PhD University of Edinburgh 1992

Phase-Only Optical Information Processing

D. J. Potter

Chapter 7
Optical Correlation: Groundwork.

Introduction

In the preceding chapters the operation and (improved) assembly of a 16×16 pixel spatial light modulator has been described in detail. This work was directed to obtaining a device of the highest possible optical quality realizable within the timescale of this project, for evaluation as a binary phase-only optical correlation filter.

The purpose of this chapter is to lay the essential groundwork for the experimental results of the final chapter. Section one of this chapter reviews the choices one must make in order to accurately represent the optical correlation system[1] on a computer and describes the computational background of the problem. Section two details the fabrication of the target objects used in the experiment and section three evaluates the information extraction techniques which were used to characterise both simulated and experimental results. A redefinition of one such parameter, optical efficiency, based upon experimentally measurable quantities is also proposed.

1 Computational Framework

1.1 Optical Requirements

The optical phenomenon to be represented on the computer is first described in this subsection. This shall provide one or two very important observations which will considerably aid discussion of the computational approach. In the frequency plane of an optical processor, if x_t denotes physical distance, ν_x the spatial frequency along the x-axis and F_L the focal length of the transform lens used then it has already been noted that

$$\nu_x = \frac{x_t}{\lambda F_L} \qquad (1)$$

The mirrors of an SLM are spaced $\delta\nu$ apart in terms of frequency, so that upon retransformation the *sampling* interval is $\delta\nu$ and, from the sampling theorem, the replicated images will be centred [$1/(\delta\nu)$] apart in the image plane which must also be their maximum extent so as to avoid overlapping. Writing the actual expression for the image field (co-ordinates x_i) explicitly then,

$$f(x_i) = C \int_{-\infty}^{\infty} F(x_t)\, e^{-2\pi i\, [(x_t x_i)/(\lambda F_L)]}\, dx_t \qquad (2)$$

where `C' is a constant, identifies the frequency variable as

$$\nu_x = \frac{x_t}{\quad} \qquad (3)$$

$$\lambda F_L$$

which is the origin of equation 7.1. If the mirrors are separated by a physical distance Δ then

$$\delta\nu = \frac{\Delta}{\lambda F_L} \tag{4}$$

which results in a maximum object width W_o of

$$W_o = \frac{\lambda F_L}{\Delta} \tag{5}$$

In the actual experiment a 600mm focal length lens was used and the mirror spacing is 0.2mm, so that W_o=1.898mm, although considering the likely error in lens focal length it is better to set W_o=1.9mm. This choice of focal length followed visual inspection of the size of the spectrum in the frequency plane for a square object of this size, where it was estimated that roughly 70% (visual estimation) [2] of the energy lay over the active area of the SLM to be used.

The preceding calculation follows a similar line of argument as the one which led to the definition of the space-bandwith product in chapter four, but it is not the intention here to repeat that calculation. Instead, for numerical simulation, one wishes to know where the object spatial frequencies are mapped to in the frequency plane. Where, for instance, is the fundamental frequency of an object of width W_0 as just determined, mapped to ? The fundamental frequency is defined as

$$\nu_0 = \frac{1}{W_0} \tag{6}$$

which, by equation 7.1, is mapped to a physical distance from the zero frequency of exactly Δ. Thus the mirror spacing determines the maximum size of object which may be input to a standard optical processor to avoid overlapping of the replicated images. If that size of object is used, the fundamental spatial frequency is mapped to the first mirrors either side of the zero frequency (central mirror) of the SLM. This result is fundamental to all numerical simulation work and, in its simplicity, shall prove an invaluable guide.

1.2 Numerical Simulation

Numerical simulation is based around the properties of the Discrete Fourier Transform or `DFT'. The maximum spatial frequency which can be represented by a DFT is the *Nyquist* frequency, ν_{NY}=[1/(δx)] where δx is the sampling interval (in terms of physical distance) in the object space. The question of exactly *where* this frequency is located within the output array from the DFT algorithm requires a moment`s clarification. Frequently the output of the algorithm stores the information on the spectrum in what is commonly termed `computer' format which, if displayed, would not resemble the optical transform[3]. Re-arrangement is a simple matter. Once in `optical format', the zero frequency component lies at J=[N/2]+1 where J denotes the pixel number in the DFT plane as illustrated in figure 7.1.

