it is very slow. In the application of telecommunications switches there are some fault scenarios in optical networks where speed is not critical. In these sorts of networks the state of polarisation is unknown, hence the switch must be polarisation independent.

There ae also applications in optical tweezing and in wavefront control, where the polarisation state can lead to unwanted momentum effects, hence polarisation independence is sometimes useful here.

Well answered overall but quite a few of the answers rather waffled on about anisotropy rather than focusing on the orientation of the liquid crystal molecules. Part a) was answered well by most, but there were some quite cryptic answers. Only one candidate spotted that VAN materials would be needed for homeotropic alignment. For b) there were quite a few disjointed answers that did not connect the orientation of the polarisation with the planar alignment direction. Not many understood the limits of twisted nematic and a few spotted the fact that FLCs are binary. Part c) was not well answered, but it is slightly outside the main material, but it was covered in the examples class. For part d) most only got one application.

Q4 a) [25%] If two functions are multiplied together and then the FT taken, the result is the convolution of the FT of each function.

$$F_T\left[g(x,y)h(x,y)\right] = G(u,v) \otimes H(u,v)$$

If we replace *h(x,y)* or *g(x,y)* with their complex conjugate, then the result is the correlation of the two FTs.

$$F_T\left[g(x,y)^* h(x,y)\right] = G(u,v) * H(u,v)$$

Correlation gives us a means of comparing two functions. By correlating two functions with each other we can judge how similar they are in contents and structure. If one of the functions is a reference image and we explicitly know its entire structure, then we are comparing an unknown image with a reference to see if the contents of the reference image is contained anywhere in the unknown image. We define the reference image as *r(x,y)* and the unknown input image as *s(x,y)*.

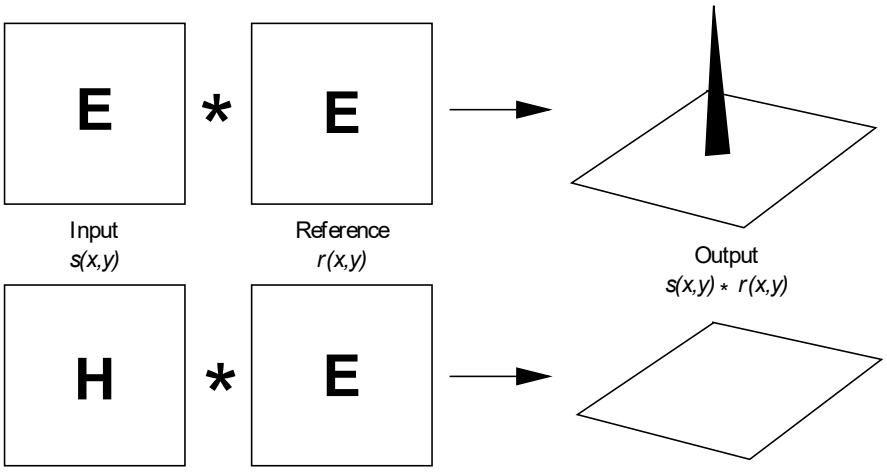

In order to perform a correlation, we need to multiply the FT of *r(x,y)* and *s(x,y)* in the Fourier domain before the final FT creates the correlation of the original images.

$$s(x,y) * r(x,y) = F_T\left[F_T\left[s(x,y)\right]F_T\left[r(x,y)\right]\right]$$

If the unknown input image *s(x,y)* contains the same information as the reference *r(x,y)*, then an autocorrelation (maximum SNR) occurs. If *s(x,y)* contains the same information as *r(x,y)* but shifted to a different location in the image, then the autocorrelation will also be shifted by a proportional amount. Hence,

$$s(x,y) = r(x - x_0, y - y_0)$$
$$F_T[s(x,y)] = S(u,v) = R(u,v)e^{-j2\pi(x_0u+y_0v)}$$

We then multiply by the FT of the reference *r(x,y)*, which we have centred on an origin or a known point of reference, and FT the product to give the correlation between the two images. The above equation has an exponential phase term, hence the autocorrelation will be shifted by the same amount as the reference object was shifted in the unknown input relative to position of the object in the reference image. With correlation we can not only detect if the object in *r(x,y)* occurs in *s(x,y)*, we can also state the position at which it occurs (if it occurs). Correlation is shift invariant.
In the input image *s(x,y)* may contain more than one reference object, in which case, both will correlate with *r(x,y)*. The result is two correlation peaks, each of which indicates if and where the reference object occurs in the input. For the sake of analysis of the correlation, we must consider each object in *s(x,y)* as a separate function, with separate positions.

