4B11 – Solutions 2022

Q1 a) [20%] A 2-D combination of pixels or apertures, used to form a diffraction pattern in the far field generated by a computer algorithm is defined as a Computer Generated Hologram and the pattern generated by the hologram if the far field is the Replay Field. The translation between the two is given by the Fourier transform.

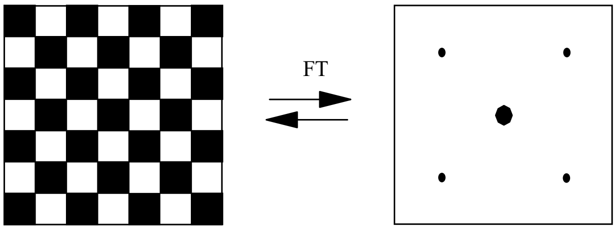

Hologram Replay Field

The far field of a single square pixel is its Fourier transform:

$$F(u,v)=\iint_{\pm\infty} f(x,y)e^{2\pi j(ux+vy)}dxdy=\int_{-a/2}^{a/2}\int_{-a/2}^{a/2} Ae^{2\pi j(ux+vy)}dxdy$$

$$=\int_{-a/2}^{a/2} e^{2\pi j(ux)}dx\int_{-a/2}^{a/2} Ae^{2\pi j(vy)}dy=\frac{A}{2\pi j}\left[\frac{e^{2\pi j(ux)}}{u}\right]_{-a/2}^{a/2}\frac{1}{2\pi j}\left[\frac{e^{2\pi j(vx)}}{v}\right]_{-a/2}^{a/2}$$

$$=\frac{A}{2\pi j}\left[\frac{e^{\pi jau}-e^{-\pi jau}}{u}\right]\frac{1}{2\pi j}\left[\frac{e^{\pi jav}-e^{-\pi jav}}{v}\right]=Aa^2\mathrm{sinc}(\pi au)\mathrm{sinc}(\pi av)$$

Assume infinite plane wave illumination, a perfect FT from the lens or free-space, no apodisation, dead space. The result is a sinc function which forms the envelope for the whole replay field of a CGH made with square pixels or apertures.

b) [30%] If the aperture is extended vertically to infinity and then the resulting stripe is then repeated horizontally every 2a positions it will form a vertically amplitude grating. If a pixellated pattern such as a grating is viewed from the end it can be modelled as a repetitive 1-D function. The repetition rate is defined by the pixel pitch or period.

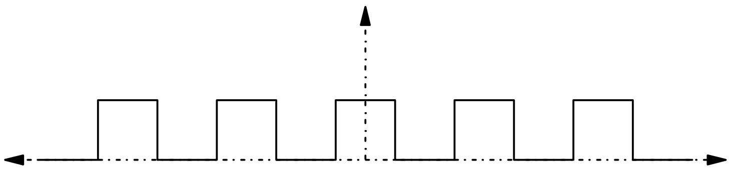

This can be expressed as a convolution of two functions.

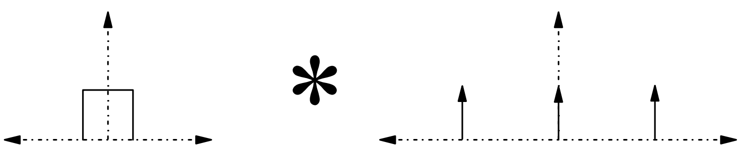

Where the delta function train represents the sampling or pixellation function and * represents a convolution. After the Fourier transform we have the replay field by Fourier analysis and the .

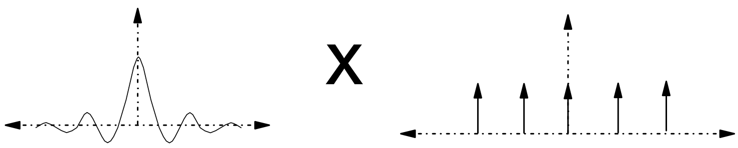

Gives the final result.

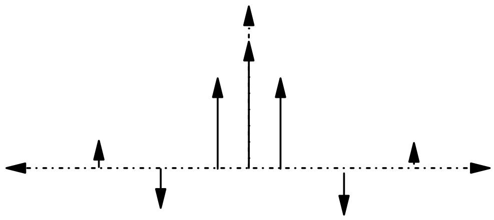

The train of delta functions has a sinc envelope and every second delta function is suppressed by the zeros of the sinc. Each symmetric pair of delta functions above represents a separate order and is repeated every odd harmonic. The envelope of the delta functions above are governed by the sinc function derived in part (a).

c) [30%] A 1-D grating $A \in [+1,-1]$ can be made by subtracting DC from a 1-D grating, $A \in [0,1]$.