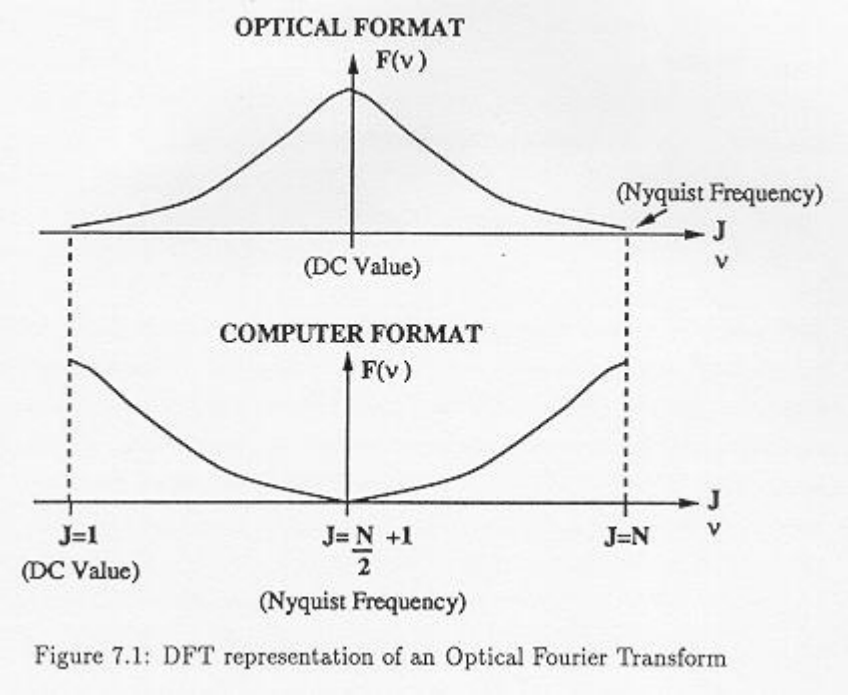

Figure 7.1: DFT representation of an Optical Fourier Transform

Figure 7.1: DFT representation of an Optical Fourier Transform

The target object may be written onto either the whole of the object space or a subsection of that space, the precise choice of which has strong implications for filter computation in the frequency plane. To see this, a fundamental property of the DFT, periodicity, shall now be studied with reference to an N point Fourier Transform, conducted in 1-D for simplicity.

Mapping in the DFT

The DFT, being the Fourier Transform of the *sampled* function, is therefore a periodic function by virtue of the sampling theorem, and this is true no matter which space - object or frequency - is used as input to the algorithm. If the object is written on all N pixels of the object space then the entire frequency space displays precisely one cycle of the DFT. In this case, the fundamental object spatial frequency is mapped to a position exactly one pixel either side of the zero frequency.

However, consider the action of the DFT algorithm when used to transform from the frequency plane to the image plane. In direct analogy to the optical calculation of equations 7.3 to 7.6, it is straightforward to show that if the mirrors are spaced n_m pixel separations apart, so that the sampling interval is n_m pixels, then the image plane shows several cycles of the DFT (each of which is an upside down image of the object function) separated by an interval of $[N/(n_m)]$. In order to avoid overlapping of the replications this must also be the maximum allowable object size n_o, and the general relationship is

$$n_o = \frac{N}{n_m} \tag{7}$$

Again, by an analogous procedure to that employed in the optical case, it is straightforward to show that an object written onto a subspace of the whole object space of size $n_o \times n_o$ has the fundamental spatial frequency (of wavelength n_o) mapped to

$$J = \left(\frac{N}{2} + 1 \right) \pm \frac{N}{n_o} \tag{8}$$

in the frequency plane. Figure 7.2 illustrates the relationship of equation 7.8 with the example of a 256 point DFT.

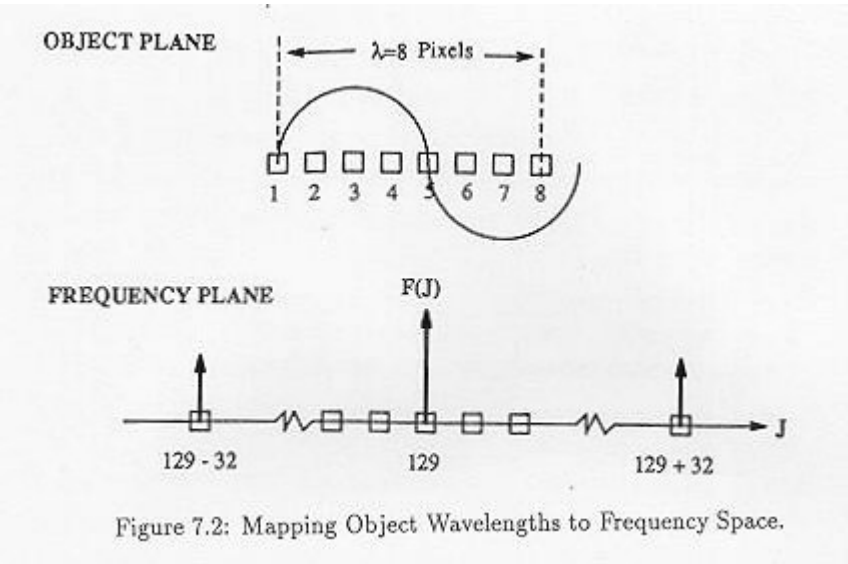

Figure 7.2: Mapping Object Wavelengths to Frequency Space.

