Optical Correlation using Pixellated Spatial Light Modulators

Steven B. Heddle

Ph.D.

University of Edinburgh

1993

Acknowledgements

There are many people who deserve acknowledgement for the help and support they gave during the course of this work. Some deserve special gratitude and I have attempted to redress the balance below. Undoubtedly I have omitted some who also deserve a special mention– rest assured that I will carry the awesome burden of responsibility to the grave.

I wish to thank my supervisors, David Vass and Dick Sillitto, for their encouragement and always well considered comments and advice. Although not officially my supervisor, Will Hossack has been particularly helpful, and has done much of the proof reading of this thesis– thanks Will. I have also received invaluable technical assistance from a number of sources, especially Andy Garrie and Eric Davidson in the Physics Department; Tom Stevenson, and Alan Gundlach in the Edinburgh Microfabrication Facility, and Paul Rodgers in Electrical Engineering, who helped with lithography; and Peter Tuffy with all round photographic wizardry and good conversation. Thanks too to Ian Chisholm, whose Fourier transforming program I have used a lot.

During my time here I have made some good friends, only to neglect them over the last year as I have toiled with this beast– cheers Gerry, Alec, Siobhan, Duncan, and Ian. Also a knowing wink to the Applied Optics veterans Doug, Mike and Ian for some excellent humour and great rants on numerous social evenings.

Finally, eternal gratitude to my greatest supporters, Mum, Dad and Donna– this thesis is for you as much as me.

Declaration

I declare that the composition of this thesis, and all the work described within it, was carried out by me, except where otherwise acknowledged.

August 1993

to Mum, Dad and Donna

Dean cruaidh d'obair
'S dean cruaidh do cluich
Agus faireach cac

Abstract

Optical correlation by filtering in the Fourier plane of an optical processor has recently received much attention, due to the advent of rapidly reconfigurable Fourier plane filters– spatial light modulators (SLMs). Many algorithms have been developed to generate filters, mostly binary in nature in line with the capabilities of current SLMs. This thesis reviews optical correlation techniques and SLM technology, before turning attention to the consequences of the pixellated structure of the Fourier plane filters generally necessary to enable arbitrary filter patterns to be written to an SLM. Conventionally the pixels are identical and regularly spaced on a rectangular grid. This is shown to lead to replication in the output plane of the correlator, and aliasing if the input exceeds dimensions related to the pixel spacing in the Fourier plane filter. If the input is also provided by an SLM, this requires the number of pixels across the filter SLM (its space-bandwidth product) to be greater than or equal to that of the input, even though the target itself may occupy only a small area of the total input scene, in order to prevent aliasing and misleading correlation results. If the replication and aliasing could be prevented, the space bandwidth product of the filter need only be matched to that of the target to be detected, rather than the entire input scene. The replicas arise through convolution with distinct spectral orders in the SLM's Fourier spectrum. A means of eliminating the spectral orders other than the zero order through randomisation of the pixel positions is presented and analysed, and implemented on transmissive matrix addressed SLMs custom designed and built for the work of this thesis. Optical correlation experiments carried out using the SLMs are presented and compared with computational simulations. These correlation experiments demonstrate the effectiveness of the pixel position randomisation in eliminating replication and aliasing in the correlator output, for both phase and amplitude modulating SLMs used as Fourier plane filters.

Contents

Acknowledgements i

Abstract iv

1 Introduction: 1
- 1.1 Optical processing . . . 2
  - 1.1.1 Parallelism and connectivity . . . 2
  - 1.1.2 Coding of information . . . 2
  - 1.1.3 Processing operations . . . 3
- 1.2 Structure of this thesis . . . 3

2 Review of optical correlation 6
- 2.1 An optical processor . . . 6
  - 2.1.1 Analysis of the optical processor . . . 7
- 2.2 Simple optical processing operations . . . 10
  - 2.2.1 Edge enhancement . . . 11
  - 2.2.2 Halftone removal . . . 13
  - 2.2.3 Feature extraction . . . 13
- 2.3 Fourier plane filters for correlation . . . 14
  - 2.3.1 Early work . . . 18
  - 2.3.2 Single modulation parameter filters . . . 19
  - 2.3.3 Binary phase only filters . . . 22
  - 2.3.4 Amplitude encoded binary phase only filters (AEBPOFs) . 26
  - 2.3.5 Fourier plane filters for scale and rotation invariant pattern recognition . . . 28

2.3.6 Joint transform correlators 32
2.4 Summary 35

3 Liquid crystals and spatial light modulators **37**
3.1 Spatial light modulation 37
3.2 Addressing 38
3.2.1 Optically addressed SLMs 38
3.2.2 Electrically addressed SLMs 41
3.2.3 Summary 46
3.3 Modulation using nematic liquid crystals 46
3.3.1 Basic concepts 47
3.3.2 Nematic liquid crystal 48
3.3.3 Amplitude modulation using the twisted nematic effect . . 50
3.4 Multiplexed drive of nematic liquid crystal SLMs 51
3.4.1 Multiplexibility 55
3.4.2 Choice of suitable liquid crystal for multiplexed drive . . . 56

4 Use of a regularly pixellated SLM **60**
4.1 Regularly pixellated amplitude filter 61
4.1.1 Analysis in 1-d 63
4.1.2 Generalisation to 2-d 72
4.1.3 Simulations for a regularly pixellated 2-d array 74
4.1.4 Discussion of results 78
4.2 Conclusions 83

5 Use of a randomly pixellated SLM **85**
5.1 Motivation to remove the higher spectral orders 86
5.2 Randomly Pixellated Amplitude Filter 87
5.2.1 Analysis in 1-d 87
5.2.2 Generalisation to 2-d 92
5.3 Specific examples of randomised arrays 94
5.3.1 Example 1: Uniformly random distribution 96
5.3.2 Example 2: Discrete position random distribution 106
5.3.3 Example 3: Special case of discrete position randomisation 114

5.3.4 Additional comments 122
5.4 Conclusions 125

6 Construction of a transmissive SLM **127**
6.1 Basic design of SLM 127
6.2 Choice of matrix addressed transmissive configuration 129
6.3 Construction of the device 131
6.3.1 Design and fabrication of the photomask 131
6.3.2 Substrate preparation 134
6.3.3 Lithography 136
6.3.4 Coatings 136
6.3.5 Assembly 140
6.3.6 Filling and sealing of device 141
6.4 Interfacing 142
6.4.1 Interface for prototype devices 143
6.4.2 Interface for devices used in an optical processor 145
6.5 Characterisation of the SLMs 147
6.5.1 Interferometric examination of the SLMs 147
6.5.2 Contrast measurement for the amplitude modulator 153
6.5.3 Physical dimensions of the SLMs 155
6.6 Summary 155

7 Correlation experiments **156**
7.1 Simulations 158
7.1.1 Simulation of AEBPOFs 162
7.1.2 Results of AEBPOF simulations 163
7.1.3 Effect of incorrectly centering the reference target 169
7.1.4 Simulation of BPOFs 170
7.1.5 Results of BPOF simulations 171
7.2 Optical implementation 177
7.2.1 The optical correlator 177
7.2.2 Alignment of components 180
7.2.3 The filter patterns 181

7.2.4 Optical results for AEBPOFs 181
7.2.5 Optical results for BPOFs 186
7.3 Conclusions and summary 193

8 Discussion **195**
8.1 Summary of the work 195
8.1.1 Consequences of pixellation 195
8.1.2 Analysis of random pixel arrays 196
8.1.3 Fabrication of transmissive matrix addressed nematic liquid crystal SLMs 197
8.2 Future work 199
8.2.1 Improvements to the SLMs 199
8.2.2 Further work with random pixel SLMs 200
8.3 Summary 203

Bibliography **205**

A Linearity of array camera **215**

B Listing of pixel.f **217**

C Publications **226**

Chapter 1

Introduction:

Spatial light modulators and optical processing

A spatial light modulator (SLM) is a device which impresses information onto an optical wavefront. This can be as simple as a photographic transparency or an aperture stop, but conventionally refers to a range of devices developed comparatively recently which can selectively modulate over a 2-dimensional region, and are reconfigurable over a short period of time (< 1 second), by means of some optical or electrical addressing scheme. The latter definition will be applied throughout this thesis.

Spatial light modulators find an important application in the field of optical processing. SLMs and optical processing will be considered in more detail in later chapters, but some description of the basic principles and motivation behind optical processing is appropriate here. With this groundwork in place, the final section of this chapter will outline the structure of this thesis, and the work contained within.

1.1 Optical processing

Optical processing is concerned with the manipulation of information encoded as a disturbance on an optical wavefront. Some of the features of optical processing are described below.

1.1.1 Parallelism and connectivity

Considering a simple imaging system where an input plane is imaged to an output plane, the process can be seen as a transfer of information between the 2-dimensional input and output planes. As each resolvable point mapped from input to output in the operation can be used to represent an information channel, clearly the transfer of information is potentially massively parallel. Importantly this parallel transfer of information can be effected through free space or any uniform transparent medium by a simple optical component (a lens) acting on all channels simultaneously. In comparison, the electrical analogue requires that each channel is assigned a wire, isolated from its neighbours– an arrangement that becomes cumbersome for even a modest degree of parallelism.

1.1.2 Coding of information

Light offers a number of convenient methods by which information can be coded. In the case of a photographic transparency illuminated by white light the information is coded in terms of the colour (wavelength) and/or intensity of the light transmitted by the transparency. When the illumination is provided by a laser (and is thus coherent and monochromatic) it must be characterised by its complex amplitude, allowing information to be coded in terms of the amplitude and/or phase of the illumination. This pertains to the field of *coherent optical processing* which encompasses the work of the thesis.

A further means of encoding information in terms of polarisation is possible, where the information is represented by the polarisation state of the light. Orthogonal polarisation states may conveniently be used in digital optical processing applications such as that of Weigelt [1]. An advantage of 'polarisation logic' is that the complement of a logic state can be formed with theoretically no loss of intensity, allowing limitless cascadibility. This is an idealisation, as in practice the polarisers, analysers and modulating devices necessary are lossy.

1.1.3 Processing operations

Many operations come under the above definition of optical processing, from 2-dimensional scale transforms, and shifts and rotations of axes routinely achieved by a lens, through the whole field of optical computing [2, 3] to adaptive optics for astronomical imaging [4]. Of particular interest in this thesis is the field of coherent optical processing which exploits the ability of a coherently illuminated system to form the Fourier transform of a 2-dimensional input plane. With access to the Fourier transform, image modification and pattern recognition operations [5, 6] may be conveniently carried out by selectively modulating its amplitude and/or phase, using some kind of spatial light modulator.

1.2 Structure of this thesis

The work of the thesis is concerned primarily with the particular Fourier plane filtering operation of **optical correlation**, and the implications of the use of a pixellated SLM as the Fourier plane filter.

Optical correlation provides a means of establishing the presence or absence of a chosen target object or objects in the 2-dimensional input to an optical processor. This is introduced in **chapter 2**, once the optical processing architecture and coordinate system used throughout has been defined. Simple image processing techniques are presented to develop the concepts of filtering spatial frequencies

and orientational features underlying optical correlation. Optical correlation using matched filters and the joint transform correlator architecture is introduced, and the evolutionary development of current matched filter based techniques described and justified.

Chapter 3 gives a brief overview of SLM technology, drawing the distinction between optically and electrically addressed SLMs, and showing why pixellated SLMs are necessary when correlating using matched filters. Particular emphasis is given to SLMs where the modulation of the light is achieved using liquid crystal. The means of achieving various forms of light modulation for nematic liquid crystal is outlined, and the electro-optic characteristics of a particular nematic liquid crystal studied for suitability when used with a multiplexed drive scheme. These results are used to fabricate the SLMs described in chapter 6, used in the correlation experiments of chapter 7.

When using a conventional regularly pixellated structure for filtering in the Fourier plane it is necessary to optimise the size of the Fourier transform to achieve the highest possible sampling frequency without incurring aliasing in the retransformed image, due to the replicas formed as a result of the regularity of the sampling. This is demonstrated in **chapter 4**, and the form of the Fourier transform of the regularly pixellated filter is analysed in detail and related to the structure of the filter. Computational and optical simulations are presented which show the effects of aliasing. A parameter of pixellated SLMs, the **space bandwidth product** or SBP, is introduced, which is equal to the number of pixels across the SLM. It is shown that as a consequence of the sampling and replication, the SBP of the input must be less than or equal to that of the filter to prevent aliasing.

Chapter 5 proposes the idea of using irregular pixel structures to suppress the higher order peaks in the Fourier spectrum of the filtering device which lead to replication and aliasing. The random pixel idea is developed from this intuitive basis into an analytic expression in 1 dimension, derived from the statistics of the random process. This is generalised to the more useful 2-dimensional case, and special cases considered. These special cases refer to the choice of probability distribution function governing the offset of the pixels from their regularly spaced

positions, and it is shown how it is possible to randomise the pixel positions such that all the higher orders can be removed, at the expense of redistributing their energy into a diffuse low level background 'noise'. The special cases are extensively simulated for 16×16 arrays, both computationally and using transparencies, and ultimately used to make up actual transmissive SLMs.

The design and construction of these devices and their electrical interfaces is described in **chapter 6**. The performances of the devices are characterised by interferometric and intensity measurements.

Correlation experiments using phase-modulating and amplitude-modulating SLMs with regular and randomised pixel arrays are carried out in **chapter 7**. The optical experiments are also simulated computationally, and compared with the simulations and theoretically expected results. The correlation experiments demonstrate the effectiveness of pixel position randomisation in removing aliases from the output– this is an important result with numerous practical benefits, as removal of the aliases means that the space bandwidth product of the filter SLM need now only be matched to that of the target to be detected rather than that of the entire input scene, as is the case with a regularly pixellated SLM. This is demonstrated using an input containing several targets whose correlation outputs would overlap if replicated.

The work of the thesis and the consequent results are summarised in the final chapter. Future work resulting from this project is suggested.

Chapter 2

Review of optical correlation

This chapter introduces the field of optical correlation as a subset of a wider range of image processing operations which can be carried out using coherent illumination on a standard Fourier optical processing bench.

This standard architecture is described and related to a system of coordinates which is used throughout the remainder of the thesis. Simple image processing operations are described, before concentrating on optical correlation techniques, which are related to the use of two-dimensional image processing devices– spatial light modulators.

2.1 An optical processor

The optical processing architecture used is shown in figure 2.1, and is well known [7, 6, 8]. This type of optical processor is generally referred to as a *4-f processor*, from the separations of the input and output planes.

For convenience, two rays are included in the diagram to show the geometrical imaging properties of this system. This geometrical viewpoint is an approximation which neglects the occurence of diffraction effects due to the wave properties of the

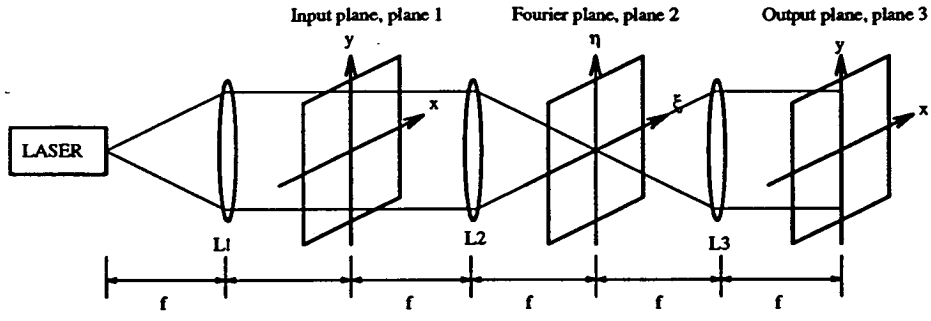

Figure 2.1: Fourier optical processor

illumination, which form the basis of the optical processing operations described later. In particular, the property whereby a lens can form a diffraction pattern of a coherently illuminated input object which corresponds to that object's Fourier transform is utilised [9].

In the architecture described, the Fourier transform of the input is formed in the back focal plane of lens L2. This plane is the conjugate plane of the illuminating point source and is referred to as the *Fourier plane.* The optical processing operations considered here are performed by placing a two dimensional filter in the Fourier plane to selectively modulate the amplitude and/or phase of regions of the input's Fourier transform.

2.1.1 Analysis of the optical processor

It is necessary to analyse the operation of the optical processor mathematically, with reference to the system of coordinates shown in figure 2.1.

Monochromatic illumination is provided by a laser, and is seen to come from a point source by means of a spatial filter assembly mounted in front of the laser. Lens L1 is mounted with the point source at its front focal point to provide collimated uniform illumination at the object plane, plane 1. The object will conventionally be in the form of a transparency or spatial light modulator, and can be represented by the two-dimensional transmission function $t_1(x, y)$ which may be complex valued. Assuming uniform illumination of unit amplitude, the scalar

light field immediately to the right of plane 1 can be represented by $t_1(x,y)e^{i\omega t}$. The $e^{i\omega t}$ is the time dependent part of the illumination, related to the wavelength by $\omega = 2\pi c/\lambda$.

At this point it is appropriate to introduce a system of notation (after Gaskill [7]) whereby the complex amplitude in plane $z = z_i$ is denoted by $u(x,y;z_i) = u_i(x,y)$. Further, the light field impinging on this plane from the left is given by $u_i^-(x,y) = u(x,y;z_i - \epsilon)$, and the light field leaving it is given by $u_i^+(x,y) = u(x,y;z_i + \epsilon)$ where ϵ is arbitrarily small, approximating that the object of interest has zero thickness. In practice the objects do have a finite thickness, but any variations in the optical path length through the object can be incorporated as a phase term in its complex transmittance. Hence in the case described so far $u_1^-(x,y) = e^{i\omega t}$, and $u_1^+(x,y) = u_1^-(x,y)t_1(x,y) = t_1(x,y)e^{i\omega t}$. This is simplified further if we drop the $e^{i\omega t}$ term, as this time dependence plays no part in the analysis.

Continuing, lens L2 forms the Fourier transform of $t_1(x,y)$ in its back focal plane, plane 2 at $z = z_2$, which is the conjugate plane to the illuminating point source [7, 8],

$$u_2^-(\xi,\eta) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} u_1^+(x,y)e^{-i2\pi(\xi x+\eta y)}dxdy. \tag{2.1}$$

The (ξ,η) are scaled spatial frequency coordinates which are used to simplify the notation. The scaling factor is such that

$$\xi = \frac{u}{\lambda f} \text{ and } \eta = \frac{v}{\lambda f} \tag{2.2}$$

where (u,v) are the spatial coordinates in plane 2 [7]. Thus $u_2^-(\xi,\eta) = \mathcal{F}[t_1(x,y)] = T_1(\xi,\eta)$ where $\mathcal{F}$ denotes the Fourier transform operator and $T_1(\xi,\eta)$ is the Fourier transform of $t_1(x,y)$. This in turn impinges on the transparency or spatial light modulator in plane 2, whose transmittance is given by $T_2(\xi,\eta)$ for consistency of notation. Thus $u_2^+(\xi,\eta) = T_1(\xi,\eta)T_2(\xi,\eta)$. Finally lens L3 forms the Fourier transform of $u_2^+(\xi,\eta)$ in its back focal plane to give

$$\begin{aligned} u_3^-(x,y) &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} u_2^+(\xi,\eta)e^{-i2\pi(\xi x+\eta y)}d\xi d\eta \\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} T_1(\xi,\eta)T_2(\xi,\eta)e^{-i2\pi(\xi x+\eta y)}d\xi d\eta \\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} T_1(\xi,\eta)e^{-i2\pi(\xi x+\eta y)}d\xi d\eta \\ &\quad * \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} T_2(\xi,\eta)e^{-i2\pi(\xi x+\eta y)}d\xi d\eta \\ &= t_1(-x,-y) * t_2(x,y) \end{aligned} \tag{2.3}$$

applying the convolution theorem, with $*$ denoting the 2-d convolution operator.

Thus we obtain a coordinate reversed image of the object convolved with the Fourier transform of the filter function.