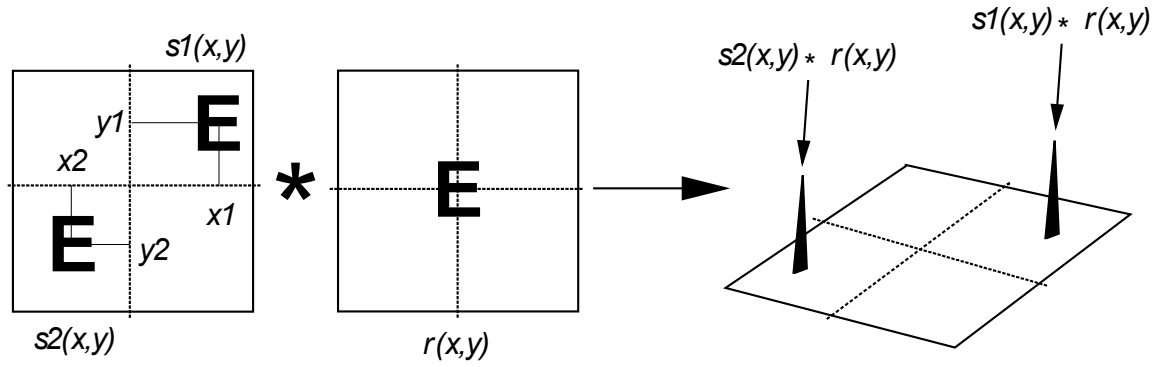

b) [25%]

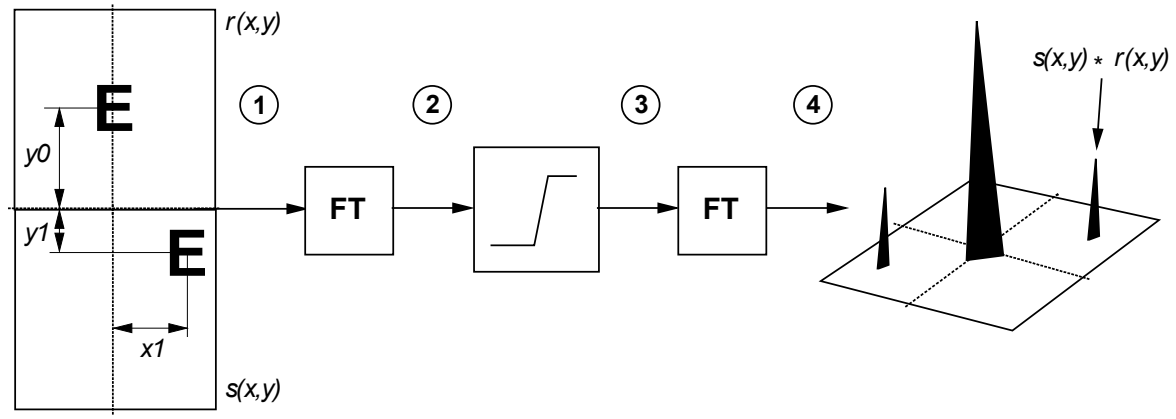

In plane 1, the input *s(x,y)* and reference *r(x,y)* are displayed side by side in an optical system and then transformed by a single lens into plane 2.

$$S(u,v)e^{-j2\pi(x_1u-y_1v)} + R(u,v)e^{-j2\pi y_0v}$$

The nonlinearity between planes 2 and 3 creates the correlation and in its simplest form can be modeled by a square law detector such as photodiode or CCD camera which takes the magnitude squared of the light falling upon it.

$$S^2(u,v) + R^2(u,v) + S(u,v)R(u,v)e^{-j2\pi\left(x_1u-\left(y_0+y_1\right)v\right)} + S(u,v)R(u,v)e^{-j2\pi\left(-x_1u+\left(y_0+y_1\right)v\right)}$$

The final plane 4 is after the second FT, with the central DC terms proportional to FT [$R^2 + S^2$] and the two symmetrical correlation peaks spaced by $(x_1, -(y_1+y_0))$ and $(-x_1, (y_1+y_0))$.

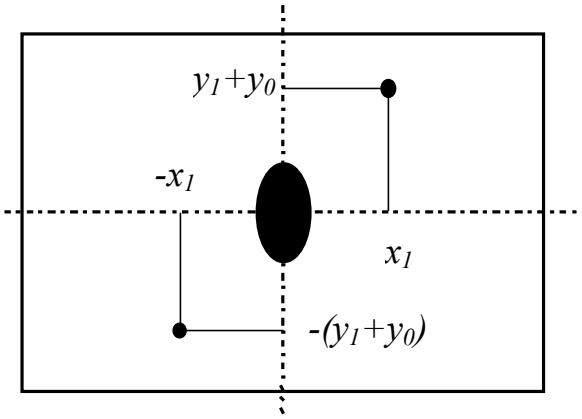

The central DC term is an unwanted source of noise, which degrades the optical system. There are always a symmetric pair of correlation peaks, so only half the output plane needs to be considered but the position of the correlation peak has to be decoded to gain the position of the reference object in $r(x,y)$ if it correlates with $s(x,y)$.

c) [30%] The JTC works on the basis of a non-linearity working on the spectrum of the input objects to create the product of the two Fourier transforms. This was modeled as a square law detector, but this gives undesirable broad peaks, but the better performance is when the spectrum is thresholded. Hence the JTC was originally built using an optically addressed SLM.