Figure 7.2: Mapping Object Wavelengths to Frequency Space.

An accurate representation of the optical system requires the mirrors to be spaced [$N/(n_0)$] pixels apart if an object subspace of n_0 pixels on a side is used. Table 7.1 lists several possible mirror separations and the resulting maximum object width.

N	n_m	$n_o=[N/(n_m)]$	$16\times n_m$
256	4	64	64
	8	32	128
	16	16	256
512	4	128	64
	8	64	128
	16	32	256
	32	16	512

Table 1: Mirror spacings n_m and corresponding object subspace size n_o (pixels).

1.3 Resolution Requirements

Equation 7.8 is analogous to the space-bandwidth product equation for the optical situation, and allows variable resolution in object and frequency spaces. The small object size allowable by the small space-bandwidth product of the SLM translates, experimentally, to a low resolution object due to the difficulty involved in fabricating test objects only 1.9mm on a side wide, which is further discussed in section two. One would further wish as many DFT pixels to cover each mirror in the frequency space as possible so that the light field is adequately sampled. Therefore in optical correlation simulations and filter calculation the limited space-bandwidth product of the 16×16 SLM naturally favours low object space and high frequency space resolutions.

The question of exactly how many mirror sample points are needed to accurately determine the characteristic phase has been much neglected in the literature. Hossack [] has suggested that the sample interval should be such that the wavefront over the mirror is at least optimally sampled. The problem is then to determine what this sampling interval actually is, which is where the Point Spread Function (PSR) of the lens enters the discussion. The smallest resolvable feature of the wavefront in the frequency plane is determined by the PSR of the Transform lens. If one were to imagine an experiment whereby the amplitude and phase of a number of points over a mirror were actually measured, it would seem reasonable to take measurements at an interval not smaller than the resolution limit of the lens which, for an abberation free lens, is given by $1.22\ \lambda\ F_{\#}$ where $F_{\#}$ is the f-number of the transform lens. Measurements made at two points separated by less than this distance will be strongly affected by the PSR of the other point. As this is the smallest feasable interval one might measure and

obtain results free from the effects of the PSR, it is suggested that this distance (or slightly greater) be the optimum sampling interval.

In the simulations used in this project a 256×256 array was used for the DFT routine which results in a 9×9 sub-array of the DFT covering each mirror region. The f-number of the transform lens was 12 and illumination with a He-Ne laser gives the smallest resolvable interval as 9.26μm. As each mirror is 100μm on a side, a 9×9 grid of sampling points provides an on-axis sample separation of 12.5μm. Given that the lens is unlikely to have been 100% aberration free, this figure shows that the wavefront in the frequency plane was sampled as finely as possible.

From table 7.2 a 256×256 DFT array size allows the computer representation of the SLM to fill the entire frequency space whilst providing an object resolution of 16×16 pixels. Although the 512×512 array size can provide better resolution in both object and frequency space, it is computationally intensive and requires an extensive amount of data storage space and was rejected in favour of the smaller, but adequate, 256×256 array.

It is desirable to centre the SLM in the Fourier plane so that a mirror lies over the zero frequency component of the spectrum. In order that this be achieved with a device having an even number of mirrors, it is neccessary to offset the device up and to the left in the frequency plane by one mirror. The relative sizes of object and frequency subspaces used are shown in figure 7.3. Therefore the leftmost column and uppermost row of the SLM are of no use in practical filtering operations. From the choice of transform scaling (lens focal length) very little energy should lie over the outer regions of the SLM so that it is considered that any effects on the correlation from a binary phase value chosen from a reduced data set on these mirrors will be further minimised.