There are several points to note here:

1. The reversal of the coordinates for $t_1(x,y)$ is a consequence of L2 and L3 performing successive forward Fourier transforms, rather than a forward FT followed by an inverse FT.

2. If $T_2(\xi,\eta) = 1$, i.e. no modulation of light takes place at the filter plane, then $u_3^- = t_1(-x,-y) * \delta(0,0) = t_1(-x,-y)$ (where $\delta(0,0)$ is a Dirac delta-function at the origin) and the result is as expected from geometrical optics. The sizes of the lenses are assumed to be large compared to the input object, and the extent of the point spread function of the lens can be ignored.

3. Given 2. , it is apparent that the system has unit magnification. This is as a consequence of the scaling between spatial frequency and physical distances for lenses 2 and 3 being taken to be the same, i.e. there is an assumption that L2 and L3 have the same focal length. This need not be the case, and magnification may be achieved by choosing lenses with differing focal lengths.

4. For the case described above, the Fourier transforms performed lie on planes 2 and 3 as described. This is as a consequence of the 2-d scenes being transformed lying in the front focal planes of the respective transforming lenses. If this is not the case, a multiplying quadratic phase factor will be included in the Fourier transform, indicating that the exact Fourier transform lies on

the cap of a sphere. In many cases the radius of the sphere will be sufficiently large for it to be approximated well by a plane, near the optical axis [9].

5. The retransformed image is convolved with the transform of the filter function. Therefore the choice of filter function will critically affect the form and quality of the retransformed image.

With the mathematical principles of the operation of the optical processor established, some simple processing operations are described prior to consideration of optical correlation techniques.

2.2 Simple optical processing operations

Using the architecture described above, the amplitude and/or phase of the Fourier transform of the input object may be selectively modulated by use of a suitable Fourier plane filter. In addition, the polarisation state of the light may be modulated, an operation which finds application in the field of digital optical logic [1], but which is not considered further here.

This section considers a range of coherent optical processing operations which use simple aperture stops in the Fourier plane to block out certain components of the input object's spatial frequency spectrum and thus modify the resulting output image. These operations date from the period at which lasers were becoming readily available as a source of coherent illumination. The effect of simple aperture stops can of course be duplicated by a photographic transparency consisting of opaque or uniformly transmitting regions, or by use of a binary spatial light modulator modulating only amplitude, and with high contrast between the two binary states. Such a spatial light modulator is used in the optical demonstrations which follow in this section.

The optical processing operations described are edge enhancement of binary input

objects, removal of halftone dot structure from input scenes, and selective removal of features which contain a distinct directional structure from an input. These and other operations are described more fully in a review article by Birch [10] and in Fourier Optics texts such as those of Steward, Reynolds et al., and Stark [6, 5, 11]. The overall approach of using a Fourier optical processor to process 2-d information, stated in analogous language to that of linear electrical systems, dates back to the work of Cutrona et al. [12]

2.2.1 Edge enhancement

This operation serves to enhance the visibility of object details such as lines and edges, while removing areas of uniform transmittance in the object from the image. Edge enhancement is carried out by blocking off the lowest spatial frequency components of the Fourier transform of the input object, which are characteristic of the larger features in the input, as might be inferred from the reciprocal relationship of the variables in Fourier transform pairs.

In particular, the zero spatial frequency component of the object's transform (corresponding to any *d.c. level* present in the object, to borrow the terminology of electrical systems) is removed. This is the dominant consequence of the stop, as removing the zero spatial frequency component from the object spectrum results in an image with an amplitude distribution with zero mean. Thus the modified image is made up of positive and negative amplitudes, compared to the object, which contained only amplitudes ≥ 0. The presence of the negative amplitudes becomes significant when the object is observed, as optical detectors (the eye included) respond to light intensity, the squared modulus of the light amplitude. Thus in the modified image both positive and negative amplitudes will give rise to positive intensities, with a zero of intensity where the amplitude is zero - this occurs at points where the amplitude in the original object falls from above to below the d.c. level and vice-versa, and is a characteristic of edges. Hence, in the case of a binary object, the modified image should have reduced contrast between the light and dark regions when compared to the object, and with dark lines at

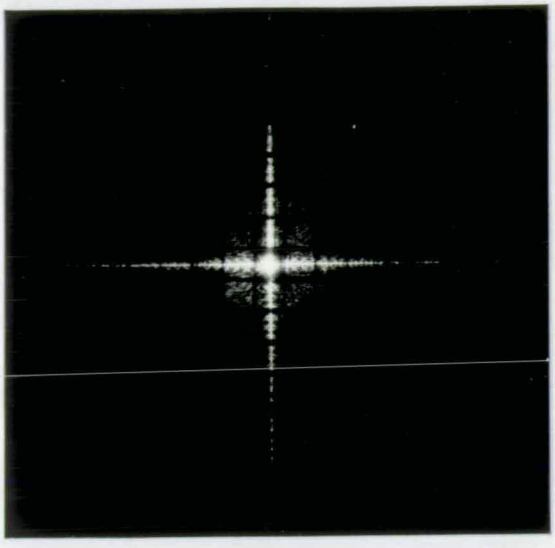

a) Fourier plane, with d.c. blocked by SLM

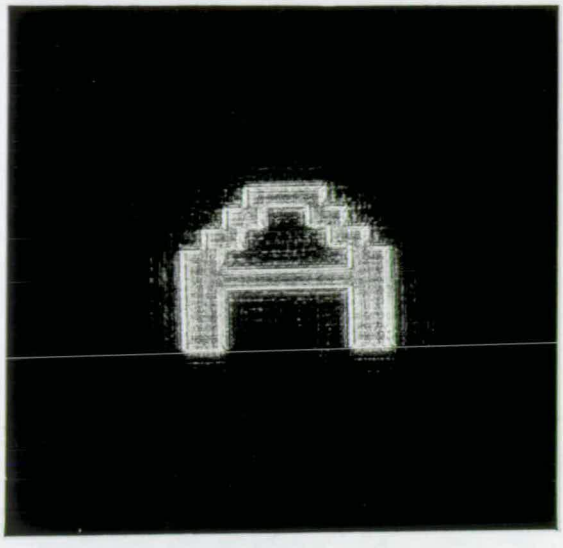

b) Output showing edge enhancement of letter A

Figure 2.2: Optical demonstration of edge enhancement

edges.

The actual situation is slightly more complicated than this, as the filter stops low frequencies other than the zero spatial frequency, and exhibits a high frequency cut off at its extremities. As noted, the transform of the filter is convolved with the input to give the image, which can be used to explain the symmetrical distribution of light and dark bands around the dark lines denoting edges [10, page 68]. This was demonstrated optically using the optical processor configuration described, with a spatial light modulator as the filter, and with a letter A defined on a high contrast transparency as input. The results are shown in figure 2.2. The spatial light modulator used was an early prototype of the amplitude modulating device described in chapter 6, so a significant fraction of the light is still transmitted through the single 'off' pixel in the centre. Nevertheless, edge enhancement is clearly demonstrated.

2.2.2 Halftone removal

Whereas edge enhancement is basically effected by high-pass filtering of the object's spatial frequency spectrum, halftone removal is a low-pass filtering operation.

Removal of the halftone dot structure from, for example, newspaper pictures is possible because the dot structure in the picture has a higher spatial frequency than the features represented in the resulting halftoned picture. A transparency of a halftoned picture used as input to an optical processor will give rise to a spatial frequency spectrum in the Fourier plane of a series of regularly spaced bright spots arising from the regular dot structure of the halftone pattern, with the spatial frequency spectrum of the underlying image centred on the central bright spot at zero spatial frequency[1], at lower spatial frequencies than the neighbouring bright spots. By choosing a stop, such as an adjustable iris, to pass only the central spot and the low frequencies up to half the distance to the neighbouring spots, the high spatial frequency information characteristic of the halftone process can be removed to give an image free from halftone dots.

The bright spots are an example of *spectral orders*, arising from sampling of the picture at regular intervals by the halftone process. This effect is considered in detail in chapter 4.

2.2.3 Feature extraction

The previous two examples have described optical processing operations where the parameter which decides whether or not a spatial frequency component is allowed to pass or not is its distance from the centre of coordinates in the Fourier plane. In other words, if the (ξ, η) coordinate system of the Fourier plane was replaced

[1]Due to convolution, the spatial frequency spectrum of the underlying image will also be replicated at each of the bright spots

by a standard (r, θ) polar coordinate system , the important parameter thus far would be r.

An application where θ becomes the important parameter is where an input object contains distinct orientational features which are required to be removed, or emphasised, such as in the 'house' image shown in figure 2.3a).

As can be seen, the various parts of the house such as the roof, door, windows etc. are characterised by sets of parallel lines at differing orientations. Each of these sets of lines behaves like a diffraction grating, giving rise to a spectrum with a number of spectral orders on a line perpendicular to the orientation of the lines in the object, the various spectra combining to give the overall spectrum of the house shown in figure 2.3b). This is a photograph of the Fourier plane of the optical processor with the house as input, and the prototype SLM as the Fourier plane filter when all pixels are 'on'.

By blocking all the spectral orders except for those lying on a line through the origin at angle θ, any features of the house made up of lines orientated at $(\theta+90)^\circ$ will be emphasised as the other features are eliminated. Similarly the features of the house orientated at $(\theta+90)^\circ$ can be eliminated by selectively blocking out the orders in the spectrum lying along an line at θ°, as is shown in figures 2.3c) and d) for $\theta = 0^0$. Note that the stop does not have a gap in it to allow the zero order to pass, resulting in some edge enhancement of the remaining features.

2.3 Fourier plane filters for correlation

The previous section has shown how very simple filter patterns can be employed to carry out image processing operations optically, which would be computationally intensive if carried out digitally. With more complicated filter patterns, which may spatially modulate phase and/or amplitude over the filter, it is possible to carry out more sophisticated operations such as pattern recognition. The approach used here is to use the Fourier transformation properties of the optical processor

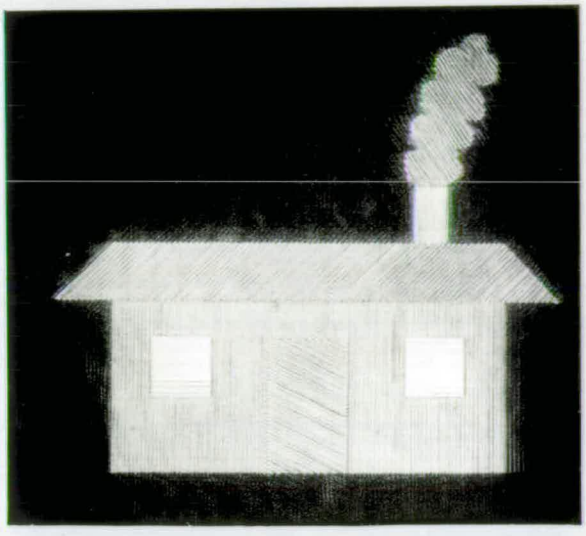

a) Input 'house' image

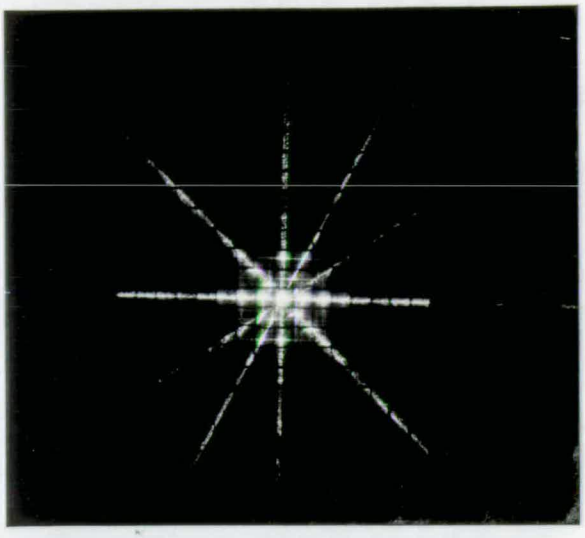

b) Fourier spectrum of house

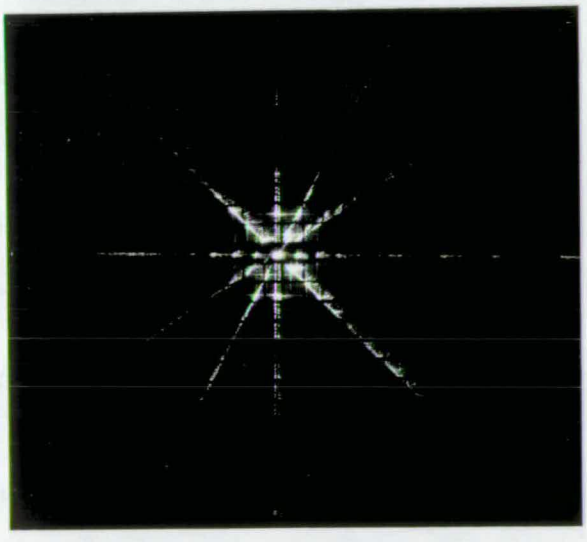

c) Filtering of spectrum by SLM

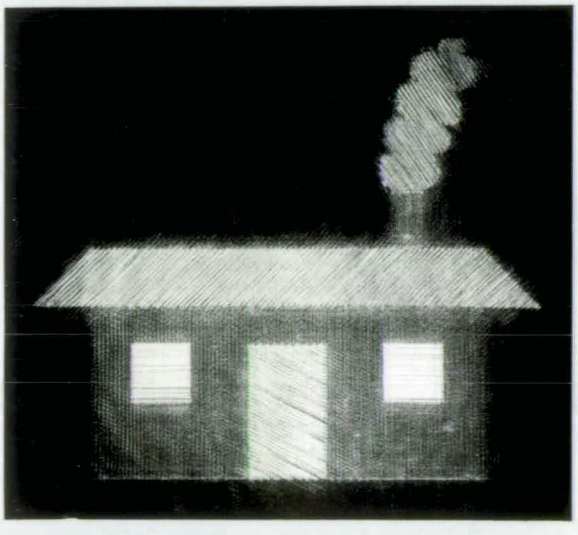

d) Resulting output of optical processor

Figure 2.3: Extraction of orientational features from an object

to achieve an output which is related to the correlation between a chosen 2-d function (the *target*) and the input, by use of a filter specific to the target.

In cases where the input contains the target an autocorrelation signal can be obtained. This idea is introduced as follows: consider a target at Plane 1, with transmittance $g(x,y)$ which Fourier transforms to $G(\xi,\eta) = |G(\xi,\eta)|e^{i\phi(\xi,\eta)}$, where ϕ is the phase of the Fourier transform. If a filter is made such that $G^*(\xi,\eta) = |G(\xi,\eta)|e^{-i\phi(\xi,\eta)}$ and placed in the Fourier plane then, as stated earlier, the output of the processor will be the convolution of the input with the Fourier transform of the filter, subject to a coordinate reversal i.e.

$$\begin{aligned} u_3(x,y) &= \mathcal{F}[G(\xi,\eta)G^*(\xi,\eta)] \\ &= g(-x,-y) * g^*(x,y) \end{aligned} \tag{2.4}$$

using the convolution theorem, and recalling that the transform of a conjugate involves an additional change of sign in the transform variable [7, page 195]. This expression for $u_3(x,y)$ can be rewritten as $u_3(x,y) = g(-x,-y) \star g^*(-x,-y)$, where $\star$ denotes the correlation operator, and where the complex correlation of $f(x,y)$ with $g(x,y)$ is defined as

$$f(x,y) \star g^*(x,y) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(\alpha,\beta)g(\alpha - x, \beta - y)d\alpha d\beta \tag{2.5}$$

[7, page 172]. Both forms are equivalent and equal to the complex autocorrelation of $g(-x,-y)$ denoted by $\gamma_g(-x,-y)$, which has the important property that its modulus is a maximum at the origin. Similarly, if some non target image with transmittance $f(x,y)$ is placed at Plane 1 and filtered by the same filter the output of the processor will be

$$\gamma_{fg}(-x,-y) = f(-x,-y) \star g^*(-x,-y) \tag{2.6}$$

where γ_{fg} denotes the *cross correlation* of functions f and g. Thus, when we view the output, the presence of the target is marked by a bright peak at the location of the target in the input field (subject to a reversal of coordinates, due to the two successive forward Fourier transforms), characteristic of the autocorrelation. A more physical interpretation of the autocorrelation is that the filter $G^*(\xi,\eta) = |G(x,y)|e^{-i\phi(\xi,\eta)}$ serves to cancel out the differences in phase in the wavefront impinging on Plane 2, resulting in a plane wave leaving Plane 2 which will focus down to a spot one focal length after lens L3– the autocorrelation peak in the image plane, Plane 3.

The technique can be generalised to larger input scenes containing several sets of objects, one of which is to be identified, e.g.

$$\begin{aligned} t_1(x,y) = u_1^+(x,y) &= f_1(x+a_1, y+b_1) + f_2(x+a_2, y+b_2) \\ &\quad + g(x+c_1, y+d_1) + g(x+c_2, y+d_2) \end{aligned} \tag{2.7}$$

where the $g(...)$ are the target objects to be detected using the filter described already, and $f_1(..)$ and $f_2(..)$ are non target objects. In this case the output of the optical processor will be

$$\begin{aligned} u_3^-(x,y) &= f_1(-(x+a_1), -(y+b_1)) \star g^*(-x,-y) \\ &\quad + f_2(-(x+a_2), -(y+b_2)) \star g^*(-x,-y) \\ &\quad + g(-(x+c_1), -(y+d_1)) \star g^*(-x,-y) \\ &\quad + g(-(x+c_2), -(y+d_2)) \star g^*(-x,-y) \\ &= \gamma_{f_1 g}(-(x-a_1), -(y-b_1)) + \gamma_{f_2 g}(-(x+a_2), -(y+b_2)) \\ &\quad + \gamma_{gg}(-(x+c_1), -(y+d_1)) + \gamma_{gg}(-(x+c_2), -(y+d_2)), \end{aligned} \tag{2.8}$$

using the *shift invariant* properties of the Fourier transform, which states that if either function is shifted by an amount (x_0, y_0) the resulting convolution is

shifted by the same amount but unchanged in magnitude or form. Thus as well as indicating the presence or absence of a target in the input, this optical correlator can be used to determine the position of a target or targets in the input, relative to that of the reference object used to set up the filter.

The example output described above is the superposition of four correlation functions - two autocorrelations, two cross correlations. Depending on the nature of the input functions, the cross-correlations may have a higher value than the autocorrelation peak (if the non-target objects pass more light than the target, for example), and the autocorrelation peak itself need not be particularly sharp or distinctive, reducing the ability of the correlator to detect the targets. These problems can be alleviated by edge-enhancing the input and reference, to reduce cross-correlation with uniformly transmitting areas conveying little or no shape information.

Further problems arise when the input target is scaled or rotated with respect to the reference– again, there are ways of compensating for the lack of scale and rotation invariance, which will be mentioned later.

Correlation has been introduced here in terms of an ideal complex filter, which in practice is non-trivial to obtain, particularly when implemented on a spatial light modulator. The SLM is particularly suited as the Fourier plane filter in an optical correlator due to its ability to be rapidly updated to display another filter, but generally is binary in operation and modulates amplitude or phase. Much work has been carried out in developing this basic Fourier plane filter concept to enable optical correlation to be carried out effectively, both in developing filters and algorithms to determine filter patterns, and this is outlined here.

2.3.1 Early work

The complex filter described above, the complex conjugate of the amplitude spectrum of the target to be detected, is an example of the matched filter postulated by Cutrona et al. in 1960 [12]. In their paper they describe how the phase and

amplitude variations required of the filter could be achieved as density variations in photographic film– in practice the phase variations could not be controlled independently, and it was not until 1964 that Vander Lugt developed a technique by which arbitrary amplitude and phase could be encoded onto a filter [13], in the form of a Fourier transform hologram of the target recorded on a holographic plate. More recently matched filters have been achieved using computer generated holograms, such as those achieved by the phase detour method of Lohmann and Paris [14]. These filters have the advantage over photographic holograms of greater ease of fabrication due to their binary amplitude transmittance.