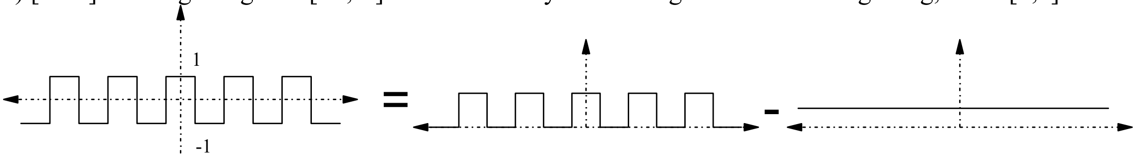

Hence in the Fourier domain we have:

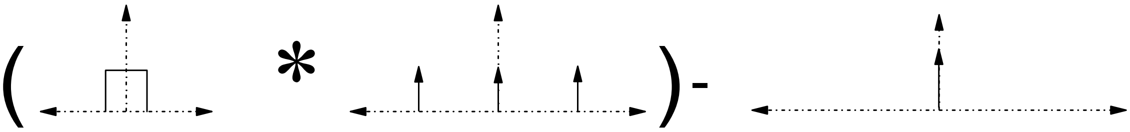

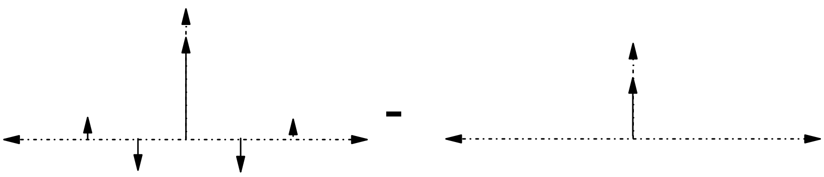

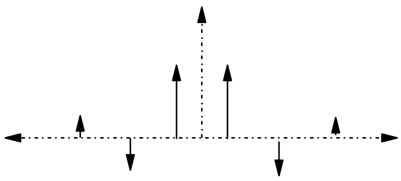

The sinc envelope means that the 41% of the amplitude is in the first order. The amplitude grating will block light, hence the estimated efficiency = 50% (amplitude)*0.41 = 20%. The binary phase grating diffracts all of the light, hence the far field diffraction pattern (no zero order) and the efficiency is now 41%.

d) [20%] A grating is an ideal candidate for a routing hologram as it is the optimal efficiency that can be achieved with binary phase (41%). Hence the loss in the switch will be minimised. Also, the noise from the grating is very discrete, which means that the crosstalk is very low in positions between the orders and their harmonics. A more subtle advantage of gratings is that they are simple to calculate and can be stored very efficiently. A SLM can be addressed directly with a grating without having to store any data other than periodicity.

The main drawback with gratings as routing holograms is that they are limited by their periodicity to discrete values. This has two main effects. 1) there are a limited number of locations in the RPF where light can be routed to, hence there are a limited number of fibre channels. 2) The number of channels is also limited by the crosstalk as the chances of picking up an unwanted harmonic or order gets higher with the number of gratings/holograms.

A compromise are Damman gratings which are more complicated than periodic gratings but can be calculated very easily and have discrete noise properties.

[Q1 was generally well answered, but not as well as expected as it is a fairly common structured question. A lot just quoted the result with no derivation. a) answered well, although quite a few forgot assumptions and only about half mentioned it was the sinc envelope for the RPF. b) was bookwork, and most did well, but quite a few were very messy and hard to follow. Very few pointed out that the grating was 2D and that the other dimension going to infinity gave a delta function. Not many picked the sinc envelop from part a). c) also well answered but again messy. Some good efficiency arguments. d) a wide variety of answers, most based on the lack of zero order. A lot argued that binary phase had 180degree symmetry, btu did not realise that binary amplitude would have the same symmetry argument.]