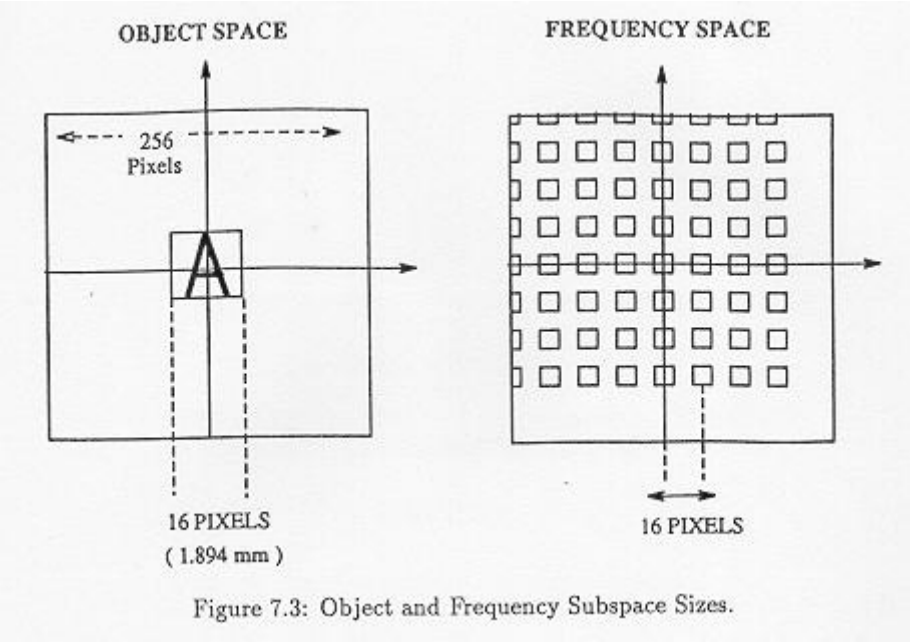

Figure 7.3: Object and Frequency Subspace Sizes.

Figure 7.3: Object and Frequency Subspace Sizes.

Mirror Boundaries

In the simulation, SLM mirror rows are numbered from zero to fifteen starting from the top row, and columns from zero to fifteen starting from the left. This convention follows the way a FORTRAN binary file is written where the first 2-D array element written to a file happens to be the top left one, proceeding a row at a time until the bottom right element is written. The central mirror thus lies at location (8,8) in this co-ordinate system. The procedure used to isolate the data to be used in calculating the binary phase of each mirror requires the pixel boundaries of each mirror to be determined.

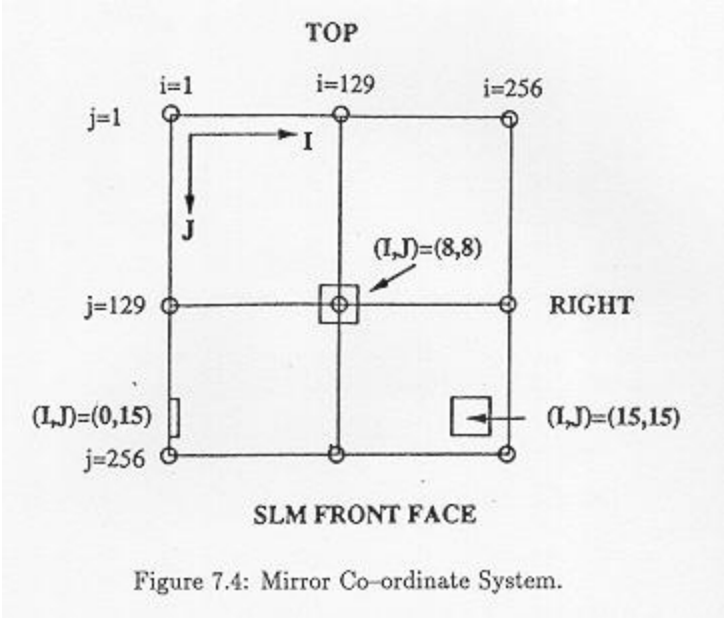

Figure 7.4: Mirror Co-ordinate System.

Figure 7.4: Mirror Co-ordinate System.

A FORTRAN program, incorporating many of the equations included thus far in this chapter was written to this end. The single parameter input is the focal length of the transform lens, which defines maximum object width and the frequency-distance relationship of the Fourier plane uniquely. As the physical positions of the mirrors are known precisely, it is a simple matter to increment the pixel number in the discretised distance-frequency relation and thus find out if which pixel numbers lie within which mirror. The program MIRMAKE.FOR is included, together with a list of actual mirror pixel locations, in appendix eight. A 50% fill factor cannot be attained in this simulation for the following reasons:

1. **In order that the computational representation of the SLM be self consistent each mirror should be identical in size to the central mirror.**

2. **For the central mirror to have a sample point at the origin of frequency *and* the mirror to be symmetrical about the origin the central mirror must have an odd number of sample points.**

3. **Mirror centres are determined by the object subspace coverage n_0 by equation 7.8 and are 16 pixels apart in this simulation. The closest situation to a 50% fill factor attainable is for the mirror to have either 9 or 7 pixels - one more or less than 8 pixels by requirement 2 above. From the sampling theorem, sinc interpolation between the datapoints would form a continuous reconstruction of the SLM, the mirror boundaries of which would lie midway between the outer mirror pixels and the first circuitry pixels. Comparison of the mirror widths so determined with those of the actual SLM reveals that a 9 pixel mirror has boundaries lying fractionally within those of the physical device. For this reason the 7 pixel mirror simulation was not used.**