A further complex filter which has been investigated is the inverse filter, which using the notation above is defined as

$$T_2(\xi,\eta) = \frac{1}{G(\xi,\eta)}. \tag{2.9}$$

This filter also gives rise to a sharp correlation peak in the output, but is undefined at points where $G(\xi,\eta) = 0$. This problem has been addressed by a number of methods, such as restricting the filter to a range of spatial frequencies up to the first singularity [15], or by adding a constant or function A, say, to $G(\xi,\eta)$ such that $G(\xi,\eta) + A > 0$ [16].

2.3.2 Single modulation parameter filters

Because of the difficulty of controlling both phase and amplitude in complex filters and the high resolution necessary when using a spatial light modulator to display a computer generated hologram, the emphasis has shifted towards filters with just a single modulation parameter, phase or amplitude.

The importance of phase in signals has been demonstrated by Oppenheim and Lim in their paper of the same name [17]. This paper shows that many of the important features of an image are preserved if only the phase information in

the Fourier spectrum is retained, while retaining only the amplitude information does not yield a recognisable image on further Fourier transformation. Further, images can often be reconstructed with remarkable fidelity to the originals when their particular amplitude spectra are replaced by a spectrum derived statistically from an ensemble of images. The implication is that the spectral information which distinguishes between objects is contained in the phase spectrum, with the amplitude spectrum differing little between objects. This is borne out by noting that in many cases the amplitude spectrum approximates to a Gaussian distribution decreasing towards the high spatial frequencies. This in turn suggests that discarding amplitude information and replacing it by some constant value will lead to the higher frequencies being accentuated when the modified spectrum is used to reconstruct the object, giving the effect of edge enhancement [17]. As noted earlier this is advantageous in sharpening correlation peaks and reducing cross correlation.

Another attraction of phase only filters in addition to reduced complexity is that 100% diffraction efficiency is theoretically attainable if the amplitude transmittance of the filter is unity [18], an important point given the low diffraction efficiency of holographic filters.

Phase only filtering has been demonstrated successfully by many authors, with favourable results when compared with conventional matched filters. Horner and Gianino have performed correlation simulations using two binary alphanumeric characters as the inputs, comparing the classical matched filter described above with the phase only and amplitude only counterparts derived from it. This demonstrated the superior performance of the phase-only filter (POF) in terms of discrimination between the characters, height and sharpness of correlation peaks, and peak-to-sidelobe ratio for the correlation peaks. This performance was maintained in the presence of additive noise in the input, whereas the already very poor performance of the amplitude-only filter deteriorated rapidly with increasing noise.

As noted by Horner and Gianino, this performance is not gained at no cost, as the emphasis of the high spatial frequency components in the Fourier plane leads

to high frequency noise components in the output, resulting in a poorer signal to noise ratio in the correlation plane than for the classical matched filter. This problem has been addressed by Kumar and Bahri [19] who describe a modified 'Optimal phase only filter' with signal to noise performance better than that of the classical matched filter. This is achieved by limiting the passband of the filter to some level determined by using the cutoff frequency as a parameter which can be varied to maximise some performance indicator of 'goodness of correlation'– in this case signal to noise ratio.

It should be noted that definitions of signal to noise ratio are not unique. The definition used thus far, and which is implicit unless otherwise stated, is that defined most clearly by Farn and Goodman [20]. Stated simply, this is

$$SNR = \frac{\text{Maximum correlation peak intensity}}{\text{Mean square value of output noise}}. \tag{2.10}$$

An alternative definition is offered by Horner and Bartelt [21]:

$$SNR = \frac{\text{Maximum correlation peak intensity}}{\begin{array}{l}\text{rms of intensity response outside a 50\%} \\ \text{threshold level in the correlation plane.}\end{array}} \tag{2.11}$$

This is justified as being tailored for practical correlation assessment. Under this definition a narrow correlation peak with low sidelobes (as desired) will yield a high SNR. Note that using this definition phase-only filters have better SNR performance than the matched filter [21].

Clearly different performance indicators can be indicated– in the previous example of Kumar and Bahri, correlation peak sharpness decreases with improved SNR– and sophisticated algorithms have been developed which optimise the filter with respect to a number of indicators, using information about the form of the input signal and limitations of the filter medium [22].

2.3.3 Binary phase only filters

The inability of most conventional SLMs to represent a continuously variable phase spectrum leads to consideration of quantisation of the phase spectra to levels which can be represented accurately by the available SLMs. In practice this generally requires the phase spectrum to be quantised to only two levels, due to the binary mode of operation of most electrically addressed SLMs. Ternary (three level) phase quantisation has been implemented using magneto-optic SLMs [23] but will not be considered here.

Effects of binarisation

The effects of phase quantisation of the Fourier spectrum has been investigated by Goodman and Silvestri [24], who note that the image obtained by inverse Fourier transformation of the modified spectrum is made up of an attenuated version of the original object (the *primary image*) plus contributions from *false images*, whose number and positioning are related to the number of quantisation levels.

With respect to binarisation, one false image may appear which may overlap with the primary image. If we choose the binarised phase values to be 0 and π, we obtain a real transmission function with values -1 or +1. By using the symmetry property of the Fourier transform which states that if an object $g(x, y)$ consists of only real values its transform $G(\xi, \eta)$ will be Hermitian, i.e. its real part will be even and its imaginary part will be odd [7, page 192], it is apparent that the false image is present to provide the proper symmetry. Though the choice of a function with real values -1 and +1 may seem arbitrary, in practice binary phase only filters (BPOFs) are often designed to realise this as it provides the largest possible modulation depth (π). Further, Horner and Leger note [25] that any binary phase function may be expressed as a real binary function with values -1 or +1 which has been multiplied by a complex constant and added to another complex constant.

The symmetry result shows that it is possible to construct a BPOF with no quantisation error if the reference target is a Hermitian function, e.g. a real symmetric object centred on the origin.

Alternatively the filter can be generated from any given target by calculating its phase spectrum and quantising the phase to two values using some threshold condition, with the consequence that the input object will correlate with the false image in the filter's spectrum in addition to with the primary image.

Binarisation methods

Using a real input function $g(x,y)$ as the reference target, Psaltis et al. [26], constructed a binary phase-only filter (BPOF) by calculating the Fourier transform $G(x,y)$ and quantising to obtain the filter $G_{Psal}(\xi,\eta)$ according to

$$G_{Psal}(\xi,\eta) = \begin{cases} +1 & \text{if } \mathcal{R}e[G(\xi,\eta)] > 0 \\ -1 & \text{otherwise} \end{cases} \tag{2.12}$$

whereas Horner and Bartelt [21] obtained a filter $G_{Horn}(\xi,\eta)$ according to

$$G_{Horn}(\xi,\eta) = \begin{cases} +1 & \text{if } \mathcal{I}m[G(\xi,\eta)] \geq 0 \\ -1 & \text{otherwise,} \end{cases} \tag{2.13}$$

with $\mathcal{R}e[G(\xi,\eta)]$ and $\mathcal{I}m[G(\xi,\eta)]$ being the real and imaginary parts of $G(\xi,\eta)$ respectively.

As noted by Cottrell et al. [27] these binarisation algorithms are closely related to the cosine transform, $C(\xi,\eta)$, and sine transform, $S(\xi,\eta)$, of $g(x,y)$, where

$$G_{Psal}(\xi,\eta) = \begin{cases} +1 & \text{if } C(\xi,\eta) > 0 \\ -1 & \text{otherwise} \end{cases}, \tag{2.14}$$

with a false image contribution equal to $g(x,y)$, and

$$G_{Horn}(\xi,\eta) = \begin{cases} +1 & \text{if } S(\xi,\eta) \geq 0 \\ -1 & \text{otherwise} \end{cases} , \tag{2.15}$$

with a false image contribution equal to $-g(x,y)$.

In both cases the false image results in correlation signals from a version of the target rotated through 180^0 that are equally as strong as that from the target itself– however, in the case of the sine version of the filter the false image contribution to the correlation is opposite in phase to the primary image contribution, whereas the two contributions are in phase for the cosine version of Psaltis et al. An immediate consequence of this is that the cosine BPOF lends itself to detecting targets with twofold rotational symmetry, with the filter calculated using a reference placed relative to the origin such that the primary and false images coincide exactly. This has the effect of doubling the amplitude of the correlation peak due to addition of the primary and false image correlations. A similar argument shows that the same improvement in performance can be obtained using a sine BPOF to detect antisymmetric targets– the antiphase contribution from the false image correlation comes from a false image itself opposite in phase to the primary image, giving a net addition.

For unsymmetric targets the choice of filter is less clear cut. One option is to modify the reference input so that it is symmetric, by placing the original unsymmetric reference in one half of the input plane opposite its rotated counterpart in the other half, as suggested by Horner and Leger [25]. This is illustrated in figure 2.4 which also illustrates the main drawback of this idea– if the original unsymmetric input occupies NxN pixels on a spatial light modulator at the input to the optical processor the modified input will require $2N$x$2N$ pixels– a doubling of the space bandwidth product (SBP) which is neither desirable or often possible. A further consequence is a concomitant increase in SBP necessary in the Fourier plane filter, a problem which is discussed in detail in Chapter 7.

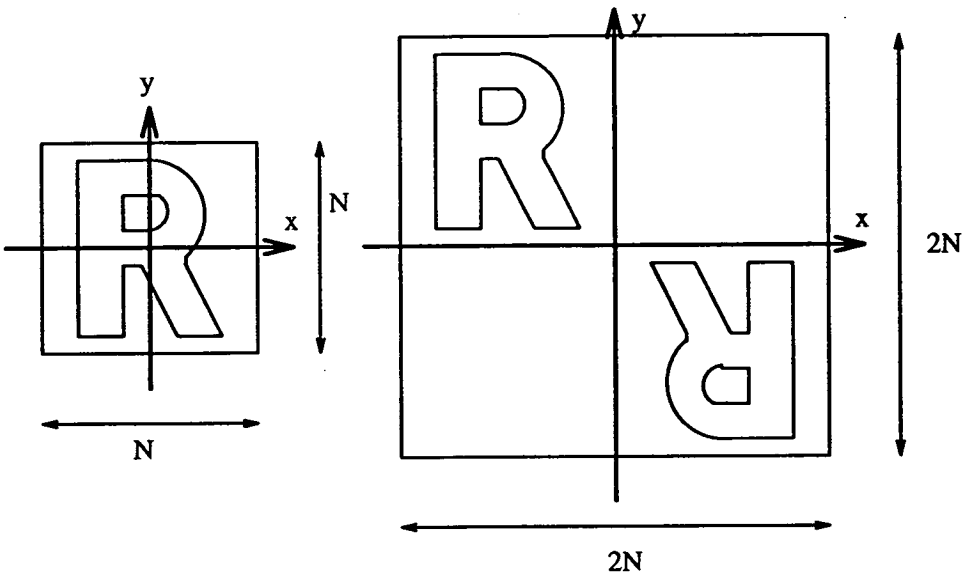

Figure 2.4: Construction of symmetrical reference input

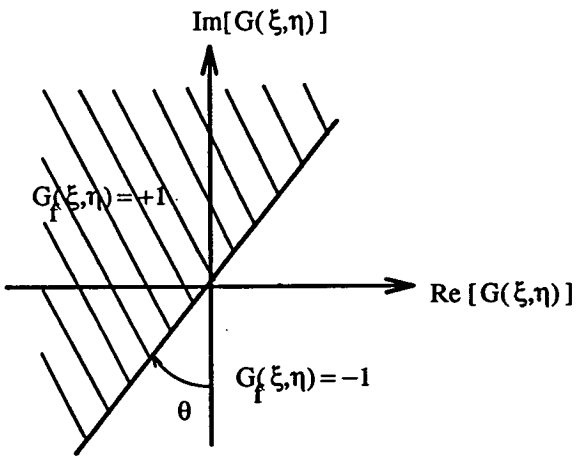

Figure 2.5: Complex plane representation of threshold condition

An alternative option is to use a threshold condition for the binarisation which results in a filter which is related to some linear combination of both sine and cosine transforms.

The threshold condition can be represented as a straight line bisecting the complex plane as shown in figure 2.5

Noting that the $\theta = 0^0$ and $\theta = 90^0$ corresponds to the cosine and sine BPOFs

respectively, an obvious compromise takes $\theta = 45^0$ as the threshold angle. This in turn corresponds to the Hartley transform BPOF described by Cottrell et al. [27], with the Hartley transform defined by

$$H(\xi,\eta) = C(\xi,\eta) + S(\xi,\eta) \tag{2.16}$$

Thus

$$G_{Cott}(\xi,\eta) = \begin{cases} +1 & \text{if } H(\xi,\eta) \geq 0 \\ -1 & \text{otherwise} \end{cases}, \tag{2.17}$$

Whereas cosine and sine BPOFs are clearly strong candidates for detecting symmetric and antisymmetric targets respectively, the case for the Hartley transform based BPOF for detecting general non-symmetric targets is less certain as the target may be predominantly odd or predominantly even. Nevertheless, the signal to noise performance of the binarised Hartley phase-only filter has been shown to be within a factor of 4 of the optimised POF of Kumar and Bahri mentioned earlier [28], and has been demonstrated to reduce the correlations with the rotated version of the target [27].

As in the case of continuously varying POFs, optimisation of the BPOF can be carried out by varying the filter's properties (such as threshold angle, phase retardances) to maximise some parameter, such as SNR [20]. It turns out that for symmetrical targets Psaltis' cosine BPOF is the optimal BPOF.

2.3.4 Amplitude encoded binary phase only filters (AEBPOFs)

Amplitude encoded phase only filters have been developed to exhibit advantages of POFs or BPOFs (e.g. good discrimination and sharp correlation peaks) using amplitude modulators such as film or inexpensive liquid crystal televisions. As proposed by Flavin and Horner [29] this involves encoding the phase information

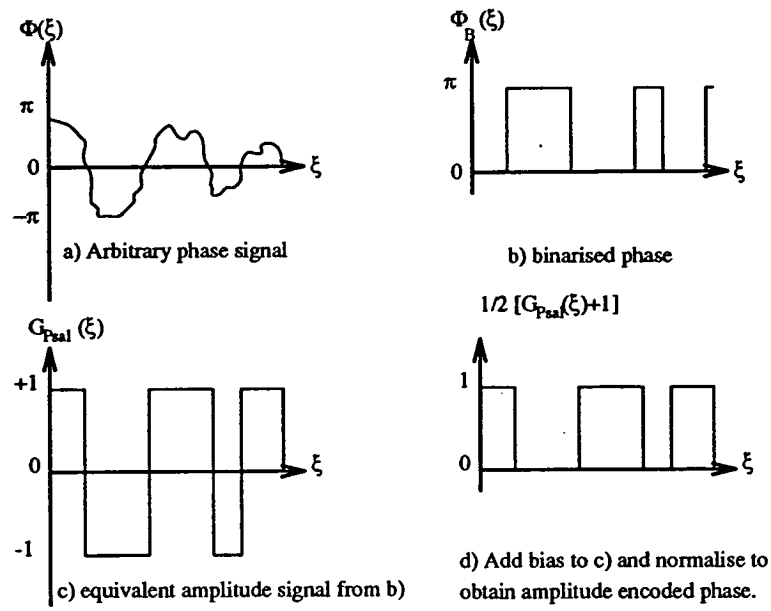

Figure 2.6: Coding of phase as amplitude

as amplitude information by adding a d.c. bias signal to a binarised phase pattern and normalising, as shown in figure 2.6.

As shown in the figure, the binarisation of the phase is performed using the rule proposed by Psaltis et al. to yield a filter function with two real values. Writing this function as $G_{Psal}(\xi,\eta)$, it is apparent that when a d.c. bias of 1 is added to remove any negative values and normalisation takes place to give a function appropriate to a passive filter, the resultant filter function will be

$$T_2(\xi,\eta) = G_{AEBPOF}(\xi,\eta) = 1/2[G_{Psal}(\xi,\eta) + 1]. \tag{2.18}$$

As desired, this function is clearly closely related to the cosine BPOF from which it is derived, though there is no reason in principle why any binarisation routine binarising to the states 0 and π could not be employed equally straightforwardly. This close relation is manifest in the output of the processor, which, assuming an input of $f(x,y)$, will be

$$u_3^+(x,y) = \frac{1}{2}\, g_{Psal}(x,y) * f(-x,-y) + \frac{1}{2}\, \delta(0,0) * f(-x,-y)$$
$$= \frac{1}{2}\, g_{Psal}(x,y) * f(-x,-y) + \frac{1}{2}\, f(-x,-y). \qquad (2.19)$$

Thus the output is the same as that of the cosine BPOF (multiplied by a factor of 1/2) except for the addition of a reproduction of the input function. As the latter contribution is known it can conceivably be subtracted digitally from the output, though Flavin and Horner note that typically the correlation peak will be much larger than the contribution from the input image. Their simulations also show that the intensity of the correlation peak for the AEBPOF is approximately 1/4 that of the corresponding BPOF as expected, and that the signal to noise performance of the two filters is similar (190 for the AEBPOF and 228 for the BPOF). Here the definition of the SNR is that of Horner and Bartelt.

Clearly the AEBPOF is a useful class of filter: Chapter 7 presents simulated and optical results using filters from this class.

2.3.5 Fourier plane filters for scale and rotation invariant pattern recognition

The algorithms and filters described thus far have taken no account of distortions in the input scene such as variations in scale or orientation of the target(s) with respect to the reference object. In some applications these distortions can be engineered out by registering the objects to be correlated in a given position and orientation, such as checking for defects in bottles on a production line [30], but otherwise their consequences must be faced, and possible solutions to any degradation in the filter's correlation performance evaluated.

The degradation in performance is in fact profound– when using a matched filter to detect a target in a detailed image the SNR has been seen to drop from 30dB to 3dB with only a 2% scale change, and from 30dB to 3dB with a 3.5^0 rotation [31].

Clearly this will severely restrict the applicability of the optical correlator unless solutions can be found. This problem has received much attention, and many modified Fourier plane filters proposed which offer varying degrees of distortion invariance. These will not be reviewed in detail, but some of the more popular solutions are introduced here.

Coordinate transformations

This class of filter specifically deals with scale and rotation distortions by transforming the input object into a space where rotational and and scale distortions in the input have no effect, then performing the filtering operation in that space.

A coordinate transformation which is suitable and which can be performed optically is the polar-Mellin transform [31], effected by a transformation to polar coordinates followed by a log transform along the radial axis. This can be achieved by addition of a preprocessing stage to the optical processor consisting of a holographic element and lenses performing the required transformation [32, 33].

The net result of such an arrangement is that the scale and rotation distortions are mapped to translations in a plane at the input of a shift invariant optical processor, thus removing the effect of the distortions on correlation peak magnitudes. However the polar-Mellin transform itself is not shift invariant, necessitating either centering of the input relative to the reference, or performing the coordinate transformation on the magnitude Fourier transform of the input, which contains no phase information about spatial offsets in the input [32]. A further drawback is that the SBP required to represent accurately the polar Mellin transform is much larger than that of the original input function– for an input function with a possible 200% scale change, to represent the polar Mellin transform with 2% accuracy requires around 5 times the SBP [34].

Synthetic discriminant functions

An alternative technique is to synthesise a filter which detects a set of objects made up of distorted versions of the target, with 3-d distortions possible in addition to simple variations of scale and rotation. In the original formulation of Hester and Casasent [35] each object in the *training set* (i.e. a set of a limited number of different perspective views of the target) is described as an orthonormal basis function expansion, e.g.

$$f(x) = \sum_i a_i \omega_i \tag{2.20}$$

where f is a member of the training set $\{f_k\}$, a_i is an expansion coefficient and $\{\omega_i\}$ is the set of basis functions. The 1-d notation is used for simplicity. An 'average' reference function, the synthetic discriminant function, is calculated which is a linear combination of the basis function expansions, which is itself a basis function expansion, e.g.

$$\hat{h}(x) = \sum_k c_k f_k. \tag{2.21}$$

The filter $H^*(\xi)$ is then calculated and used in the Fourier plane of the standard optical processor.