Q2 a) [35%] The calimatic molecular shape 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 and the short axis the ordinary n_o axis. The difference between the two is the birefringence. $\Delta n = n_e - n_o$

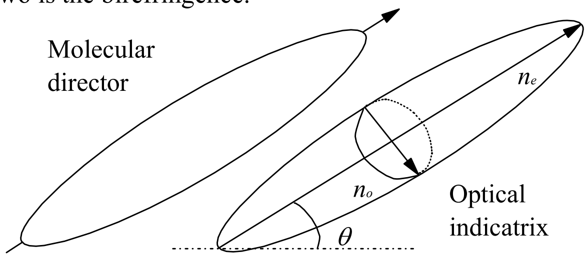

The combination of the flow allowing the molecules to move when an electric field is applied and the optical anisotropy means that we can effectively rotate the axes of the indicatrix as the molecules move, creating a moveable wave plate or optical retarder. This along with polarising optics makes the basis of most liquid crystal intensity and phase modulation characteristics.

We can then calculate the retardance Γ of the liquid crystal layer for a given cell thickness d and wavelength λ. In the unswitched pixel, the molecules are aligned parallel to the cell electrodes and the electric field will see a refractive index n_e. When switched into the homeotropic state the LC molecules are perpendicular n_o the cell walls and the E field now sees n_o. If we compare the two pixels after passing through the layer of LC then one E field will have travelled faster than the other. As a result there will be a phase difference (or retardance) Γ between the pixels.

$$\Gamma = \frac{2\pi\ t}{\lambda}\left(n_e - n_o\right)$$

If we have an optical indicatrix oriented at an angle of θ 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.

$$n(\theta) = \frac{n_o n_e}{\left(n_e^{\ 2}\sin^2\theta + n_o^{\ 2}\cos^2\theta\right)^{1/2}}$$

We can then calculate the retardance Γ of the liquid crystal layer for a given cell thickness d and wavelength λ. The retardance is the phase difference between a wave passing through the short axis and the wave passing through the material oriented and an angle θ.

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

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.

Assumptions: Monochromatic, coherent light, all the LC rotate with the angle θ, The E field is aligned parallel to the long axis of the LC,

b) [20%] In the generalised Jones matrix, the orientation of the extraordinary axis is defined with respect to the y axis. If the extraordinary axis is parallel to the x-axis this corresponds to an angle of $\psi = 3\pi/2$ (positive angle s are clockwise from the y axis). For quarter wave $\Gamma = \pi/2$. This gives the Jones matrix.

$$e^{\pi/4} = \frac{1}{\sqrt{2}}(1+j) \qquad W_{1/4} = \frac{1}{\sqrt{2}}\begin{bmatrix}(1+j) & 0 \\ 0 & (1-j)\end{bmatrix}$$

c) [30%] Set up the Jones matrix for the horizontally polarised light.

$$Y = \begin{bmatrix}1 \\ 0\end{bmatrix}$$

Take the product with the quarter wave plate matrix (twice) in series

$$Y = \frac{1}{\sqrt{2}}\begin{bmatrix}(1+j) & 0 \\ 0 & (1-j)\end{bmatrix}\frac{1}{\sqrt{2}}\begin{bmatrix}(1+j) & 0 \\ 0 & (1-j)\end{bmatrix}\begin{bmatrix}1 \\ 0\end{bmatrix}$$

Which gives the result (working expected in the answer)

$$Y = \begin{bmatrix}j \\ 0\end{bmatrix}$$

If we repeat the calculation with the vertically polarised light

$$Y = \begin{bmatrix}0 \\ 1\end{bmatrix}$$

We get the result (with working)

$$Y = \begin{bmatrix}0 \\ -j\end{bmatrix}$$

The two waveplates in fact combine to make a single half wave plate equivalent as each input polarisation state is maintained from linear to linear. If the pair were rotated, then the output polarisation would be rotated but remain linear which is a form of phase modulation (polarisation dependent). In the context of a nematic LC this is a means of negating polarisation dependence. If a quarter wave plate and a mirror are combined with a nematic LC, then the reflection through the quarter wave plate acts like a half wave plate and rotates the polarisation with respect to the nematic LC molecules.

d) [15%] When using a nematic LC as a phase modulator, the range of 0 to 2π is expected in order to create good balanced holograms in the replay field. This puts a lot of pressure on the materials properties and the birefringence and thickness of the cell limit Γ and limit the overall phase depth of the SLM. By providing more phase depth than the required 2π, we can easily keep the balance of out hologram without putting too much pressure on the drive electronics to maintain the full range. It also avoids the more non-linear portions of the response curve and makes for a better and more controllable series of phase levels.