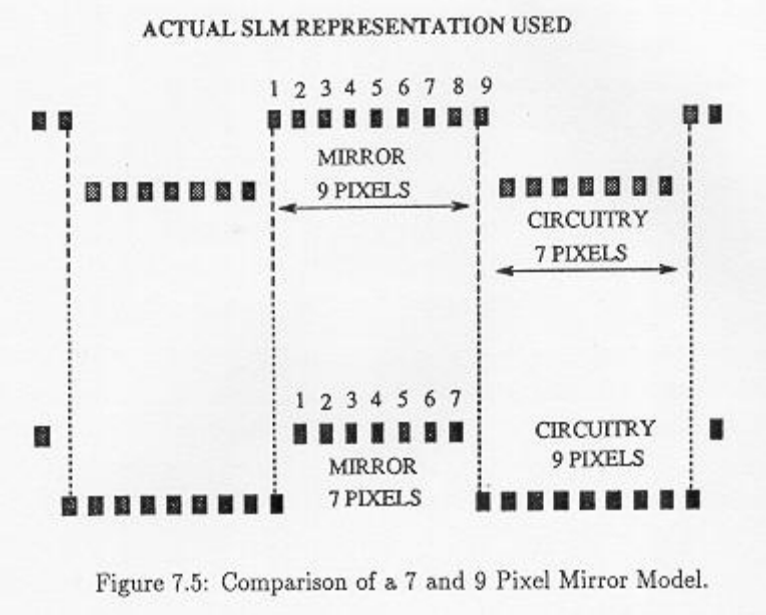

Figure 7.5: Comparison of a 7 and 9 Pixel Mirror Model.

2 Correlation Target Objects

In chapter four, it was quoted that binarised input targets resulted in the best autocorrelation results for a binary phase-only filter. The limited space-bandwidth product of the device used in this project, coupled with the need for accurate computational modelling of the physical target, suggest that low resolution binary target objects be used. Binary objects are easy to create as a data file with zeros or ones, and are further easy to fabricate, as this section shall describe.

It is common for assertions to be made on the merits of particular correlation systems based on the simulation results using but a single target object. Whilst it is quite plausible that the generalisations extrapolated are true to a high degree, this approach would not suit the requirements of this project. Several algorithms will be evaluated both using simulated and experimental results. There is much debate about the suitability of certain filter algorithms as shall be discussed more fully in chapter eight. Here it is noted that the degree of symmetry of a target object may radically affect the resulting auto-correlation, so that one should bear in mind that the particular object used may be biasing the results.

The target objects used were photographic transparencies and it is likely that the fabricated target objects as used in this project suffer from a degree of phase noise. In order to account for this, and to guard against drawing conclusions based on a single object, it is proposed that a range of different target objects be used for each filter algorithm under observation and that the *overall* properties of the resulting correlation set be compared. With larger SLM arrays, such as the 50×50 device, it will be possible to present a separate target in each quadrant of the object space and obtain both cross and auto-correlations simultaneously. However, due to the very small size of the object allowed, 1.9mm, only one distinct target can be described per object space with the 16×16 SLM. The work of this chapter therefore very much paves the ground for the bigger arrays to come.

2.1 Fabrication

A total of seven letters and symbols, shown in figure 7.6, were used in this experiment. Each symbol was written onto a background of a 16×16 pixel set and in order that each transmit an identical amount of light, comprised of 76 ones against a background of zeros.

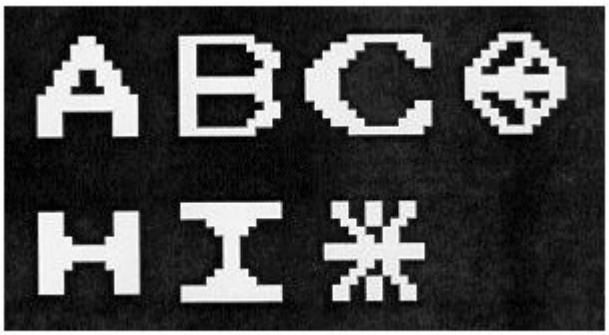

Figure 7.6: The Seven Target Objects.

Figure 7.6: The Seven Target Objects.

For the required degree of resolution of object and frequency space discussed earlier, the object is represented on a small array of discretised points. Such is the computer representation of a continuous object, but for the experiment the mathematical data point information must be turned into a physical spatial variation in transmission. Laser plotters assign a finite physical dimension to each input data point they receive and are thus an obvious choice for the realisation of the targets. The procedure for object fabrication is to scale up the object to a suitable size for display on a laser plotter, using black to white inversion before plotting so that a black symbol on a white background results. This is required in the next stage - photoreduction - carried out by the Dept. of Physics photographer Peter Tuffy, so that the photographic negative of the reduced symbol has a transmitting symbol on a black background. High contrast `lithographic' film was used to minimise the effects of density non-uniformities in the laser printed output.