The original formulation of SDFs was reported to have poor discrimination ability [36], and variants have been developed to produce sharp correlation peaks with low sidelobes, such as minimum average correlation energy (MACE) filters [37, 38].

In general these have assumed the existence of complex filters and have thus seldom been implemented optically. Horner and Gianino have however applied their phase only filter concept to the SDF filter, to achieve improved performance over the classical matched filter based SDF in terms of correlation peak height ($\sim 49\times$ higher intensity) and sharpness ($\sim 10\times$ smaller correlation spot area)

[39]. A further improvement in peak height has been achieved with an SDF-BPOF constructed using a method which takes account of the filter modulation [40].

Simulated annealing

In common with the previous SDF method for the sythesis of filters, simulated annealing requires the use of a training set of objects encompassing the expected distortions in the target. The simulated annealing method however is dependent on a sophisticated optimisation algorithm which allows the inclusion of a set of non-target objects whose cross correlations when using the final filter are to be minimised, whilst maximising the correlations of the targets. Clearly some compromise must reached to allow this, and maximise the discrimination between the target and non-target objects.

Conceptually, simulated annealing achieves this by an iterative process whereby random changes are made to the filter (e.g. switching a pixel in a BPOF from one state to the other) and the resulting changes in the correlation performance measured– if the performance is improved, the changes are accepted unconditionally; if the performance is degraded, the changes *may* be accepted provided the degradation is not *too much*; otherwise the filter is returned to its orginal state. The process is then iterated, with the allowable degradation in performance for the changes to be accepted being gradually reduced, until the filter converges to a steady state.

The analogy of this process with annealing becomes apparent when the process is analysed in a little more detail. The correlator is considered as a system with an energy value E which is to be minimised by cooling from a high temperature to a low temperature, repeatedly perturbing the system variables in a random manner as the system cools. The variables in this case are the changes which can be realised in a practical filter. These changes affect the energy of the system, which is dependent on the correlation values of the objects in the training set in some way, the dependence determined by what is to be optimised. The change in

energy, ΔE, is measured and the perturbations allowed to stand if ΔE is negative. The changes *may* be allowed to stand if ΔE is less than some positive value which is not *too much*, determined by the acceptance probability $P(\Delta E)$, equal to

$$P(\Delta E) = \frac{1}{1 + e^{(\Delta E/T)}}. \tag{2.22}$$

This shows the dependence on the temperature parameter T. The temperature should decrease slowly enough so that the system does not get trapped in local energy minima [41].

Kim and Guest have applied the simulated annealing algorithm to BPOFs, using filters generated to discriminate between the letters R and P. In simulations they used training sets of 5 letters P in a range of fonts, and five letters R in a similar range of fonts, and successfully classified a set of 10 letters R and P which did not belong to the training sets [41]. Similar impressive performance has been demonstrated optically, with greatly improved discrimination ability when compared with the corresponding cosine BPOF [42].

The applicability of the simulated annealing algorithm to distortions of scale in addition to the small variations in font has also been demonstrated [43]. The general applicability of the method to problems where the distortions can be encompassed in a training set make it attractive, especially as it could be further applied using a real optical system to optimise an imperfect SLM for specific correlation tasks.

2.3.6 Joint transform correlators

The joint transform correlator (JTC) is a competing correlator architecture, described and implemented as long ago as 1966 [44], which has become popular recently with the advent of real-time reconfigurable SLMs. Joint transform correlators share similarities with the basic correlator architecture described already,

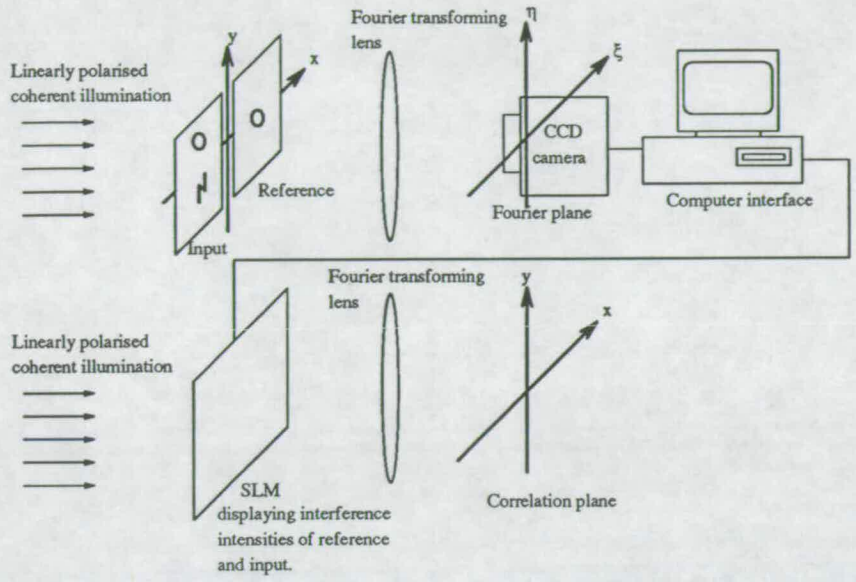

Figure 2.7: Joint transform correlator schematic (after Javidi and Odeh).

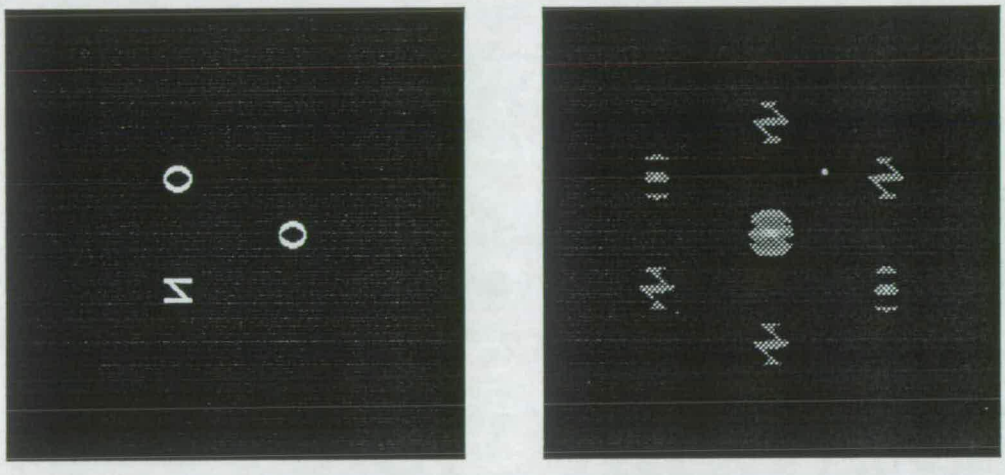

a) Input, showing reference O on right

b) Output, showing zero order and correlation peaks

Figure 2.8: Simulation of joint transform correlation

such as the formation of two successive Fourier transforms between input and output. However they differ in several crucial respects as is apparent with reference to the schematic shown in figure 2.7 (after Javidi and Odeh [45]).

The input scene and the target with which it is to be compared are presented side by side at the input, and both Fourier transformed simultaneously so that their transforms' interference intensities may be recorded at the Fourier plane by some square law detector, such as a CCD array camera. The resulting intensity spectrum (the *joint power spectrum*) is then redisplayed on a uniformly illuminated SLM and retransformed to give an output containing a dominant on-axis peak, and lower peaks corresponding to cross-correlations between the constituent parts of the image, such as between the reference on the one side of the input plane and any corresponding targets which may exist in the other. The positions in the output of the cross-correlation peaks relative to the on-axis peak are indicative of the relative positions of the features in the input which correlate and give rise to them.

That the output should be of this form is obvious when we realise that the nonlinear operation of squaring the interference pattern is equivalent to multiplying the Fourier transform of the combined input by a filter which is its complex conjugate. Thus the JTC is a special case of the matched filter correlator which correlates the input with itself, giving rise to the strong correlation peak on-axis, and the smaller off-axis cross-correlation peaks which are of interest.

The main advantages of the JTC are that it does not require filters to be calculated as they are effectively formed by updating the input [46], and that the joint power spectrum redisplayed in the second half of the JTC is real, requiring an SLM with only a single modulation parameter.

The JTC also suffers disadvantages however. We have performed a simulation of joint transform correlation (shown in figure 2.8) using a letter O as the reference, to be compared with the remainder of the input which contains an O and an N. The output of this simulation shows six correlation peaks due to each letter correlating with every other, and shows how the interpretation of the output becomes highly complicated with even a small amount of clutter [47]– with the corresponding classical matched filter correlator the input letters O and R would correlate with only the reference letter O, giving just two correlation peaks whose positions could be related unambiguously to the relative positions of reference and target.

Further disadvantages include the need to use an area of the input SLM to display the reference, reducing the amount of the SLM available for displaying the input scene (equivalently, the space-bandwidth product), and the high dynamic range of the joint power spectrum to be displayed in the second part of the correlator. The latter problem has led to the binarised or bipolar joint transform correlator of Javidi and Odeh [45] which quantises the joint power spectrum to two levels, +1 and -1, which can be displayed on a binary phase only SLM. This gives a vastly improved optical efficiency (with correlation peaks between target and reference $> 10^5$ times larger) and better discrimination between target and non-target objects when compared with the 'classical' JTC, but has been shown to give rise to spurious correlation peaks, affecting its viability in practical use [48].

2.4 Summary

This chapter has introduced the standard Fourier plane filtering architecture and established a system of coordinates which will be applied to it throughout this thesis. The technique of Fourier plane filtering has been shown to be useful in applications which cover a wide range of complexity, from simple annular stops to implement edge enhancement or halftone removal, through more complicated spatial filters to extract or emphasise orientational features, to the detection of objects by optical correlation using complex filters. The convenience of reconfigurable SLMs for the latter application has been noted, and concepts of optical correlation introduced: shift invariance of the Fourier transform; matched filtering; the importance of phase in signals which allows the use of POFs with high diffraction efficiency; the ability to quantise, at the cost of a 'false image' which will also correlate with the input; the relation of binarisation schemes to the symmetry of the target; and sensitivity to distortions. Several techniques have been introduced which attempt to overcome the lack of distortion invariance of the correlator– these will be pursued no further, but the applicability of simulated annealing to a wide range of optimisation problems should be noted, particularly as it can take into account the imperfections which generally exist in currently available SLMs. Joint transform correlation has been considered, but is wasteful

of space-bandwidth product, and is difficult to interpret when the input scene is cluttered.

Binary phase only filters provide a practical means of optical correlation, and can be readily implemented using available SLMs. The use of these, and the closely related amplitude encoded BPOFs, will be described in Chapter 7.

Chapter 3

Liquid crystals and spatial light modulators

In this chapter a brief review of available SLM technology shows the prevalence of liquid crystal as a light modulating medium. Modulation of light using nematic liquid crystals is introduced, and a specific nematic liquid liquid crystal is assessed for use in a transmissive electrically addressed SLM. First, however, we consider the modes of modulation possible, and the manner in which a pattern of modulation may be imposed upon the the SLM (i.e addressing).

More detailed reviews of SLM technology can be found in papers by Casasent [49], Fisher and Lee [50], and Johnson et al.[51]. Reports on the newest devices can be found in the relevant topical issues of Applied Optics, and Optical Society of America Technical Digests on Spatial Light Modulators and Applications [52].

3.1 Spatial light modulation

The mechanism by which an SLM modulates light, the multiplication of the scalar amplitude of the light incident on the SLM by the SLM's transmission function, is stated for the coherently illuminated case by the equation

$$u_2^+(\xi,\eta) = T_1(\xi,\eta) \times T_2(\xi,\eta), \tag{3.23}$$

where $T_1(\xi,\eta)$ is the complex amplitude of the illumination incident on the SLM, $T_2(\xi,\eta)$ is the complex transmission function of the SLM, and $u_2^+(\xi,\eta)$ is the modulated light [1]. Ideally the SLM should allow control of both amplitude and phase at each point over its area, but the difficulty of independent control of each parameter has led to SLM research concentrating on amplitude-only or phase-only modulators. Even with these constraints, continuous modulation of either phase or amplitude at each point of the SLM is often not possible, with the modulation being quantised to usually two discrete levels. Whether or not the SLM is binary or continuous is closely related to its mode of addressing.

3.2 Addressing

The manner in which a pattern of modulation is imposed on an SLM makes it fall into one of two categories: optically addressed or electrically addressed. These two forms of addressing are shown schematically in figure 3.1.

3.2.1 Optically addressed SLMs

For the optically addressed SLM the pattern of modulation is impressed by a write beam on one side of the device which is sensitive to light. The relation of the pattern of modulation to the pattern of the write beam is dependent on the transfer characteristic of the SLM, which varies from device to device. The modulation is then impressed upon a 'read beam' of light incident on the other side of the SLM. Consequently optically addressed SLMs are generally reflective devices, with a dielectric mirror isolating the 'read' and 'write' beams.

[1] Note that if the SLM is a reflective device as opposed to transmissive, the effect is simply to replace $T_2(\xi,\eta)$ with $R_2(\xi,\eta)$, the complex reflectance of the SLM.

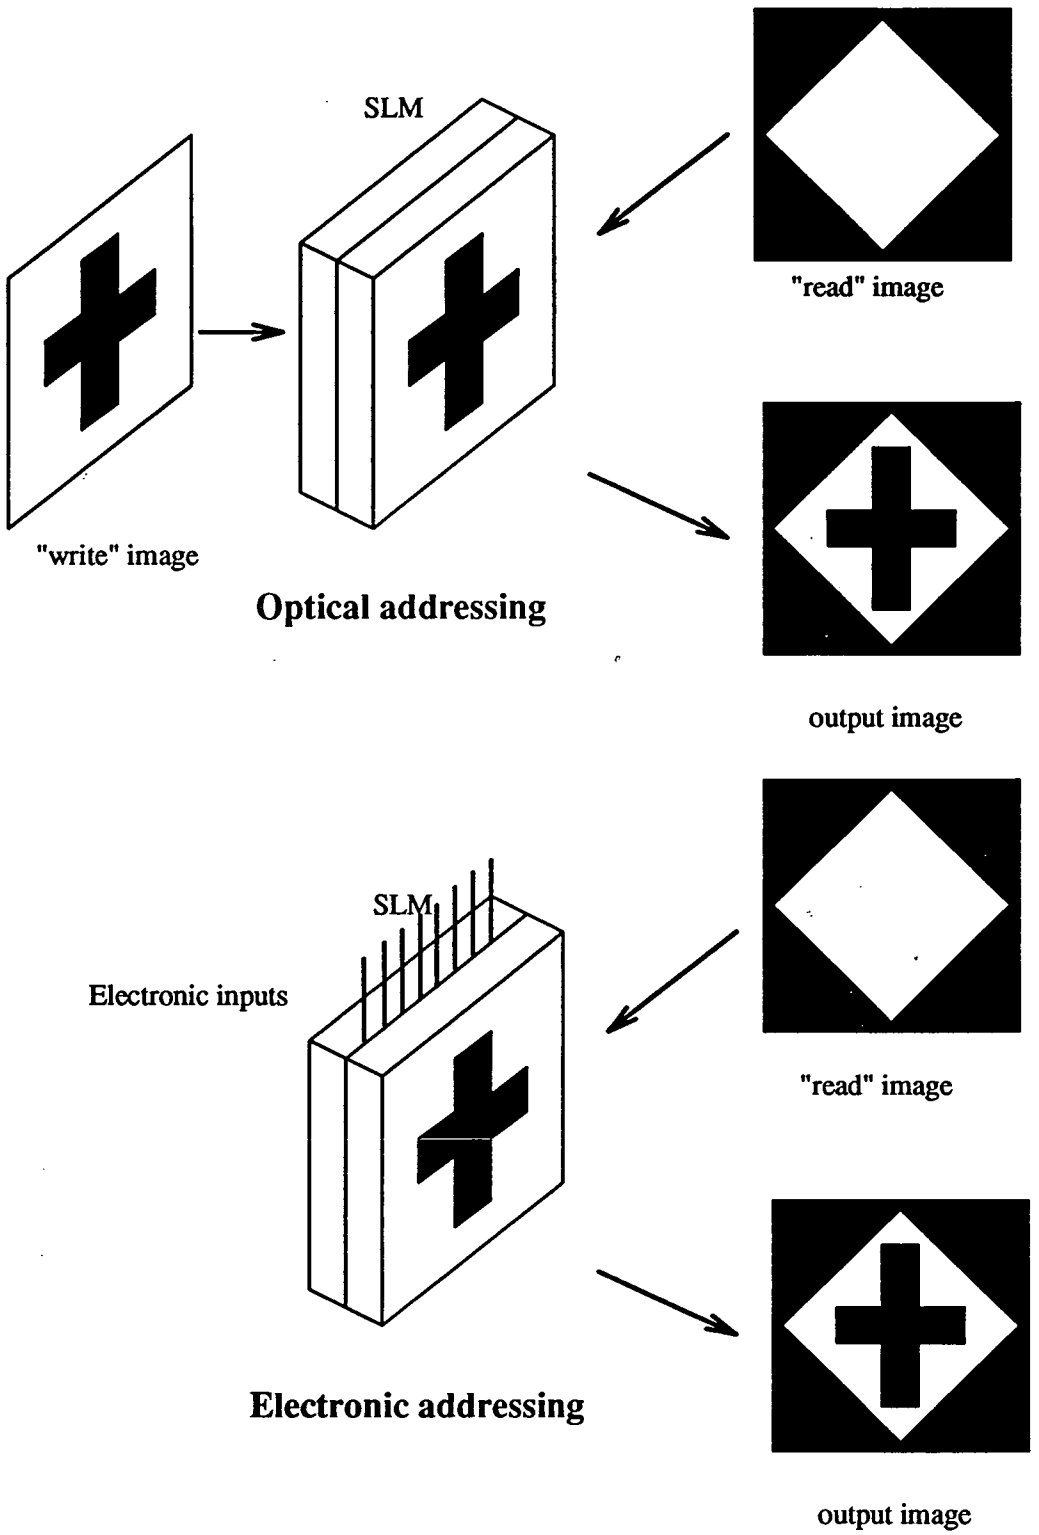

Figure 3.1: Schematic showing optical and electrical addressing (after Ranshaw)

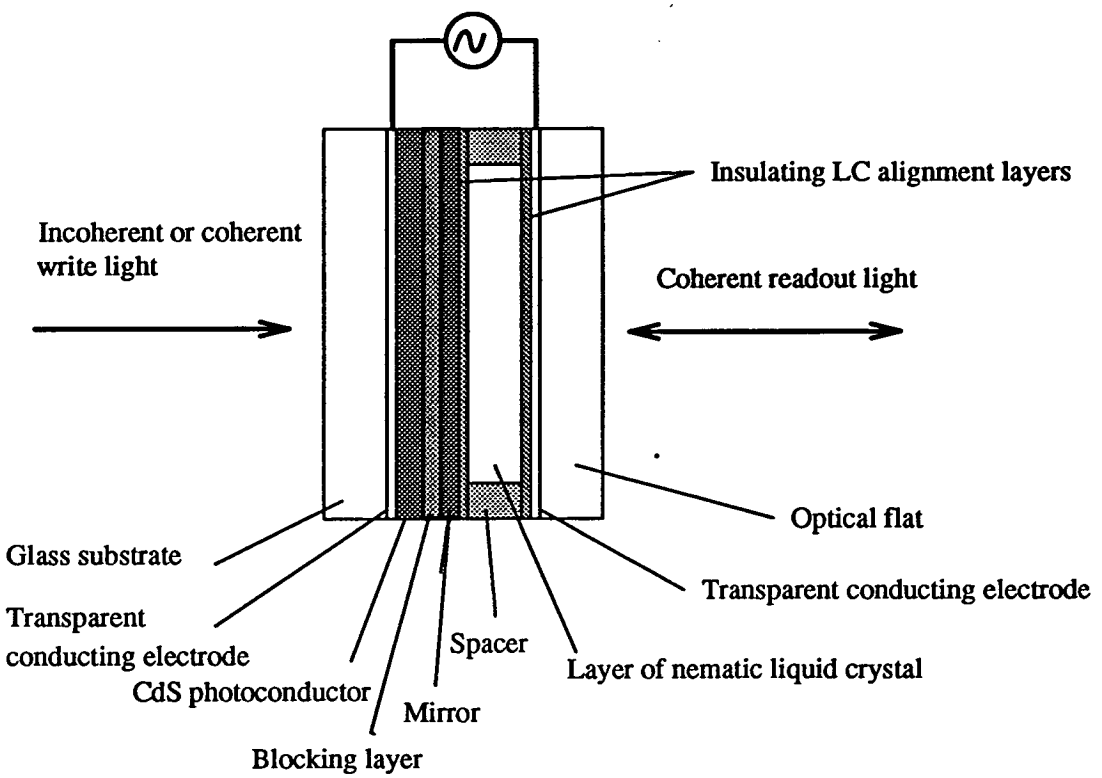

Figure 3.2: The liquid crystal light valve

Liquid crystal light valves

The Hughes Liquid Crystal Light Valve (LCLV) [53] is one of the earliest (1975) successful reconfigurable SLMs, and is the forerunner of many similar devices. The basic structure of the LCLV is shown in figure 3.2.