[Q2 well answered overall which was pleasing as it was the most challenging of the exam. a) was NLC bookwork with a few silly errors. Many forgot that the indicatrix is in fact a 3D structure and that it was polarisation sensitive, hence it assumes the input polarisation parallel to n_e. b) well answered although only one candidate realised that the angle was -π/2 or 3π/3 (not π/2) but this did not effect the result. c) most got the result right unless they got b) wrong, but this was carried through ok. Very impressed by the fact that the analysis was interpreted to be a solution to polarisation sensitivity, especially after many had failed to point this out in the assumptions of part a). d) well answered with most getting the importance of the linear region in the Γ characteristic.]

Q3. a) [30%] the *1/f* JTC is based on the classical JTC but without the limitations of the nonlinear optical components (such as an OASLM). If the optical system of the classical JTC is split between the photodetector and the modulator in the OASLM, then the JTC just becomes two Fourier transforms and two image captures and 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 choice of binary phase means that the SLM should be using FLC which gives a much faster frame rate ad therefore correlation speed.

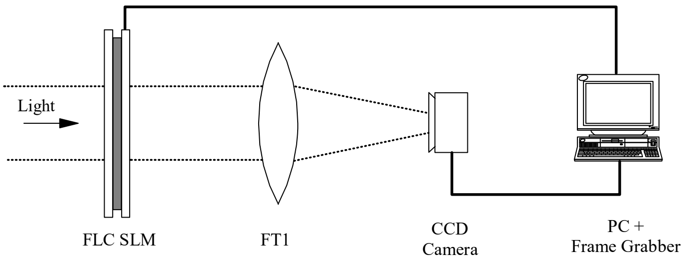

The input and reference images are displayed side by side on a FLC SLM as in a full JTC. The SLM is illuminated by a collimated laser beam and the images are Fourier transformed by a single lens in its focal plane. This spectrum is then imaged onto a CCD camera. The spectrum is then non-linearly processed in software 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.

If the spectrum was directly Fourier transformed, the result would be the two symmetrical correlation peaks characteristic of the JTC along with a huge zero order located in the centre of the output plane. The quality of the correlation peaks and the zero order can be improved by non-linearly processing the spectrum that also suits the available FLC 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.

A variety of processing schemes have been tried. The first technique was to use a basic median threshold filter on the spectrum, but this gave a very poor bitwise average after the threshold to binary phase. A 3x3 convolution binarisation scheme was also tried and gave reasonable results, but the processing time was slow due to the complexity of the filtering algorithm. The success of this algorithm is due to fact that the 3x3 convolution is a form of edge enhancement, which enhances the spectrum and the noise in the dark areas.

One of the main limitations of this architecture is actually the camera and it has limited frame rate (at low cost) and has limited dynamic range (usually 8-10bits max).

b) [25%] This is the basis of a head tracking system. A camera is mounted on the helmet of the pilot and is pointing at the cockpit area, hence the pilot's view of the cockpit is fed frame by frame to the correlator. Each frame is fed in sequence through the correlator, with the first frame being used as the reference and the next one as the input, then the next one becomes the reference and the third the input and so on… Thus motion between frames can be tracked as long as there is an overlap of common information in each frame from the helmet.
If the pilot's view (and therefore head position) does not change from one frame to the next, then the correlation between frames will be in the centre of the output plane. If the pilot's head rotates, then the view of the cockpit will shift from one frame to the next and the correlation peaks will also shift proportional to how much of the pilot's view has shifted from one frame to the next. This is shown below along with the correlations from an unshifted and shifted pair of frames.