Accurate photoreduction, however, has three limitations. Firstly, there is always a loss in resolution involved in such a process so that as large an initial object as possible should be used to compensate for this. Secondly, photoreduction was available by integer steps only so that the initial object used must be as close as possible to an integer multiple, `K', in width of 1.898mm[4]. Thirdly, there is a maximum size of object that can be conveniently photoreduced. Laser plotting also has a limitation in that the spatial width assigned to a single data pixel is fixed. Therefore, the fabrication problem can be reduced to two stages:

1. **The integer value `K' of the photoreduction process is varied and a list of acceptable input dimensions suitable for photoreduction, D_K, obtained. On the plotter used, each data point is represented by a square of side [1/76]". In metric units, mm, the actual width of the scaled data block 16m pixels on a side is then**

$$D_m = 16m \times 0.334 \qquad (9)$$

   As the physical size attributed to each data pixel is so small (≅ 0.33mm) one would wish a relatively large value of m so that each data pixel is mapped to a dimension large enough that the resolution loss upon photoreduction does not significantly degrade the target object appearance. Without resort to calculation it was thought that m ≥ 3 would be a wise precaution.

2. **For each value of K, the data scaling variable `m' which gives Laser plotter output of dimension D_m closest to D_m is found.**

$$D_K = 1.898 \times K \qquad (10)$$

Table 7.2 lists several values of D_K and best agreeing m value. Many more combinations were tried than those listed here though the data serves to illustrate the nature of the procedure.

K	D_K	Best m	D_m	Final Object Size	% Error
10	18.984mm	3	16.042mm	1.604mm	15.5%
11	20.8824mm	4	21.389mm	1.944mm	2.42%

12	22.7808mm	4	21.389mm	1.782mm	6.12%
14	26.577mm	5	26.734mm	1.910mm	0.61%

Table 2: Possible Scaling Parameter Combinations.

Assuming perfect photoreduction, the least error incurred in fabricating an object which maps 16 data pixels to the ideal object width of 1.898mm occurs when the data is scaled up by a factor of 5 for laser plotting, then photoreduced by a factor of 14. These factors were used in the production of all seven target objects. However, an additional small scaling error was introduced in the fabrication stage so that the final target objects were found to be in the region of 1% larger than their ideal extent instead of 0.6% as expected from the above analysis. It is thought probable that this resulted during the photoreduction stage. Considering the relatively large mirror size of the 16×16 SLM, the slightly smaller scale spectrum of the target objects was not envisaged as a source of concern.

2.2 Alternative Target Fabrication.

Photographic transparencies suffer from inherent phase noise which cannot be eliminated. As mentioned in chapter one, a relief image proportional to photographic density commonly occurs on transparencies [15]. The use of binary amplitude transmittance objects minimises this effect as much as possible, though it will undoubtably have some (hopefully slight) detrimental effect on the system performance as a phase-only correlator. Alternative target fabrication techniques which do not suffer from such noise were examined. One such technique was to expose a layer of photoresist over a copper film (on a glass substrate) through the photographic transparency. Once etched with acid, the exposed regions were removed. The copper was sputtered onto the glass which allowed such features as the centre of the letter `A' to remain in place after etching, which in copper sheet would just drop out. If successful, the glass would be replaced by optically flat glass to greatly reduce phase noise in the target.

Unfortunately the resolution obtained by this method was very much worse than that of the photographic transparency, and the technique was abandoned. A further technique, known as `spark etching', physically cuts out a shape from a thin sheet of metal with a moving probe held at an extremely high potential relative to the sheet. The resolution obtainable with this device was at least an order of magnitude too low for the targets required, and again the centres of letters, etc. would drop out, requiring that it too be eliminated. In summary,

1. **Seven target objects, represented within a 16×16 array in the computer object space, were realised as photographic transparencies for input to an optical correlator system.**

2. **The target objects suffer from both a slight scaling error and the inherent phase noise associated with transparencies.**

3. **Experimental results should be taken as baseline due to the very small space-bandwidth product of the SLM used, the inherent phase noise of the targets used and the scaling error identified in this section.**

3 Information Extraction

In computer simulation of correlation algorithms it is neccessary to compare the quality of correlation peak with results obtained by other researchers. This final section deals mainly with the parameters chosen to characterise the correlation peaks of this project.

3.1 Characterisation Parameters.

In chapter 4, two commonly used parameters which characterise the performance of correlations resulting from a specified filter were introduced, Horner efficiency η_H and the signal to noise ratio SNR. Both these parameters have been used extensively by researchers to compare the performance of particular filter computation algorithms as determined by computer simulation. Optical efficiency provides a measure of how much energy is actually diverted into the region of the correlation and is therefore a parameter of importance to this project.