The cross-section shows how the LCLV is made up from layers of photoconductor, dielectric mirror (or mirror and light blocking layer) and liquid crystal sandwiched between a pair of transparent electrodes. Modulation is possible due to the electro-optic response of the liquid crystal, discussed later. The amplitude of the voltage dropped across the the LC layer is varied spatially by the write beam, the intensity of which varies the impedance of the photosensitive CdS layer. Depending on the response of the liquid crystal to applied voltage, it is possible to achieve a linear relation between write intensity and degree of modulation, thus allowing continuous modulation of phase or amplitude. Alternatively, the electro-optic response

of the liquid crystal may exhibit a sharp threshold characteristic leading to binary or near-binary modulation. Note that the write light need not be coherent, thus the LCLV may act as an incoherent-to-coherent converter.

The first Hughes LCLV was not pixellated, with resolution limited by diffusion of charge carriers in the photoconductor. Later devices such as the Hughes LCLV from 1985 [54] and an LCLV device from GEC-Marconi [55] replaced what was effectively a large area photodiode on the 'write' side of the LCLV with a pixellated array of microdiodes, to prevent charge from diffusing beyond the pixels. The pixellation is extremely fine– the GEC-Marconi SLM has pixels on a 14μm pitch over a diameter of approximately 20mm, giving a space-bandwidth product of approximately 2000. This should be compared with the SBP of electrically addressed devices. At this point it is sufficient to note that the space-bandwidth product is simply the number of pixels across the SLM. Chapter 4 contains a fuller description of this important quantity.

3.2.2 Electrically addressed SLMs

With an electrically addressed SLM, the pattern of modulation across its area occurs in response to some set of external electrical signals. Such devices are almost always pixellated with electrical connection to the pixel made by electrodes, but a notable exception is the electron-beam addressed lithium niobate SLM from Optron Systems Inc. [56]. In this latter device the pattern of modulation is written onto the $LiNbO_3$ crystal by an electron gun, offering good grey scale performance but low resolution and frame rates (the rate at which the entire SLM pattern can be updated) compared with other electrically addressed SLMs. We will consider briefly four of the more established technologies.

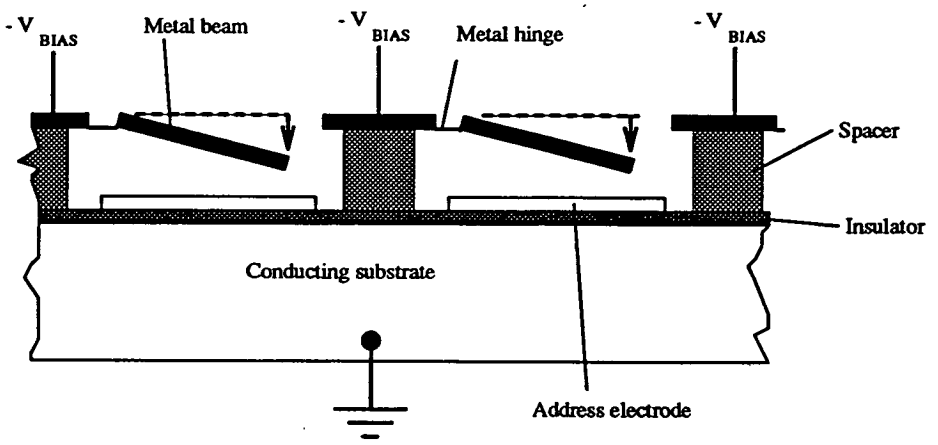

Figure 3.3: Principle of the Deformable Mirror Device (after Hornbeck). Each pixel consists of a hinged mirror (or mirrors) which deflect in response to an electric field.

The magnetooptic SLM

The original magnetooptic SLM (MOSLM) [57] consisted of an array of electrically addressable pixels of a bismuth garnet film, which affected the linear polarisation of transmitted light by means of the Faraday effect. Though capable of high switching speeds (frame rates of ~1kHz), widespread use of such devices has been prevented by low transmission, and failures due to magnetic forces and heating as a result of the necessary high drive currents [58]. Research continues into this technology, and recently a new reflection mode MOSLM has been presented [59] with 128×128 pixels and drive currents 3-5 times lower than for the earlier transmissive device.

The deformable mirror device

This technology has been pursued for a number of years, primarily by Texas Instruments, and is reviewed concisely by Hornbeck [60]. The principle is shown in figure 3.3.

The modulation is achieved by means of a mirror which deflects in response to

an electrical input, and with a suitable cantilever arrangement the reflected light can be modulated in direction and/or phase. Although Texas Instruments have demonstrated considerable technological achievements in producing arrays of up to 128×128 pixels, with pixels switching in 10-20μs, these devices have yet to become available to the general scientific community.

Multiple quantum well devices

These devices consist of pixellated multilayer semiconductor structures (typically many periods of InGaAs/InP, grown by molecular beam epitaxy [61]) forming the wells, whose electronic properties are altered by application of an electric field. In particular the application of the field leads to a shift in the absorption edge, allowing modulation of light tuned to appropriate wavelengths when the multilayer structure is grown on a dielectric mirror. The shift in the position of the absorption edge is a manifestation if the quantum confined Stark effect. As the shift of the absorption edge is of the order of 10nm, the optical bandwidth of the modulator is correspondingly small, necessitating the use of a highly stable light source. The resulting modulators are very small (quantum well structures 50Åwide have been reported [62]) and can be scaled to large arrays when the switching voltages are applied by a matrix addressing scheme. Fast switching is also a feature– AT&T Bell Labs have demonstrated digital optical systems using multiple quantum well self-electrooptic effect devices (SEEDs) being clocked at a frequency of 1MHz [63].

Liquid crystal spatial light modulators

Electrically addressed liquid crystal SLMs can be divided into two categories– devices developed from display technology, and liquid-crystal-over-silicon SLMs of the type pioneered at the University of Edinburgh. For both categories, in common with the LCLV, modulation is achieved as a result of the liquid crystal's birefringence and response to applied voltage.

Devices developed from liquid crystal displays

This category encompasses transmissive matrix addressed devices such as the liquid crystal television (LCTV) [64]. Here applied voltage per pixel is the difference in voltage on the overlapping electrodes either side of the liquid crystal which define the pixel. Matrix addressing of the pixels is discussed in section 3.4.

LCTVs have been directly adapted for use in optical processors, but have low contrast and a frame rate of the order of tens of hertz, limited by the nematic liquid crystal used. Poor uniformity and optical flatness are also problems, but simple techniques are available which offer improvements in performance [65, 66].

Matrix addressed SLMs using ferroelectric liquid crystal are a development of this technology, and have been fabricated in arrays as large as 1280×1120 pixels. Despite the fast switching times of ferroelectric liquid crystals (>1000 times faster than nematics, which switch in times of the order of 100ms) the rates at which an entire SLM can be updated are still low because of the matrix addressing [51].

Liquid-crystal-over-silicon SLMs

This class of SLMs is reviewed in more detail in a paper by Johnson, McKnight and Underwood [51].

These devices, as their name suggests, are made up from a thin layer of liquid crystal sandwiched between a cover glass and a custom VLSI silicon backplane, on which a regular array of mirrors is defined. An electrical signal can be applied to the metal mirrors using the circuitry on the silicon, which, in conjunction with a transparent electrode on the cover glass, applies a voltage to the liquid crystal over the mirrors. Because of the regular mirror structure the device is inherently pixellated, as is apparent in figure 3.4.

Over the last nine years work on this technology at the University of Edinburgh has produced a number of innovative devices. The first such device, a 16×16 pixel amplitude modulator with a memory cell at each pixel, was designed and built

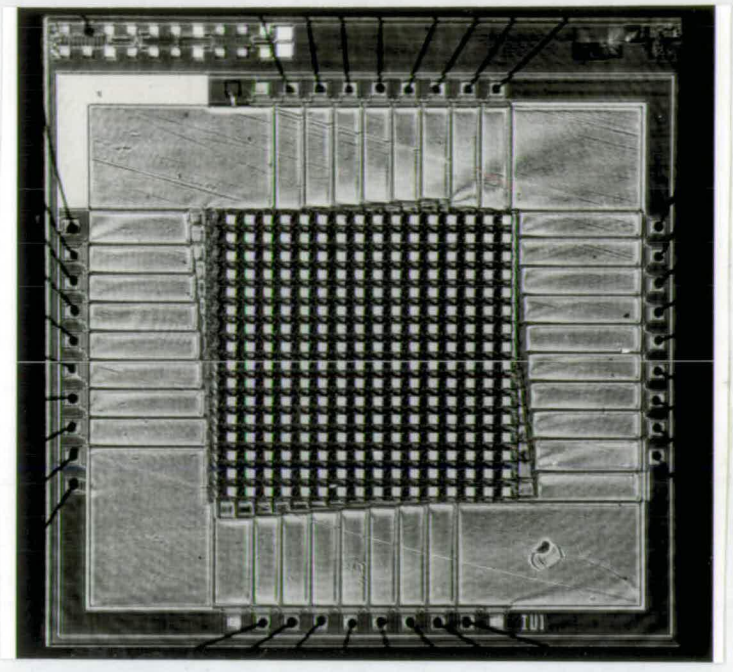

Figure 3.4: Liquid-crystal-over-silicon SLM. The figure shows the packaged IC before the cover glass and liquid crystal is added (courtesy of M.Ranshaw).

by Underwood for his PhD thesis [67], and subsequently modified for use as a phase modulator by Ranshaw [68]. McKnight in 1988 presented a 50×50 pixel SLM [69], and more recently a 176×176 pixel SLM has been fabricated [70]. This latter SLM has one of the largest numbers of pixels on a device of this kind to date, and with the pixels on a pitch of 30μm approaches the resolution of optically addressed SLMs.

Originally this technology was attractive as a means of achieving high frame rates, as the addressing does not rely on multiplexing, but on the use of on-chip logic on the silicon. More recently the role of the silicon has expanded to give rise to so-called 'smart' SLMs [71]– devices where at each pixel there is logic, memory or sensor functionality attached to the mirror. 'Smart' SLM research is currently being pursued by groups at the University of Edinburgh [72] and the University of Colorado at Boulder [73].

In common with other liquid crystal based SLMs, this technology has been spurred on by improvements in liquid crystal materials, with earlier devices using nematic liquid crystal, and later devices using ferroelectric liquid crystal.

3.2.3 Summary

Optically addressed SLMs have several attractive features such as high space-bandwidth product, grey scale capability, and the ability to perform simple image processing operations such as image addition and edge enhancement directly. Against them are their high cost and problems with uniformity over their areas. Of more relevance to this project is the inability to write directly an arbitrary input to them such as computed matched filter. For this reason an optically addressed SLM often is used as a simple incoherent to coherent converter at the input to optical processors, or as a square law detector in the transform plane of a joint transform correlator.

With electrical addressing any pattern can be written to the SLM, subject to the SLM's characteristics. These will be typically a relatively low space-bandwidth product, pixellation, and binary phase or amplitude modulation. Despite these limitations, the flexibility of the addressing is ideal for the use of sophisticated Fourier plane filter algorithms in matched filtering applications, or for arbitrary input an optical processor. Smart liquid-crystal-over-silicon devices offer the promise of increased functionality through processing elements at each pixel.

3.3 Modulation using nematic liquid crystals

Custom SLMs designed and fabricated during this project are described later in the thesis (chapter 6). As the SLMs required only 16×16 pixels, the simple reliable technology of matrix addressed nematic liquid crystal displays was adopted. Here we describe briefly the mechanism by which nematic liquid crystals may be used to modulate amplitude and phase, and show how it is possible to select pixels independently using a multiplexed drive scheme. A much more detailed description of the interactions of these liquid crystals with light and applied fields may be found in the PhD thesis of MacGregor [74] or standard texts such as Blinov [75] or Chandrasekhar [76]. A review article by Clark et al. [77] provides a useful introduction to nematic liquid crystal devices.

Although ferroelectric liquid crystals have the desirable properties of speed and bistability (resulting in inherently binary modes of operation unless time multiplexing takes place) the actual fabrication of liquid crystal cells and SLMs is much less reliable than for nematics, due mainly to the extremely thin ($\sim 3\mu$m) liquid crystal layer necessary. A detailed introduction to the use of ferroelectric liquid crystals is provided by Skarp and Handschy [78].

3.3.1 Basic concepts

Liquid crystal is the name given to a class of materials which can exist in phases which possess some degree of positional or orientational order, yet retain the flow properties of fluids. These phases, lying between those of a crystalline solid and an amorphous liquid (*mesophases*), may only exist over a certain temperature range, outwith which transitions to other mesophases may occur. Such behaviour is described as *thermotropic*. The differing mesophases are distinguished by varying degrees of positional order. Two such thermotropic mesophases are the *nematic*, with orientational but no positional order, and the more ordered *smectic* C^* mesophase of ferroelectric liquid crystals.

The liquid crystals considered here are constituted of long 'rod-like' organic molecules, exhibiting anisotropic behaviour with particular regard to optical and electrical properties. In order that a physical system composed of these molecules may be characterised succinctly with regard to the molecules' ordering and anisotropic behaviour, it is useful to introduce the concept of the *director*, $\underline{\mathrm{n}}$. The director is a unit vector used to specify the local average molecular orientation, which in many liquid crystal mesophases is parallel to the time averaged orientation of the long molecular axis.

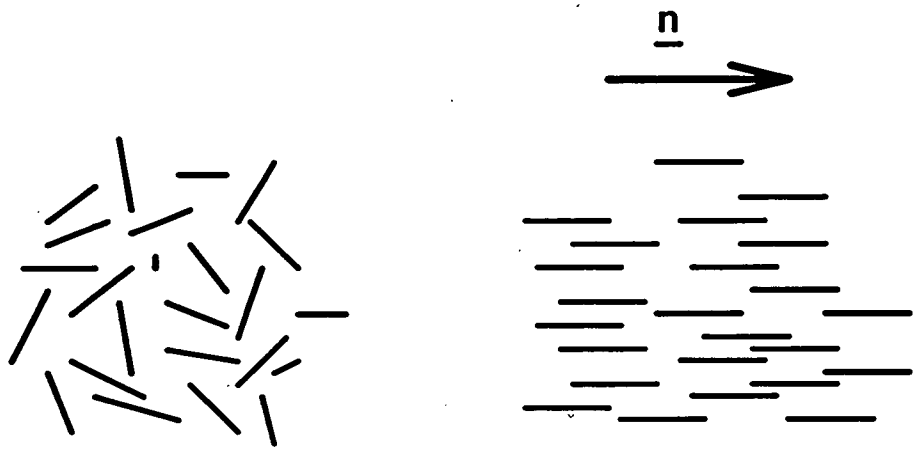

Figure 3.5: Orientational order of nematic mesophase, compared with an isotropic mesophase of the same molecules. Note how the nematic molecules tend to orientate themselves in the same direction.

3.3.2 Nematic liquid crystal

Figure 3.5 shows schematically how the molecules in a nematic liquid crystal tend to orientate themselves in the same direction, in the absence of external forces.

Nematic liquid crystal is a uniaxial medium, with both its dielectric permittivity and refractive index having different values perpendicular to the director from those parallel to the director. Thus the electrical and optical responses are coupled, and can be used to obtain electrically addressed modulation of light.

The orientational order can be altered by bounding surfaces, with a treatment which forces the liquid crystal molecules at the surface to align preferentially in some direction. Many surface treatments have been tried, giving varying degrees of alignment, and involving complicated interactions with the liquid crystal. A review of alignment treatments is carried out by Cognard [79]. We simply note that a suitable *alignment layer* may be achieved by rubbing microgrooves into a layer of polyvinyl alcohol (described in chapter 6), or by evaporating a thin ($\sim$ 140 Å) layer of silicon oxide onto the glass substrate at an oblique angle.

Phase modulation by Field Induced Birefringence

We consider the case where a nematic liquid crystal is constrained between two glass plates with a surface treatment which causes the liquid crystal at each surface to align in the same preferred direction, and with the director everywhere parallel to the surface (*homogeneous alignment*). We also assume that the liquid crystal has a positive dielectric anisotropy i.e.

$$\Delta\epsilon = \Delta\epsilon_{parallel} - \Delta\epsilon_{perpendicular} > 0 \tag{3.24}$$

where directions are relative to the director. This has the effect that the director will tend to align parallel to an applied electrical field. Using this reponse to an applied voltage (usually a.c. to prevent electrochemical degradation of the liquid crystal), we obtain a variable retarder as a result of the liquid crystal's birefringence. Light polarised parallel and perpendicular to the director at the surface sees material with refractive indices n_e and n_o respectively, when no voltage is applied. However, as the rms voltage across the the plates increases, the molecules tilt in reponse to the applied field until they are perpendicular to the surface of the plates, presenting a refractive index of n_o to any incident polarisation. Thus the ordinary optic axis of the retarder has a refractive index fixed at n_o, while the extraordinary axis can vary between n_e and n_o as a function of voltage. Hence we obtain a phase-only modulator by polarising the input light parallel to the extraordinary axis, with a depth of modulation specified by the rms voltage of the 'on' state. This is equal to

$$\begin{aligned} \Delta\phi(V) &= \frac{2\pi}{\lambda}(n(V) - n_o)d \\ &= \frac{2\pi}{\lambda}\Delta n(V)\, d, \end{aligned} \tag{3.25}$$

where d is the geometrical thickness of the liquid crystal cell and $\Delta\phi(V)$ is the depth of modulation. Figure 3.6 shows the form of the variation of this *field*

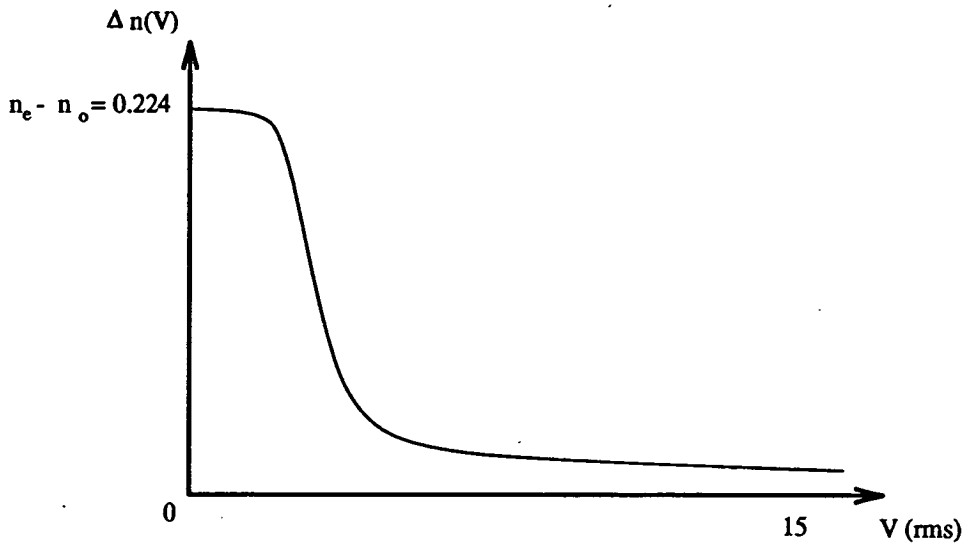

Figure 3.6: Variation of field induced birefringence with applied voltage for BDH E7.

induced birefringence (FIB) for the E7 liquid crystal from BDH. The high birefringence of this material ($n_e - n_o = 0.224$, [80]) allows modulation of π in cells $\sim 3\mu$m thick, using a HeNe laser for illumination. Field induced birefringence is used in the phase modulators of chapters 6 and 7.