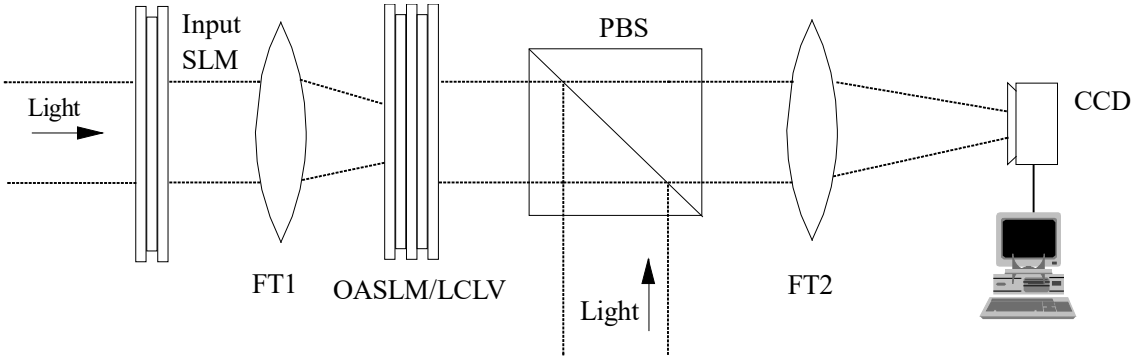

A better layout is to exploit the symmetry about the OALSM or non-linearity. If the optical system is split at this point, then the JTC just becomes two Fourier transforms and in fact can

be done with a single laser, SLM and camera by doing two passes through the Fourier transform lens. This is know as the 1/f JTC.

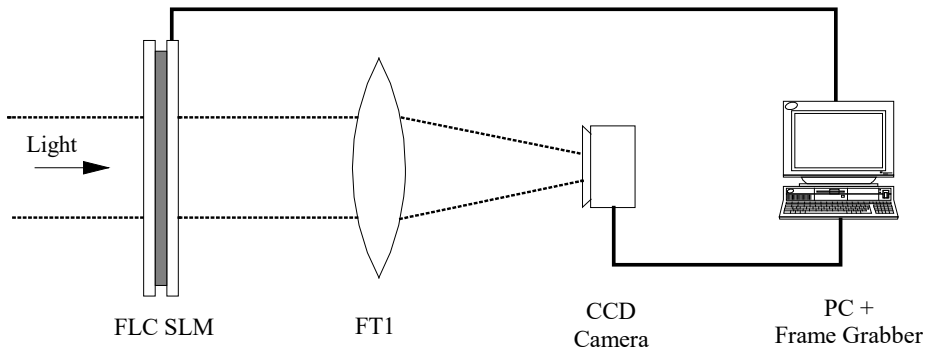

The input and reference images are displayed side by side on a FLC SLM as in a full JTC. The SLM is illuminated by a collimated laser beam and the images are Fourier transformed by a single lens in its focal plane. This spectrum is then imaged onto a CCD camera. The spectrum is then non-linearly processed by computer before being displayed onto the SLM again to form the correlation information. The 1/f JTC is a two-pass system, using the same lens to perform the second Fourier transform.

The quality of the correlation peaks and the zero order can be improved by non-linearly processing the spectrum that also suits the available SLM technologies. Some of the best results have been reported by using binary thresholds on the spectrum to improve the correlation peaks. A binarised spectrum produces good sharp correlation peaks and reduced zero order. If the binarised spectrum is converted to binary phase modulation [-1,+1], then the zero order can be reduced to around the height of the correlation peaks.

Advantages, simple, cheap and easy to align. Inherent power due to a programmable non-linearity

Disadvantages, slow (two pass), requires input preprocessing

d) [20%] In the architecture above, the SLM and lens are both sources of imperfection along with the alignment of the optical system. These will all lead to various forms of aberration which can be corrected using adaptive optics. If the flatness of the SLM and the les and alignment are known the aberrations can be evaluated and then out into the algorithm that will calculate the two passes of the JTC. This is not trivial, but it can be done once and then used as a look up table to compensate for any further images being processed by the system.

This question was probably the easiest of the four and is a common theme. This being said, there were still some quite convoluted answers to what should have been standard bookwork. Part a) was surprisingly variable in its answers, with many confusing different aspects of the correlation process. Only a handful stated that the function had to be real and many did not provide the full expression for the correlation of the functions rather than their spectra. Part b) was a more standard section with most doing well. Part c) was also well answered with most deriving the 1/f JTC from the symmetry and explaining its foibles. Part d) had many different answers and there was quite a spread of quite novel suggestions.