Frequently computer simulations are such that the phase binarisation algorithm is applied to each spectral data pixel of the frequency plane [33], [46], [52]. As such, a single correlation peak results in the image plane and it is a straightforward task to determine the peak height, average noise level of the background and so forth.

Optical Efficiency

As a reminder, the precise specification of what the Horner optical efficiency η_H measure is

1. **HORNER efficiency η_H is defined as the ratio of the energy within the 50% of peak threshold level to the total energy falling on the output plane of the correlator.**

Whilst this parameter serves well to characterise a single correlation peak in the image plane, the use of pixellated spatial light modulators causes replication of the correlations and thus limits the usefulness of this definition for simulations which mimic physical spatial filters. In fact, Horner efficiency cannot be used as a parameter as this measure assumes only one correlation peak. One might think that a more lenient interpretation of η_H would measure energy in the zero order replication relative to all the other replications. However, for practical, reflective pixellated filters the mirror shape determines the spectral envelope function and thus the relative distribution of energy in the image plane. It is for this reason that an experimental definition of optical efficiency is required.

Awwal et al [52] discuss several figures of merit for an optical correlator and propose a slightly different measure of optical efficiency than the popular definition of Horner.

1. **Optical efficiency η_A, as proposed by Awwal et al, is defined as the ratio of the peak correlation energy to the total energy of the beam leaving the frequency space.**

The clause about the beam energy as opposed to energy in the image plane represents a move towards a more experimental definition of the problem, where it is recognised that physical systems always suffer from a degree of attenuation. It is speculated that an underlying assumption of this definition is that most correlation peaks are sharp enough to be represented by a sharply peaked Gaussian function. In such a case, the peak energy would represent the energy present in the entire peak region quite accurately. Although this reduces the amount of calculation required to characterise a peak, the same clause applies to this definition as to that of Horner, in that it cannot in general handle replications.

A New Definition

Ideally one would desire to use the same definition of optical efficiency for both simulation and experiment. If the correlations are sharp enough, each replicated correlation peak should be surrounded by an area which is very much devoid of any light. Therefore it is proposed that

1. **An easily measurable optical efficiency, η_P, is defined as the ratio of the peak energy to the total energy contained within the spatial extent that the target object would cover in one replication present in the image plane.**

Notice that this definition becomes indistinguishable from that of Awwal et al in the limit of zero attenuation and zero filter pixellation. Although there will be limitations to the validity of this definition, it is nonetheless highly workable and will lead to meaningful comparison between filter algorithms in the next chapter.

Signal to Noise Ratio

The reflective mode SLM used in this project has large areas of highly reflecting circuitry in between the active light modulating regions. Also, the glass cover cube causes a very strong `front face' reflection which manifests itself as a bright, slightly out of focus, inverted image of the object in the image plane. There are thus three

primary sources of noise in the image:

1. **Light from the front face reflection.**
2. **Light scattered from the circuitry.**
3. **Filter dependent scatter of light out of each replicated peak centre.**

It is the last term which is commonly measured as the noise term, and is a complementary measure to the optical efficiency, and which Horner refers to in his definition of SNR as

1. **The SNR is defined as the ratio of the peak correlation value to the rms noise outside of a 50% of peak threshold level.**

However, the noise arising from the physical structure of the filter cannot be removed and thus makes measurement of SNR impractical. Although SNR *could* be measured in the simulations alone, this detracts from the primary purpose of the project which is to find an efficient filter algorithm with which to test the correlation abilities of the 16×16 SLM. Therefore it was decided not to attempt an estimate of this parameter in favour of more worthy and measurable ones.

Peak Sharpness

As mentioned above, the Awwal definition of optical efficiency - and the definition proposed here, η_P, are most suited to characterisation of a set of sharply peaked correlation functions. Peak sharpness ranges from being slight (most classical matched auto-correlations) to very sharp indeed (phase-only matched filters) and would thus seem an important characterisation parameter. In this project, the standard deviation σ of the best fitting 2-D Gaussian function is used as a measure of peak sharpness. For the computer simulations, the Gaussian was centred upon the peak maximum and used data up to 8 pixels away in both x and y directions in its calculation. The radial distance `r' of each point (i,j) was computed and the value of σ which minimised the error function

$$E(\sigma) = \sum_i \sum_j \left(e^{-[1/2]([r(i,j)/(\sigma)])^2} - \frac{I(i,j)}{I(0,0)} \right)^2 \qquad (11)$$

used to determine the best value of σ, where I(x,y) denotes the squared amplitude of the image plane pixel value.