3.3.3 Amplitude modulation using the twisted nematic effect

In the twisted nematic effect the preferred directions of alignment at the two bounding surfaces are at 90^0 to each other. This results in the orientation of the director rotating through 90^0 as the liquid crystal cell is traversed, for homogeneous alignment and in the absence of an applied electric field. If the pitch of the resulting helix satisfies Maugin's condition ($\lambda \ll (n_e - n_o)4d$, d is thickness) the twisted structure will lead to a rotation of the plane of polarisation of plane polarised light incident either parallel or perpendicular to the front director. Maugin's condition states that this will occur for cell thicknesses of $\geq 10\mu$m, using E7. If an electric field is applied, the molecules of the liquid crystal will tend to align with the direction of the field– at sufficiently high voltages the tilt of the

molecules relative to the bounding surfaces will be such that the twist structure is not apparent, and rotation of polarisation does not take place. Thus, by aligning an analyser with its plane of polarisation parallel to that of the input polariser, light will be blocked by the analyser for no applied voltage, and transmitted when the polarisation guiding is broken down by applied voltages. This principle is shown in figure 3.7. The output may be inverted by rotating the anayser through 90^0– transmission for no applied voltage, blocking of light for applied voltages above a certain threshold.

Gooch and Tarry [81] further show that the transmission is dependent on the thickness of the liquid crystal layer, the wavelength of the illumination and the liquid crystal's birefringence. They depict this relationship in a graph of transmission against $2d\Delta n/\lambda$ (the Gooch-Tarry curve), with the best contrast between 'off' and 'on' states obtained at the minima of the curve, corresponding to an 'off' state with zero transmission. The minima are obtained when $d = N\lambda/\Delta n$, N an integer. Satisfying this, the Maugin condition, and keeping the layer as thin as possible to increase the switching speed [82] gives the optimum layer thickness using E7 and illumination at 633nm as $\sim 11.2\mu$m. It should be noted however that beyond the second minima of the Gooch-Tarry curve, the transmission of the 'off' state is never more than a few percent. In this case 11.2μm corresponds to the fourth minimum, so errors in thickness of a few microns are tolerable.

The twisted nematic effect is widely used in displays, and is used in the amplitude modulators of the later chapters.

3.4 Multiplexed drive of nematic liquid crystal SLMs

When nematic liquid crystal is used in a pixellated SLM, consideration must be given to how each pixel can be independently switched between the 'off' and 'on' states (or non-transmitting and transmitting states respectively, in the case of an amplitude modulator). Note that binary operation has been assumed. That this is a sensible mode of operation is apparent when the electro-optic response of the

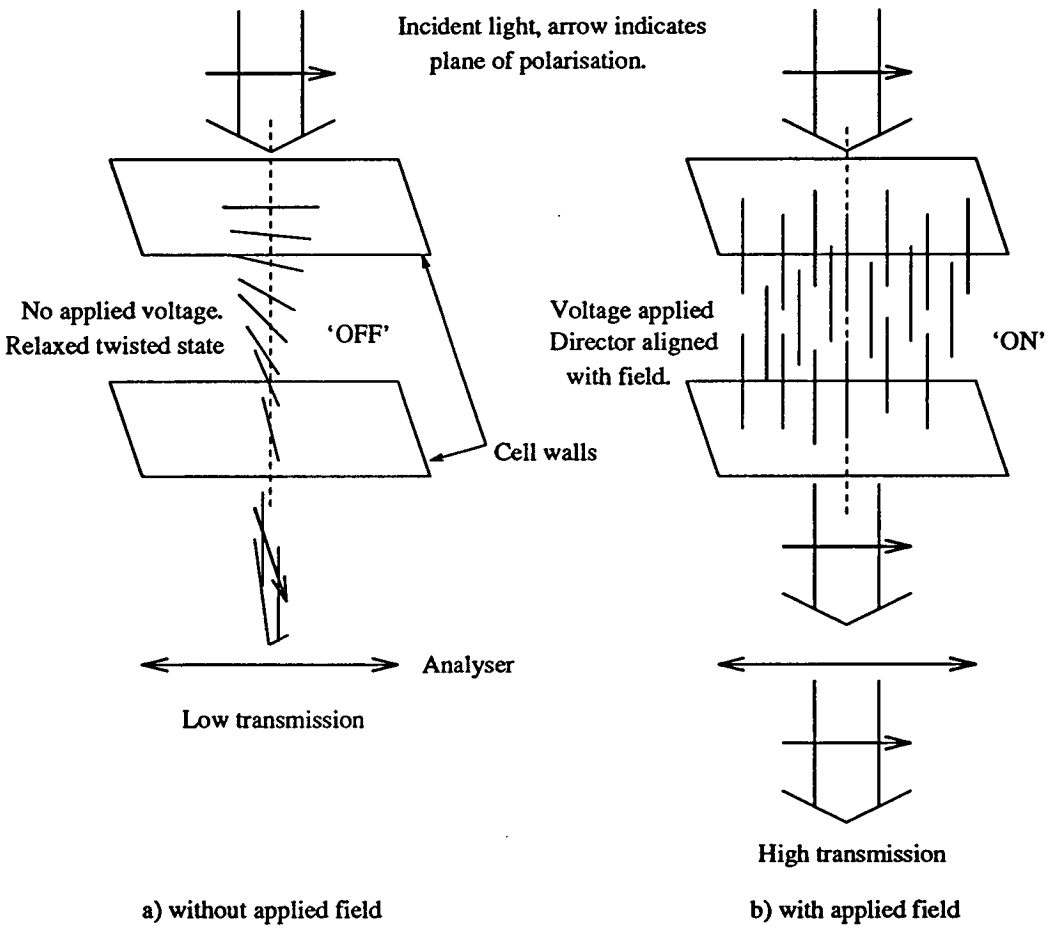

Figure 3.7: The twisted nematic effect, showing the orientation of the directors b) with, and a) without an applied electric field.

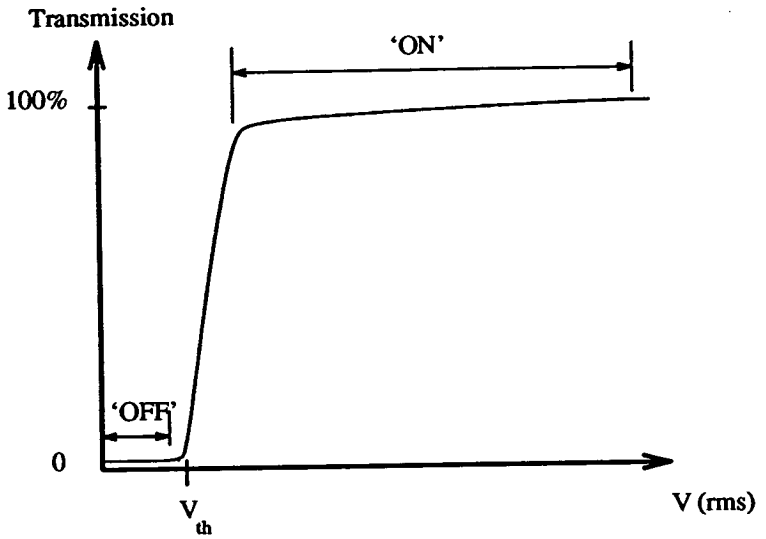

Figure 3.8: Form of the electro-optic response of a twisted nematic liquid crystal cell between aligned polarisers

nematic liquid crystal is considered. Figure 3.8 shows the typical response when the polariser and analyser are aligned– two ranges of uniform transmittance exist, with the transition between them taking place over a comparatively narrow range of voltages above a threshold voltage V_{th}.

Achieving well defined 'off' and 'on' states at each pixel is simple if they are individually driven, with an external electrical connection from each. The drawback here is that for an $N \times N$ array, N^2 external connections are needed, which soon becomes unwieldy. However, if the pixels are defined by the overlap of two sets of crossed electrodes, only $2N$ connections are present– is it possible to address each pixel independently in such an array? The answer is yes, for arrays up to a certain size. This is possible due to the sharp threshold characteristic of the liquid crystal's electro-optic response, and is illustrated by considering the drive scheme for a twisted nematic display of Alt and Pleshko [83] depicted in figure 3.9.

The figure shows 3×3 pixels from an $N \times N$ array of pixels, accessed by horizontal lines carrying 'strobe' voltages and vertical lines carrying 'data' signals. Information is supplied in parallel to the data lines, and the strobe lines scanned sequentially with a voltage pulse V_s, to build up the pattern of 'off' and 'on' pixels

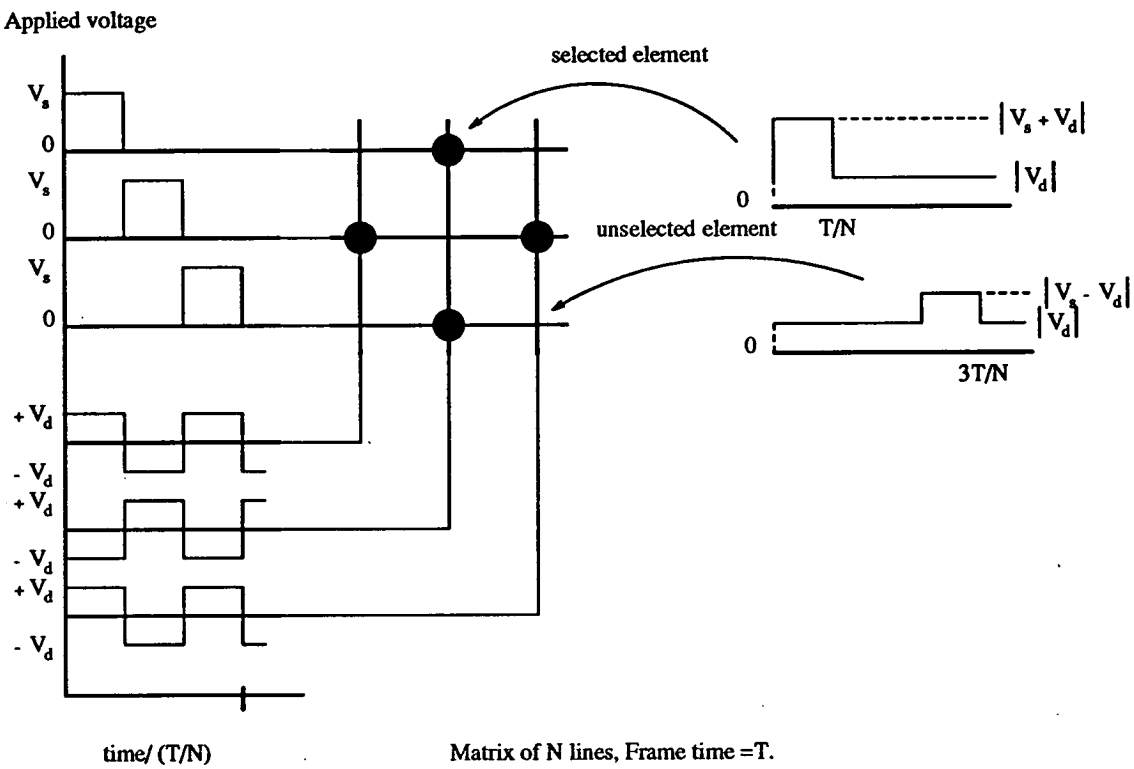

Figure 3.9: Multiplexed drive scheme of Alt and Pleshko (after Needham).

line by line. The information is entered such that a data voltage of $-V_d$ is supplied to the pixels to be switched during a strobe pulse, and $+V_d$ to the pixels whose transmission state is to be left unaltered. Thus a switched pixel recieves a pulse of magnitude $|V_s + V_d|$ while it is being addressed and pulses of magnitude $|V_d|$ while the other lines are being addressed. Similarly the unswitched pixels receive a pulse of magnitude $|V_s - V_d|$ while being addressed, and $|V_d|$ while the other lines are being addressed.

The frame period, i.e. the time taken to address all N strobed lines, is significantly shorter than the time taken for the nematic liquid crystal to switch between 'off' and 'on' states ($\sim$ 100ms), so the steady state response of each pixel is determined by the rms voltage across it over a frame period. Therefore the rms drive voltage of the 'on' pixels must lie in the 'on' region shown in figure 3.8 and the rms voltage of the 'off' pixels lie in the 'off' region, for successful independent switching of each pixel.

3.4.1 Multiplexibility

Alt and Pleshko note that the number lines which can be addressed successfully using this drive scheme is a function of the threshold characteristic of the electro-optic response, and that increasingly large numbers of lines requires an increasingly sharp threshold characteristic. Needham [84] defines a quantity M, the *voltage discrimination ratio*, as

$$M = \frac{\hat{V}_{ON}}{\hat{V}_{OFF}} \tag{3.26}$$

where $\hat{V}_{ON}$ is the rms voltage when the display is in its minimum acceptable 'on' state and $\hat{V}_{OFF}$ is the rms voltage when the display is in its minimum acceptable 'off' state. From this figure the maximum number of scanned lines possible can be estimated using

$$N_{max} = \left(\frac{M^2+1}{M^2-1}\right)^2 . \tag{3.27}$$

M can be obtained from the graph of the electro-optic response shown in figure 3.8. Analytic expressions for V_d and V_s may also be obtained in terms of M, but in practice are best established experimentally by iteration for each particular liquid crystal cell.

3.4.2 Choice of suitable liquid crystal for multiplexed drive

Reference has already been made to BDH E7, a long established commercially available liquid crystal which is often used in twisted nematic displays where each display element is directly driven. For multiplexed drive we consider instead BDH E300, a newer liquid crystal with a sharper threshold characteristic.

In common with E7, E300 has a positive dielectric anisotropy, necessary for homogeneous alignment. E300 also has a high birefringence (0.139) allowing phase modulation of π (or higher odd multiples of π) to be carried out in cells of the order of 10μm thick. A cell thickness of 13.7μm corresponds to the third minima of the Gooch-Tarry curve, so 12μm spacers are acceptable.

The BDH datasheet for E300 states that it is suitable for 16-way multiplexing, without reference to contrast ratio. A more quantitative figure may be obtained from quoted figures for the voltages at which 90% transmission and 10% transmission are obtained. These values give a voltage discrimination ratio of 1.34, indicating that contrast of 9:1 may be obtained for up to 12 multiplexed lines, using equation 3.27. This estimate of the multiplexibility was confirmed in the following experiment, to determine the form of the electro-optic response of E300.

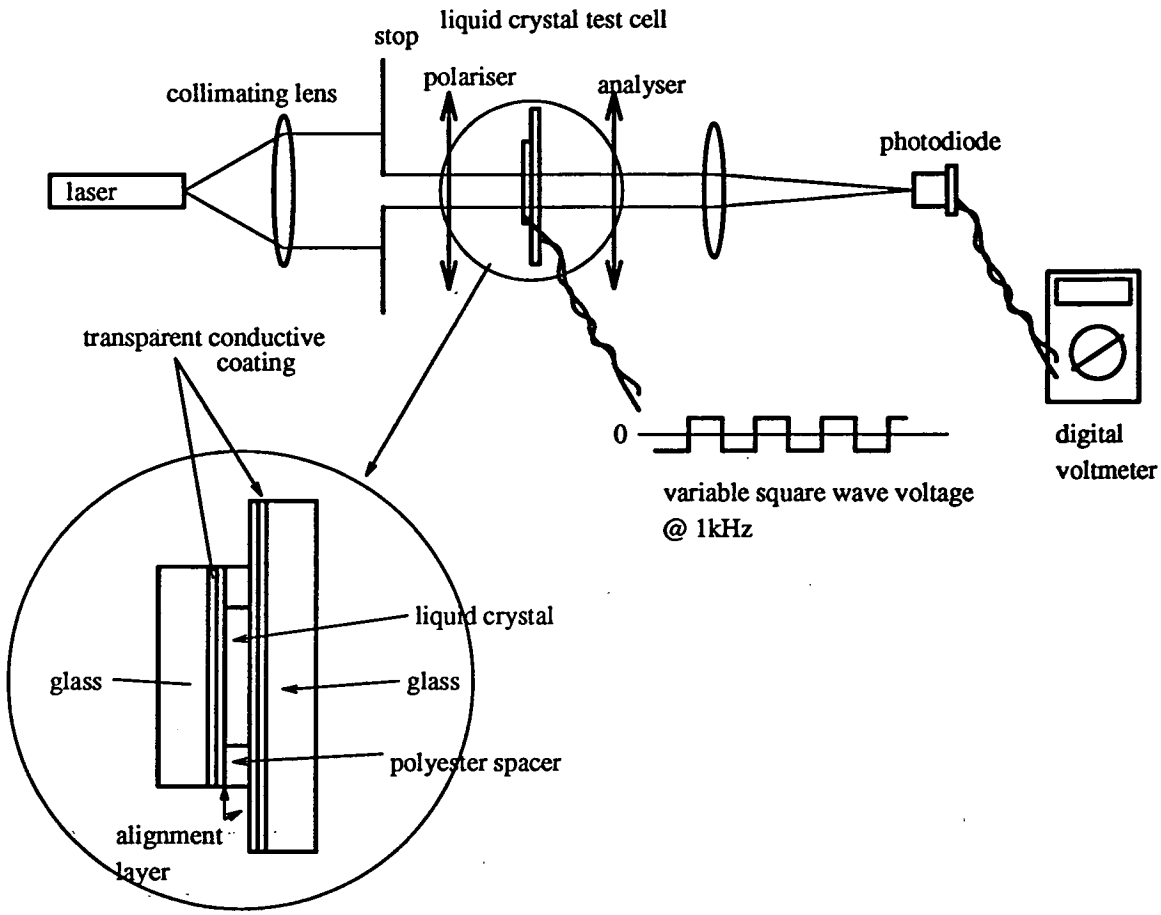

Figure 3.10: Experimental arrangement to determine the electro-optic response of a liquid crystal

Experimental determination of electro-optic response curve

This was done using an E300 liquid crystal test cell. The experimental apparatus is shown in figure 3.10. The equipment and method follows that of Ranshaw [68].

The test cell was made by sandwiching a layer of liquid crystal between two glass plates, each with a sputtered coating of indium tin oxide (ITO) followed by an evaporated coating of silicon monoxide. The ITO acts as the necessary transparent conducting electrodes by which the drive voltages are applied to the liquid crystal, and the silicon monoxide is evaporated on at 60^0 to the plate normal to a thickness of 140Å[79] to provide an alignment layer. The cell thickness was 12μm.

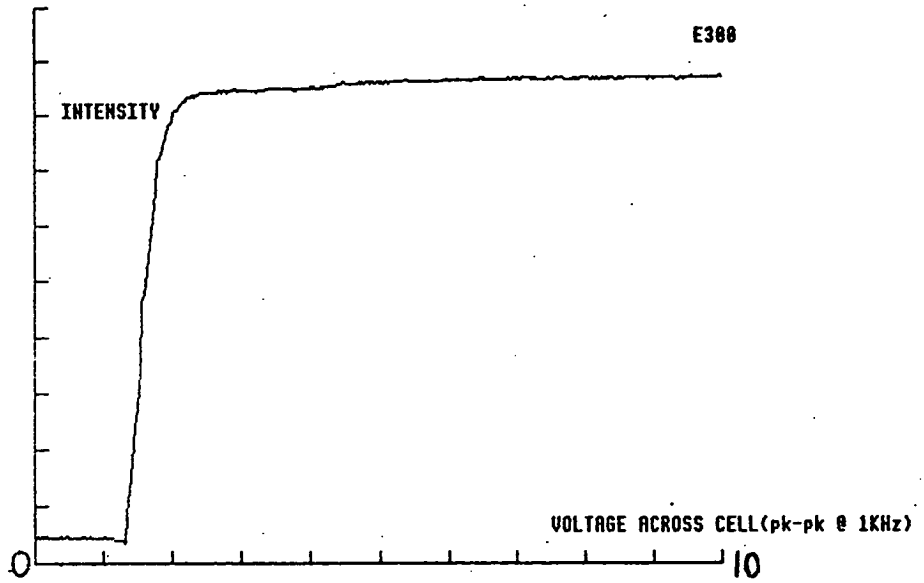

Figure 3.11: Electro-optic response of E300. 12μm thick cell placed between parallel aligned polaroids.