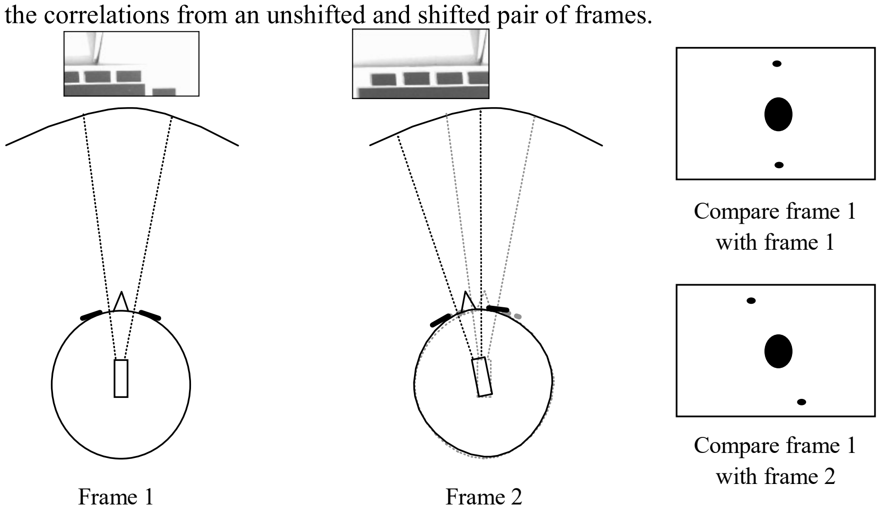

The position of the correlation peaks tells us how much one frame has moved from the previous one when the pilot's head rotates, hence by measuring the distance and knowing the dimension between the pilot and the cockpit, we can sense how much the head has rotated and therefore track it.

c) [30%] The view of the pilot from one frame to the next must overlap by more than 20% in order for the correlator to produce reliable peaks.

The 20º field of view of the camera corresponds to roughly 2*1.3*tan(20) = 0.45m of the cockpit panel. The 20% overlap area is 0.09m along the cockpit panel.

This corresponds to 90mm which is the maximum distance per correlation cycle along the cockpit panel the view can take before it goes beyond the 20% overlap limit.

The 90mm corresponds to an angular resolution of 4º about the pilot's head centre per correlation cycle.

If the correlator operates at 50 correlations per second, then the maximum angular speed of the pilots head must not exceed 200º per second (3.5 rads^{-1}) to keep accurate track of the position.

Assumptions – pilots head motion is purely rotational and not up and down, there is sufficient detail in the cockpit for the overlapping images to correlate

d) [15%] When the pilot looks up through the canopy they will be looking at either the sky or the ground depending which way up they are. The sky is rather feature-less for correlation, hence the head tracker will struggle to register the position of the pilot's view. If they are looking at the ground, then it will be moving wrt the canopy, hence it will confuse the correlation process even further.

The simplest solution is to make sure that the view of the camera on the pilot's head never leaves the cockpit console panel during any of the pilot's possible head positions, hence it will always have a fixed reference to correlate from frame to frame.

A second solution is to put a pattern on the canopy which can be picked up by the camera, but not by the pilot. Perhaps something holographic?

[Q3 less popular question, but well answered. a) mostly book work and most got the details of the 1/f JTC right. A few forgot about the two passes and a few mentioned chequerboard encoding. b) mostly done well, although quite a few could not explain how the cockpit view translated into correlation shifts via the head rotation. A few random diagram, bit most did well. c) surprisingly few got this right. Most made wrong assumptions about the rotation overlap as an angle rather than a percentage of the image. d) ok given it was off piste. A few got the fact that the view would be sky.]

Q4 a) [25%]