The experimental data was captured using a CCD array camera. As with many such cameras, the pixels are slightly rectangular which means that the CCD pixel spacing is different along the x and y axes of the device. Consequently, captured images appear contracted in the x-direction when displayed as a square array and in order to present the Gaussian fitting routine with correct positional information one must multiply all x-coordinates by a factor of 1.42. Due to the larger scale of the image obtained from the CCD camera, as set against the 256×256 array upon which the data is written, the Gaussian fitting procedure used fewer data points than did the simulations. Further, inspection of the actual correlation peaks showed a very good approximation to a Gaussian over the central region of the peak, as figure 7.7 illustrates, but a poor agreement towards the outskirts of the correlation. In order that this region not degrade the fit obtained by more central data, a 5×5 data block was found to give the best fitting Gaussian for the peak as a whole.

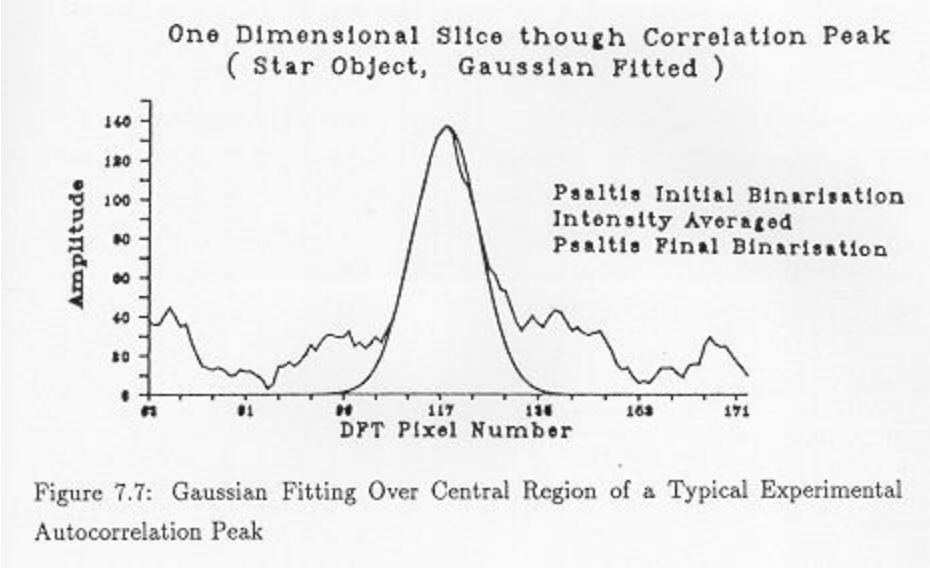

Figure 7.7: Gaussian Fitting Over Central Region of a Typical Experimental Autocorrelation Peak

Figure 7.7: Gaussian Fitting Over Central Region of a Typical Experimental Autocorrelation Peak

3.2 Chapter Review

The computational framework allowing precise simulation of the optical correlation system has been described. Preference was given to frequency space resolution over object space resolution because

1. **The small space-bandwidth product of the 16×16 SLM required a small object to obtain a reasonable frequency bandwidth.**

2. **To physically manufacture such a small object required it be rather course - of low spatial resolution - to begin with.**

3. **Filter calculation, it seems logical to assume, would be aided by having as much information on the spectral distribution over each mirror - high resolution - as is conveniently obtained by computational Fast Fourier Transformation.**

If, as is commonly speculated, the correlation efficiency of a filter depends on certain target object characteristics, such as degree of symmetry, filter evaluation should base itself around results obtained from a variety of different objects. In so doing, the general trends observed in correlations arising from use of a particular filter may be picked out. The manufacture of seven target objects has been described in detail and an error analysis performed to determine the effects of object width scaling error on SLM-expected spectrum registration. A small scaling error does exist in the objects used but is expected to be acceptable due to the large `capture' area of the mirrors on the 16×16 SLM used in the experiment.

Several parameters commonly used in the characterisation of correlation results have been evaluated with respect to their applicability to characterisation of real data. In particular, the optical efficiency of the correlator has been redefined to cover the use of pixellated spatial filters. The next, and final, chapter describes the filter computation algorithms used and compares the results of computer simulation with actual data obtained by using the 16×16 SLM as a binary phase-only optical correlator.

Footnotes:

[1] **A review of the basic optical processing system may be found in appendix one.**

[2] **The influence of scale of transform on correlation peak intensity is a subject requiring further study. There may be a general optimum relationship with percentage energy of the high spatial frequencies excluded. This work is being continued in the Applied Optics Group.**

[3]**This is a straightforward consequence of the fact that the computed spectrum is of form**

$$F(j) = \sum_{k=1}^{N} f(k)\, e^{-i\,2\pi[\,jk/N]} \quad (12)$$

An `optical' transform format is obtained by multiplying the k'th object data point by -1^k, eliminating the need for re-arrangement altogether, where k ranges from 1 to N^2.

[4]**The precise, three significant figure value was used here although it could be argued that 1.9mm is equally valid given the usual tolerences of lens focal lengths**

File translated from T_EX by T_TH, version 3.02.
On 27 Oct 2001, 23:41.