A small uniform area of the cell was illuminated by collimated light from a helium-neon laser. The cell was placed between parallel polaroids aligned with the silicon monoxide evaporation direction on the side of the cell nearest the laser. A square wave voltage at 1kHz was placed across the cell, and increased in steps from 0 to 10V peak-to-peak by means of a stepper motor attached to a variable resistor. The transmitted intensity was measured using a linear photodiode, with a liquid crystal settling time of 200ms allowed between each intensity measurement and subsequent increase in the applied voltage. The resulting curve is shown in figure 3.11.

The voltages recorded at 10% and 90% transmission (allowing for subtraction of background) were (1.40 ± 0.02)V and (1.92 ± 0.02)V respectively. These figures were used to calculate a voltage discrimination ratio of (1.37 ± 0.03), and consequent multiplexibility of (11 ± 3) lines with better than 9:1 contrast, in agreement with the published figures. Some scaling error may have been present in the voltage measurements, but this should not affect the ratio. The curve shows the two ranges of near-constant transmission either side of the threshold characteristic.

Although the number of lines addressable while retaining acceptable contrast seemed small, the decision was taken to use E300 in the 16×16 matrix addressed

SLM, in the absence of a liquid crystal with a sharper threshold characteristic. However, contrast ratios greatly in excess of 9:1 were measured for a 16-way multiplexed device constructed in chapter 6. This suggests that the response to the multiplexed drive waveforms may be more complex than indicated by the model used, which assumes that the response is equal to that of a square wave applied voltage with the same rms value.

Chapter 4

Use of a regularly pixellated SLM in a 4-f optical processor

Current reconfigurable SLMs fall into two categories– electronically addressable, and optically addressable. The electronically addressed devices are inherently pixellated, and pixellation is also performed in some optically addressed liquid crystal devices to prevent resolution being degraded by spreading of photo-excited charge in semiconductor layers. Therefore it is important to be aware of the effects of pixellation on the output from an optical processor, and on the transform function of the filter itself.

The choice of filter will affect the form and quality of the retransformed image. Indeed the purpose of the filter is to perform some processing operation which will alter the form of the image. However any underlying structure in the filter function will also impose itself on the image through its Fourier transform, with possible effects on image quality. In addition to any structure effects, image quality will be degraded in a more general sense due to the filter being non-ideal, such as band-limiting due to finite size, low contrast ratios in amplitude filters, phase errors in phase filters etc.

With pixellated filters there is most definitely an underlying structure. Although the output of the optical processor is the Fourier transform of the filter function convolved with the coordinate reversed image of the object, it is possible to simplify the analysis by concentrating on the filter function and its transform without

reference to a specific object. The simplest case to consider first is that of identical pixels evenly spaced on a regular 2-d grid, forming a binary amplitude filter. In the remainder of this chapter the filters are considered with all pixels 'on' unless otherwise stated. In a processing operation this will not generally be the case, but the simplification allows the analysis to proceed without reference to the processing operation. Given the huge range of widely differing filter functions possible, this is considered the best approach for the general analysis.

It emerges that for a filter with regularly spaced pixels the output of the optical processor is made up of regularly spaced replicated images. Depending on various parameters of the optical processor, including pixel pitch and size, these replicas may overlap and confuse the output. These replicas arise due to sharp peaks in the filter's Fourier spectrum, and the idea of randomising the pixels' spacing and position is suggested as a means of suppressing these *spectral orders.*

4.1 Regularly pixellated amplitude filter

Using the same coordinate system as before we postulate a filter transmission function for a filter in the Fourier plane depicted in figure 2.1–

$$T_2(\xi,\eta) = p(\xi,\eta) * \sum_{n=0}^{Q-1}\sum_{m=0}^{Q-1} \delta(\xi - n\alpha, \eta - m\gamma) \tag{4.28}$$

(Q, n, m integers), corresponding to a QxQ array of identical regularly spaced pixels described by a pixel function $p(\xi,\eta)$, spaced on a pitch of α in the ξ-direction and γ in the η-direction, and where $\delta(\xi,\eta)$ is a two-dimensional Dirac delta function. This is illustrated in figure 4.2. In practice the filter will be placed approximately symmetrically with respect to the optic axis, corresponding to the central pixel being centered on the origin if there is an odd number of pixels, or a pixel nearest the centre centered on the origin if there is an even number of pixels. The description above is used in the analysis as it leads to the same result, and allows treatment of both even and odd values of Q without altering the form of the equation. A pixel is centered on the optic axis to allow the dominant d.c.

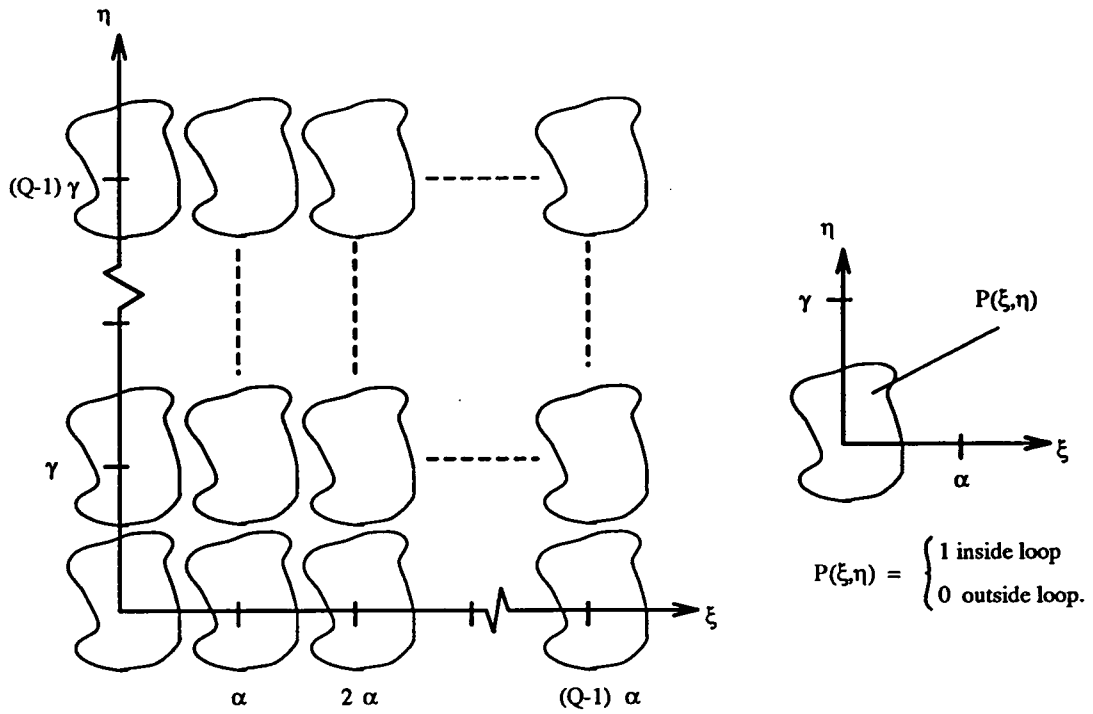

Figure 4.1: Regularly pixellated filter

part of the object's Fourier transform, falling at $u_2^-(0,0)$, to be transmitted by the filter.

We now wish to Fourier transform $T_2(\xi,\eta)$ to determine its effect on the image. From the separability of the Fourier transform it is possible to simplify $T_2(\xi,\eta)$ by considering only the 1-d case at first. Thus $T_2(\xi) = p(\xi) * \sum_{n=0}^{Q-1} \delta(\xi - n\alpha)$. The Fourier transform of this is treated analytically in two distinct ways: the first gives the spectrum of the filter as an easily calculated product for consistency with the analysis of randomised pixels carried out later; the second gives the filter spectrum in the form of a convolution with δ-functions which simplifies consideration of sampling.

4.1.1 Analysis in 1-d

Treatment 1

Fourier transforming $T_2(\xi) = p(\xi) * \sum_{n=0}^{Q-1} \delta(\xi - n\alpha)$ directly, gives

$$t_2(x) = P(x) \times \sum_{n=0}^{Q-1} e^{-i2\pi n\alpha x} \tag{4.29}$$

using the convolution and shift theorems, where $P(x)$ and $t_2(x)$ are the Fourier transforms of $p(\xi)$ and $T_2(\xi)$ respectively. This is of course equivalent to the amplitude distribution found at at plane $z = z_3$ if the filter plane $z = z_2$ is illuminated uniformly, as plane 1. Noting that $\sum_{n=0}^{Q-1} e^{-i2\pi n\alpha x} = \frac{\sin(\pi Q\alpha x)}{\sin(\pi\alpha x)} e^{-i\pi(Q-1)\alpha x}$ we can rewrite $t_2(x) = P(x) \times \frac{\sin(\pi Q\alpha x)}{\sin(\pi\alpha x)} e^{-i\pi(Q-1)\alpha x}$, and consider the terms in this expression individually:

- $e^{-i\pi(Q-1)\alpha x}$ is a phase factor arising from the unsymmetrical way the array was defined earlier– if the array was placed symmetrically about the origin such that summation was from $\frac{-(Q-1)}{2}$ to $\frac{(Q-1)}{2}$ it would disappear altogether. Further, it would be cancelled by its complex conjugate when calculating intensities in the image plane as we will do later– so it will be disregarded.

- $\frac{\sin(\pi Q\alpha x)}{\sin(\pi\alpha x)}$ is familiar from consideration of multiple slit interference and gives rise to δ–function like behaviour as Q becomes large, with the function having large peaks at the zeroes of $\sin(\pi\alpha x)$ [85, page 333]. These peaks alternate between Q and $-Q$, found by applying L'Hospital's Rule. [1] Figure 4.2 shows this function for $Q = 16, \alpha = 2$. The disturbances between the peaks oscillate at a frequency proportional to Q, and with a magnitude of order 1, independent of Q.

  Thus for larger Q the peaks will be narrower and much greater than the intermediate disturbances, approaching δ–function like behaviour, as illustrated in figure 4.3 for $Q = 176, \alpha = 2$. [2]

---

[1] $\lim_{x\to x_0} \frac{f(x)}{g(x)} = \lim_{x\to x_0} \frac{f'(x)}{g'(x)}$

$\Rightarrow \lim_{x\to n/\alpha} \frac{\sin(\pi Q\alpha x)}{\sin(\pi\alpha x)} = \lim_{x\to n/\alpha} \frac{Q\pi\alpha}{\pi\alpha} \frac{\cos(\pi Q\alpha x)}{\cos(\pi\alpha x)} = Q\,(-1)^{n(Q-1)}$

[2] The choice of $Q = 16$ and $Q = 176$ is to correspond with array sizes of spatial light modulators

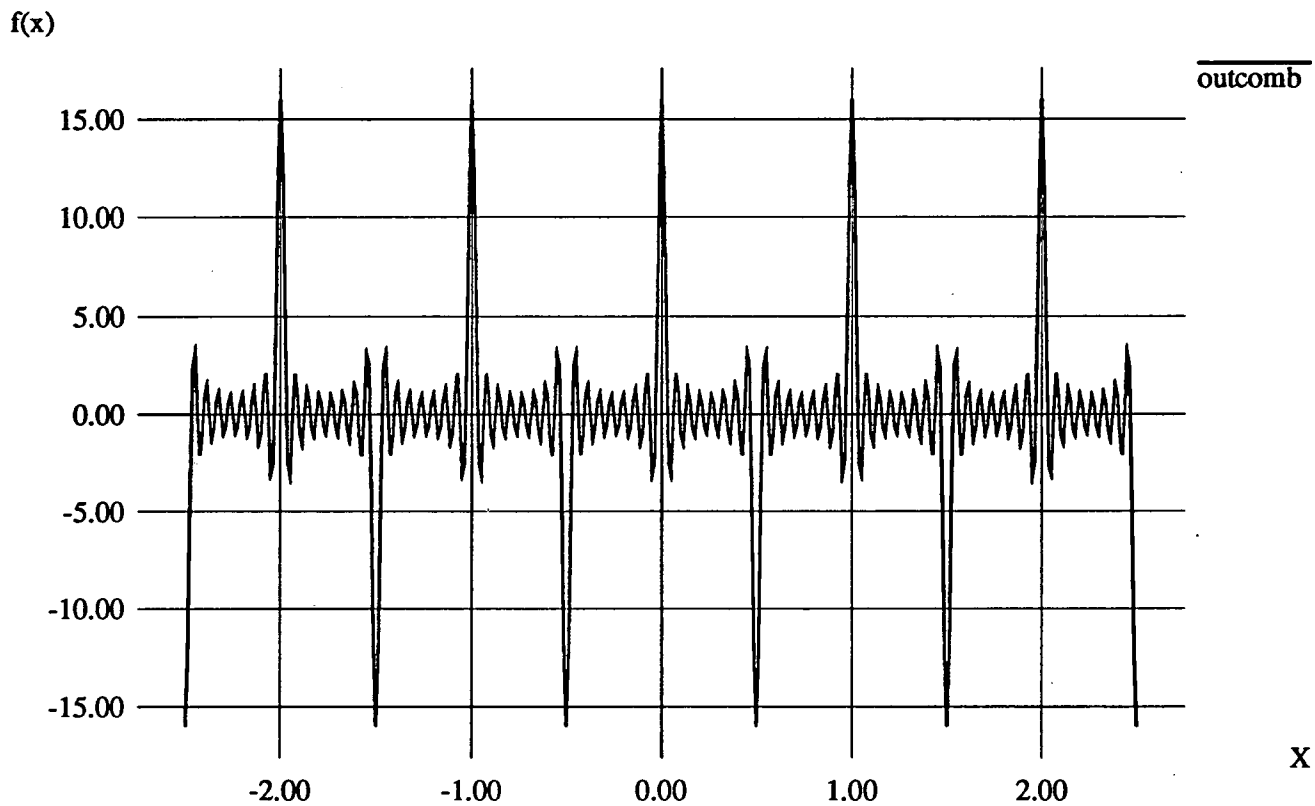

Figure 4.2: $\frac{\sin(\pi Q \alpha x)}{\sin(\pi \alpha x)}$ for $Q = 16, \alpha = 2$

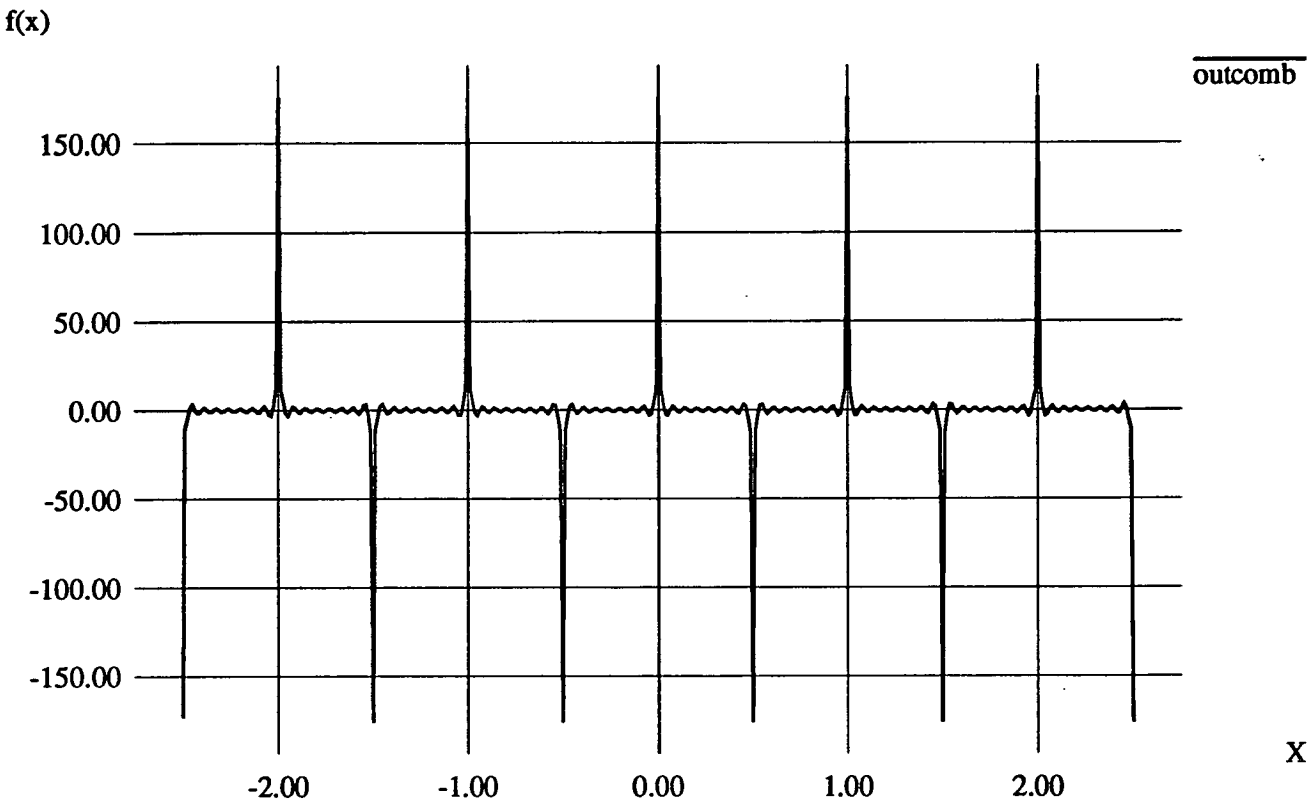

Figure 4.3: $\frac{\sin(\pi Q \alpha x)}{\sin(\pi \alpha x)}$ for $Q = 176, \alpha = 2$

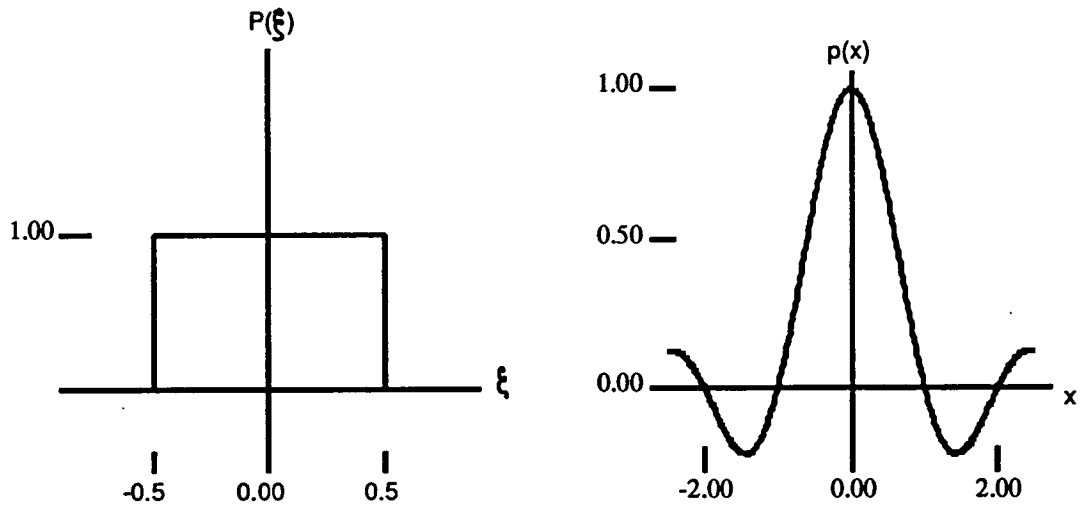

Figure 4.4: a) $p(\xi)$ and b) $P(x)$ for rectangular pixels.