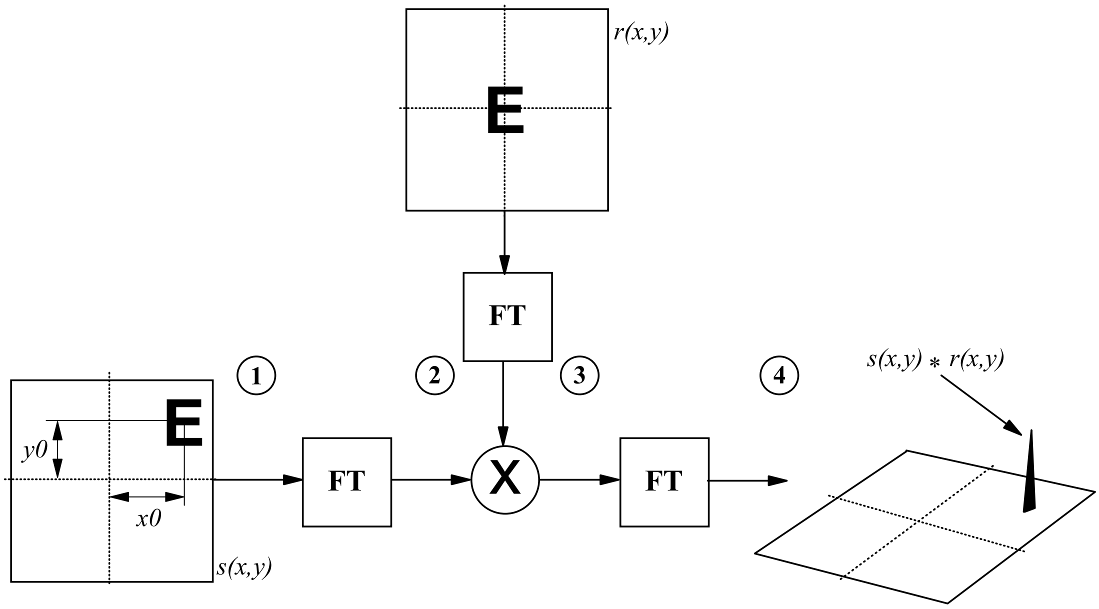

The input image *s(x,y)* is displayed in plane 1 before the FT into plane 2.

$$S(u,v)e^{-j2\pi(x_0u+y_0v)}$$

The FT of *s(x,y)* is then multiplied by the FT of the reference *r(x,y)*.

$$R(u,v)S(u,v)e^{-j2\pi(x_0u+y_0v)}$$

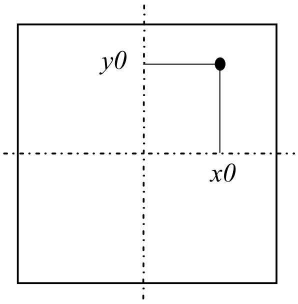

The FT of the reference is done off line on a computer and is defined as the matched filter *R(u,v)* for that particular reference *r(x,y)*. The product of the input FT and the filter then undergoes a further FT to give the correlation in plane 4.

Advantages: the object in the reference *r(x,y)* is centred in the process of generating the filter *R(u,v)*, so that if a correlation peak occurs, its position is directly proportional to the object in the input image, with no need for any decoding. Unlike in the JTC, there is only one correlation peak and there are no DC terms to degrade the correlator output and it can handle multiple objects with ease. It is also very easy to include invariance features such as SDFs or transformations into the filter to widen the functionality of the matched filter

Disadvantage: optomechanics and alignment of filter and FT are difficult. Requires 2 SLMs

b) [30%] The filter is matched to the reference via a Fourier transform such that.

$$F(u,v) = F_T\left[r(x,y)\right]$$

The best test for any matched filter is to perform an autocorrelation with the filter that has been generated. If the matched filter above is used for the reference image of a letter E, then the autocorrelation will have optimum SNR. The autocorrelation peak is very broad and has a huge SNR, as there is no appreciable noise in the outer regions of the correlation plane. Such a filter is not very useful for pattern recognition. Such a broad peak could lead to confusion when the position of the peak is to be determined. Also, similar shaped objects (such as the letter F) will correlate well with the filter leading to incorrect recognition. Another identical E which is placed in the input along with the original one will also cause problems as the correlation peak will take an extremely complex structure. Finally, the filter is a complex function and there is no technology available to display the filter in an optical system.

Great improvements can be made to the usefulness of the correlation peak, by using a phase only matched filter (POMF). The matched filter *F(u,v)* is stripped of its phase information (i.e. the phase angle of the complex data at each pixel) and this is used as the filter in the correlator.

$$F(u,v) = F_{amp}(u,v)e^{i\phi(u,v)}$$

Hence

$$F_{POMF}(u,v) = e^{i\phi(u,v)}$$

The autocorrelation for the POMF is much more desirable even though there is a reduction in the SNR due to the increase in the background noise. The correlation peak is much narrower which is due to the information which is stored in the phase of the matched filter. The POMF is the most desirable filter to use as it has good narrow peaks and still remains selective of similar structured objects (like Fs). The POMF is also a complex light modulation scheme, so the problems associated with binary phase (180º symmetry) will not occur. However, the continuous phase structure of $\phi(u,v)$ means that it cannot be easily displayed in an optical system. A POMF could be displayed using a nematic LC, but these materials are inherently slow and would limit the correlation rate.

The penalties associated with going to binary phase are greatly outweighed by the advantages gained by using FLC SLMs in the optical system. The binary phase is selected from the POMF by two thresholds δ_1 and δ_2.

$$F_{BPOMF} = \begin{cases} 0 & \delta_1 \le \phi(u,v) \le \delta_2 \\ \pi & \text{Otherwise} \end{cases}$$

1) The SNR is up to 6dB worse than the case of the POMF.
2) The filter cannot differentiate between an object and the same object rotated by 180º (due to the fact that the BPOMF is a real function).
3) The BPOMF is not as selective as the POMF due to the loss of information in the thresholding.

A FLC SLM is capable of multiple kHz frame rate which greatly improves the correlation rate.

c) [30%] The SDF provides a way of achieving a limited form of invariance to the effects of image rotation and scaling. The idea is to take a series of reference images $r_1(x,y)$ to $r_n(x.y)$ and combine them by a linear summation. The final composite SDF image contains all of the references and so will correlate with the input image. Unlike mathematical transform techniques, the SDF does not destroy information and still retains the ability to be shift (position) invariant.

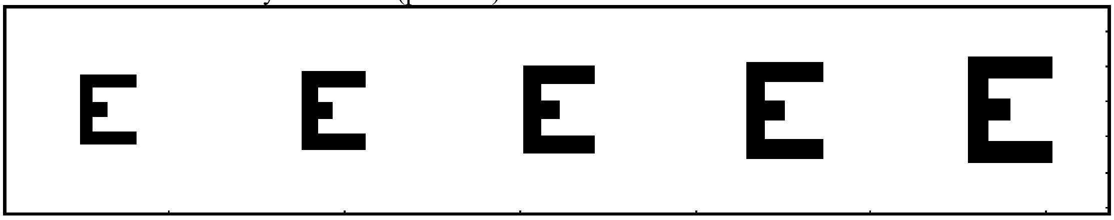

If we define a *1xn* vector **a**, whose elements are the weighting coefficients for each respective reference image, then the SDF is defined as.

$$h_{SDF}(x,y)=\sum_{n}^{i=1} a_i r_i(x,y)$$

We can easily calculate the weight coefficients from the cross correlations between all the reference images and the correlation peak value desired for each reference.

$$\mathbf{a}=\mathbf{R}^{-1}\mathbf{c}$$

Where **c** is a *nx1* vector whose elements are the desired correlation peak values for each image correlated with the SDF and **R** is the *nxn* correlation matrix. The elements of **R** are calculated from the cross correlation of the images in the reference set. Hence the element $R_{i,j}$ of the matrix **R** is the correlation between the reference images $r_i(x.y)$ and $r_j(x.y)$. One of the challenges in designing a SDF is the choice of data set used to include the invariance. The number and range of scales chosen across the set will greatly effect the ability for the filter to recognise a range of scales.

When the SDF is used in a matched filter and displayed as continuous phase only filter then the correlation peaks are reasonably close to the expected value of 1.0. If the phase only filter is then thresholded to form a binary phase only filter then the results are much worse. The correlator cannot distinguish between similar objects. This is the main drawback with the SDF, as the more references that are included, the more difficult it is to display the SDF on a modulator in a realistic optical correlator.

d) [15%] The main source of optical aberration is such a miniaturized correlator will be from the lenses, especially the first FT lens after SLM1. Once the design and build of the correlator has been fixed, a ray tracing software package such as Zemax could be used to evaluate the aberrations and then the lenses could be modified to minimize them. This would be expensive and require multiple surfaces per lens which might also limit compactness. A better technique would be to include the aberration in the design of the matched filter so that it compensates for the aberrations, especially if optimization algorithms such as the SDF are being used in the correlation filter design.

A second approach might be to measure the aberration in situ either with a wavefront sensor such as a Shack Hartmann or using the SLMs themselves to display holograms that would allow the aberrations to be ascertained. It maty even be possible to do this in a continuous fashion using the final camera of the correlator which means that the system could adapt for aberrations that might occur due top components tolerances of positional variations of the optics in the system.

[Q4 was very well answered by most as it was the least speculative. Parts a) and b) were ok although a few candidates described DBS rather than phase only filters. Most mentioned the mechanical limitations of the matched filter. c) was also well answered with most candidates knowing how the SDF worked, although less realized the penalty of converting it to a BPOMF. d) also ok, with quite a few just describing an AO system rather than how aberrations could be detected and corrected in the matched filter architecture.]