We can label these peaks as *spectral orders* with the zero order at $x = 0$.

- $P(x)$, being the transform of the 'narrow' pixel function, will give rise to a 'wide' envelope function. To quantify this, as $p(\xi)$ is more than Q times narrower than than $\sum_{n=0}^{Q-1} \delta(\xi - n\alpha)$ then $P(x)$ will be more than Q times wider than $\frac{\sin(\pi Q \alpha x)}{\sin(\pi \alpha x)}$, width in the transform plane being defined approximately as the distance from the origin to the first zero of the transforms.

  Looking at the specific case of rectangular pixels, which can be represented in 1-d by a rect function (figure 4.4a),

$$p(\xi) = \text{rect}\left(\frac{\xi}{\beta}\right) = \begin{cases} 1 & -\beta/2 \leq \xi \leq \beta/2 \\ 0 & \text{otherwise} \end{cases}$$

  then

$$P(x) = \beta \frac{\sin(\pi\beta x)}{\pi\beta x} = \beta \text{sinc}(\beta x)$$

  which is plotted for $\beta = 1$ in figure 4.4b.

Treatment 2

Again we Fourier transform $T_2(\xi) = p(\xi) * \sum_{n=0}^{Q-1} \delta(\xi - n\alpha)$. However we write the finite sum of δ-functions as a product of an infinite sum of δ-functions (the 1-d

designed by the Applied Optics Group at Edinburgh University.

comb function in Gaskill's notation) and an offset rectangular window function corresponding to the size and position of the array. Thus (after Gaskill [7, pp 60-63]),

$$\begin{aligned}\sum_{n=0}^{Q-1}\delta(\xi-n\alpha) &= \sum_{n=-\infty}^{\infty}\delta(\xi-n\alpha)\times\text{rect}\left(\frac{\xi-\frac{(Q-1)\alpha}{2}}{Q\alpha}\right)\\ &= \frac{1}{\alpha}\,\text{comb}\left(\frac{\xi}{\alpha}\right)\times\text{rect}\left(\frac{\xi-\frac{(Q-1)\alpha}{2}}{Q\alpha}\right)\end{aligned}\tag{4.30}$$

Fourier transforming both sides and using the convolution and shift theorems,

$$\begin{aligned}\sum_{n=0}^{Q-1}e^{-2\pi i n\alpha x} &= \text{comb}(\alpha x) * Q\alpha\,\text{sinc}(Q\alpha x)\,e^{-i\pi(Q-1)\alpha x}\\ &= \frac{1}{\alpha}\sum_{n=-\infty}^{\infty}\delta(x-\frac{n}{\alpha}) * Q\alpha\,\text{sinc}(Q\alpha x)\,e^{-i\pi(Q-1)\alpha x}\end{aligned}\tag{4.31}$$

Comparing with treatment 1, this shows more explicitly how the phase factor arises as a result of the unsymetrically defined array. More importantly it allows us to note that $\frac{\sin(Q\pi\alpha x)}{\sin(\pi\alpha x)} = Q\,\text{sinc}(Q\alpha x) * \sum_{n=-\infty}^{\infty}\delta(x-\frac{n}{\alpha})$, and, disregarding the phase factor, write

$$t_2(x) = P(x)\times\left(Q\,\text{sinc}(Q\alpha x) * \sum_{n=-\infty}^{\infty}\delta(x-\frac{n}{\alpha})\right)\tag{4.32}$$

Thus $t_2(x)$ can be considered as a series of δ-functions convolved with narrow sinc functions (from the transform of the array's window function) serving to broaden the peaks, all modulated by a wide envelope function (from the transform of the pixel function). Again, this shows more explicitly the origin of the spectral orders.

Sampling and Replication

Returning to equation 4.32, it is apparent that as $t_2(x)$ will be convolved with the coordinate reversed object, multiple images will be formed at plane 3 due to the broadened δ-function peaks. The magnitudes of the higher order peaks are

attenuated by $P(x)$, the transform of the pixel function. Figure 4.5 shows the amplitude spectrum $t_2(x)$ for the case $Q = 16$, $\alpha = 2$, $\beta = 1$ with rectangular pixels. Figure 4.6 shows the intensity spectrum $|t_2(x)|^2$, from which the relative intensities can be deduced.

The presence of the lower intensity replicated images at the higher orders is not necessarily a problem– indeed it may be desirable to have multiple images to fan out information to other optical processors, though the differing intensities of the higher orders makes this less useful than a Damann grating, say. If the higher order images are not desired then they can be masked off at plane 3 if they do not overlap. If they do overlap, information from orders will leak into the adjacent orders, corrupting the information contained in them, and leading to *aliasing.*

From figures 4.5 and 4.6 it is apparent that if the replicas are more than 1 unit in width they will overlap with resulting aliasing, corrupting the zero order replica. For the unit magnification system described it is immediately apparent that this sets a maximum size of 1 unit for the object in plane 1. Note however that figures 4.5 and 4.6 illustrate a special case where the higher even orders are absent due to coinciding with zeroes of $P(x)$. This would suggest that if the third order replica was taken as the desired image, it would be possible to have an object up to 2 units in size without overlapping of the images. However, as the intensity of the third order is $\left(\frac{\sin(1.5\pi\beta)}{1.5\pi\beta}\right)^2 = 0.045$ times that of the zero order it is obviously inefficient to not use the zero order, and impractical if the output is to be cascaded.

The replication and aliasing occurs as a consequence of the Fourier transform of the object being sampled in the Fourier plane (plane 2) by the discrete pixels of the spatial light modulator, and can be understood with reference to the *sampling theorem* [5, p 371], [86, p 87], [8, p189], [7, pp 266-268].

The sampling theorem states that any band-limited function can be specified exactly by its sampled values, taken at regular intervals, provided these intervals do not exceed some *critical sampling interval*, the *Nyquist interval* [7, p267]. The

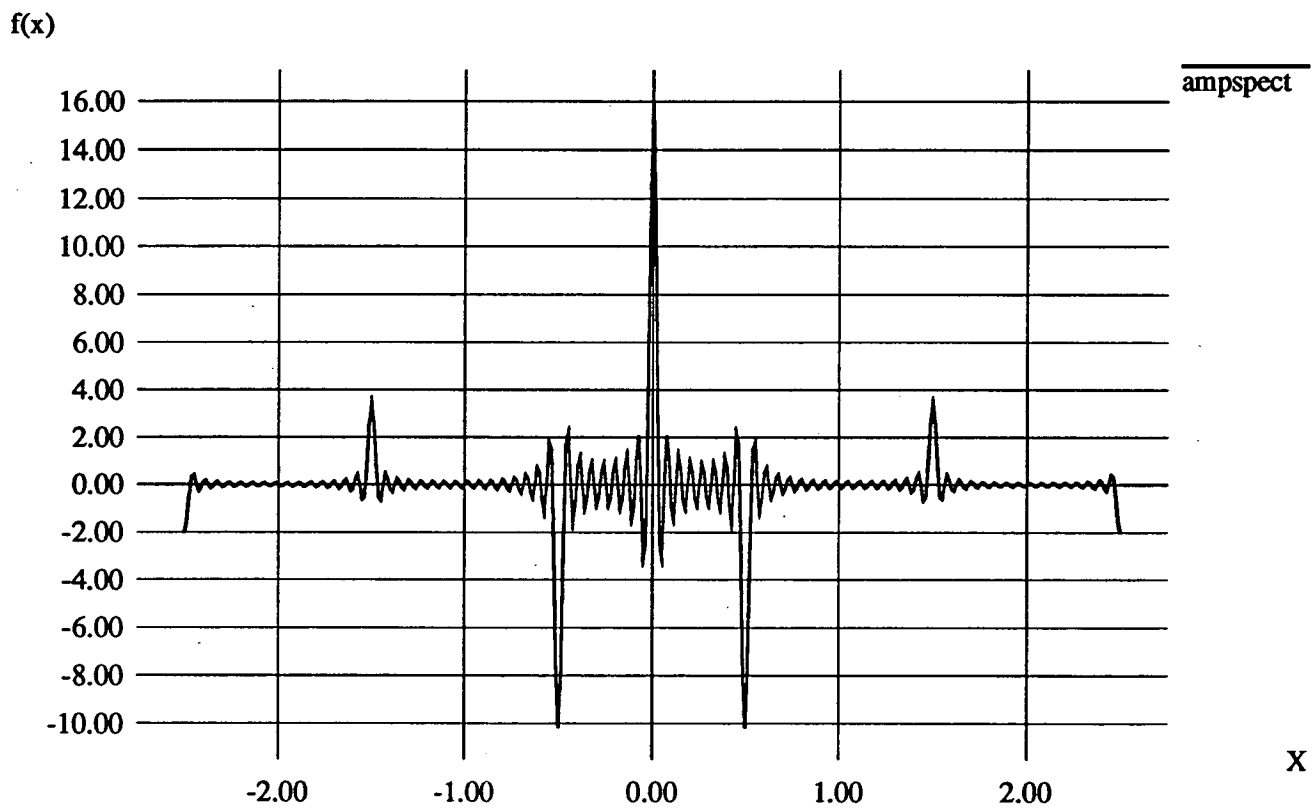

Figure 4.5: Amplitude spectrum $t_2(x)$ for $Q = 16$, $\alpha = 2$, $\beta = 1$.

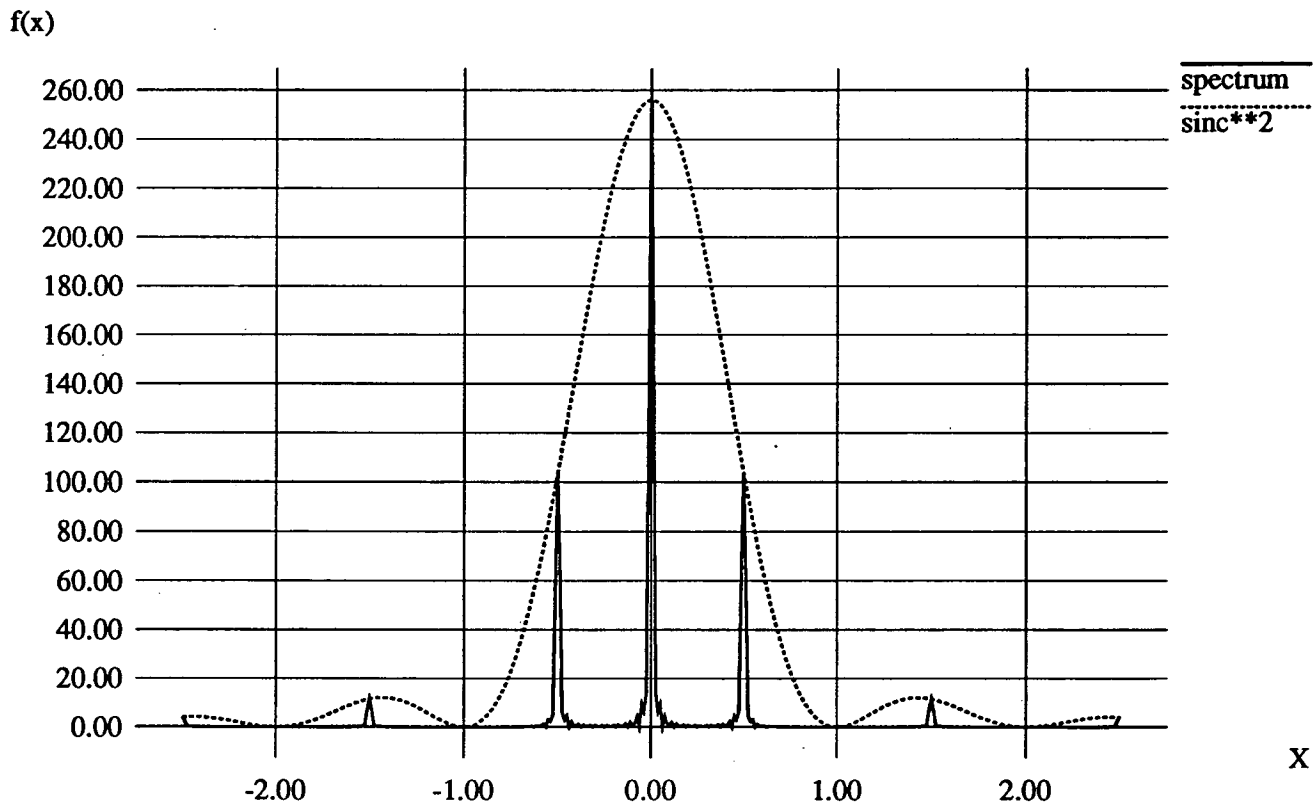

Figure 4.6: Intensity spectrum $|t_2(x)|^2$ for $Q = 16$, $\alpha = 2$, $\beta = 1$, under the dotted $Q^2\text{sinc}^2(\beta x)$ envelope.

sampling interval is the inverse of the sampling rate, so equivalently the sampled function needs to be sampled at a rate in excess of some *critical sampling rate*, the *Nyquist rate*, for it to be specified exactly by its sampled values.

The Fourier transform formed in plane 2 (with reference to figure 2.1) is sampled at a rate of $1/\alpha$ (in the same units as ξ, i.e. mm^{-1}) implying that the spectral width is α (mm) and that the spectral orders can be up to α in width before they overlap, with concomitant aliasing.

As the pixel pitch is fixed, this imposes $1/\alpha$ as the sampling rate, and α as the sampling interval. Hence the Fourier transform formed in plane 2 must be scaled such that α is its critical sampling interval i.e. so that the spectral orders are as close together as possible without overlapping. If the scale is larger than this *critical scale* the replicated images formed at plane 3 will be smaller because of the reciprocal nature of the Fourier transform, and of poorer quality due to less of the high spatial frequencies being passed by the spatial light modulator as a consequence of the increased scale of the transform. Conversely, if the scale of the transform is smaller more of the high spatial frequencies will be passed, but the replicated images formed at plane 3 will overlap rendering this gain of little use.

The scaling of the Fourier transform can be achieved either by varying the size of the object, or by varying the magnification of the transform. The second of these is achieved by altering the focal length of the first transforming lens, lens 1, and thus altering the scaling factor $1/\lambda f$ which relates ξ and η, the spatial frequency coordinates, to x_2 and y_2, the actual spatial coordinates in the Fourier plane. Thus, as $\xi = x_2/\lambda f$ and $\eta = y_2/\lambda f$, it is apparent that increasing the focal length will reduce the scale of the transform and decreasing the focal length will increase the scale of the transform.

With regard to the spatial light modulator it is apparent that some sort of trade-off is necessary– if the size of the Fourier transform is reduced, higher spatial freqencies are able to be transmitted, but with a reduction in the sampling rate meaning that the replicated images are closer together and may overlap; if the size of the Fourier transform is increased, there is a reduction in the maximum

spatial frequency able to be transmitted, but an increase in the sampling rate and greater spatial separation of the images in plane 3.

The optimum tradeoff is thus to scale the transform so that the pixel pitch, α, is the critical sampling interval. This sets the passband of the SLM as $Q\alpha$ where Q is the number of pixels across the device. It is easy to see that the product of the passband and the sampling rate (the *space-bandwidth product*) is a constant equal to Q. Thus, with reference to a device with Q_1 pixels, if we have a device with a higher number of pixels Q_2 we can scale the transform so the passband is the same and oversampling takes place with increased separation of the replicas, or scale so that critical sampling takes place and replicas are formed containing higher spatial frequencies. [3] In short– with respect to sampling, the more pixels the better.

The size of the pixels has been shown to govern the magnitude of the replicated spectral orders. It is also apparent that it is possible to remove some of the higher orders if the pixel pitch (α) is an integral multiple of the pixel width (β). In figure 4.5 α:β is 2:1 and the 2nd, 4th, 6th, 8th ... orders are removed. Similarly if α:β was 3:1 the 3rd, 6th, 9th ... orders would be missing. If α:β was 1:1 only the zero order would remain, and aliasing and replication would not take place. Here the SLM corresponds to a single large aperture equal in size to the overall size of the array, when all the pixels are transmitting. In practice this is not possible with the active backplane or matrix addressed devices described, although it is desirable to maximise the size of the pixels to attenuate the first order replicas as much as possible, and improve the diffraction efficiency of the filter.

4.1.2 Generalisation to 2-d

The generalisation of the analysis to 2-d for the filter with regularly spaced square pixels is straightforward, as the positions and sizes of the pixels in the ξ and

[3] These replicas are also convolved with a narrower $Q\,\mathrm{sinc}(Q\alpha x)$ function which contributes to a higher spatial frequency content.

η directions are independent of η and ξ respectively. Thus, substituting using $p(\xi,\eta) = p(\xi)\ p(\eta)$ and $\delta(\xi - n\alpha, \eta - m\gamma) = \delta(\xi - n\alpha)\ \delta(\eta - m\gamma)$ we see that equation 4.28 is easily separable into a product of two functions dependent only on ξ and η respectively, the transform of which is the product of their individual transforms [7, p 292], [87, p 45], i.e.

$$T_2(\xi,\eta) = \left[p(\xi) * \sum_{n=0}^{Q-1} \delta(\xi - n\alpha)\right]\left[p(\eta) * \sum_{m=0}^{Q-1} \delta(\eta - m\gamma)\right]. \tag{4.33}$$

Thus $t_2(x,y)$ is a product of two terms similar to $t_2(x)$.

Giving, for **Treatment 1**

$$t_2(x,y) = P(x,y)\frac{\sin(\pi Q\alpha x)}{\sin(\pi\alpha x)}\frac{\sin(\pi Q\gamma y)}{\sin(\pi\gamma y)} \tag{4.34}$$

and for **Treatment 2**

$$t_2(x,y) = P(x,y)\left(Q^2\mathrm{sinc}(Q\alpha x)\mathrm{sinc}(Q\gamma y) * \sum_{n=-\infty}^{\infty}\sum_{m=-\infty}^{\infty} \delta(x - \frac{n}{\alpha}, y - \frac{m}{\gamma})\right) \tag{4.35}$$

where $P(x,y) = P(x)P(y)$, $\delta(x - \frac{n}{\alpha}, y - \frac{m}{\gamma}) = \delta(x - \frac{n}{\alpha})\delta(y - \frac{m}{\gamma})$, and the phase factors due to offsetting of the array have been omitted.

For an array with square pixels of size β, $p(\xi,\eta) = \mathrm{rect}(\xi/\beta, \eta/\beta)$ and

$$P(x,y) = \beta^2\mathrm{sinc}(\beta x, \beta y) = \beta^2\mathrm{sinc}(\beta x)\mathrm{sinc}(\beta y) \tag{4.36}$$

Note that the magnitude of the zero order is then $t_2(0,0) = Q^2\beta^2$, equivalent to the transmissive area of the array, which satisfies the requirement that the definite

integral from $-\infty$ to ∞ should equal the value of its transform at the origin [8, p136], i.e. that if $F(\xi)$ and $f(x)$ are a Fourier transform pair, then

$$f(0) = \int_{-\infty}^{\infty} F(\xi)\, d\xi \tag{4.37}$$

This pixel function is used for the following simulations, giving the power spectrum of the corresponding regularly pixellated 2-d array as

$$|t_2(x,y)|^2 = \beta^4 \mathrm{sinc}^2(\beta x, \beta y)\, \frac{\sin^2(\pi Q \alpha x)}{\sin^2(\pi \alpha x)} \frac{\sin^2(\pi Q \gamma y)}{\sin^2(\pi \gamma y)}. \tag{4.38}$$

4.1.3 Simulations for a regularly pixellated 2-d array

Simulations were carried out to show how a regularly pixellated filter in the Fourier plane of an optical processor affects the output, for inputs at various scales. The simulations were carried out for a 16x16 array of regularly spaced square pixels, with α and β in the ratio 2:1.

Computational simulations

Computational simulations were carried out using 256×256 arrays ofnumbers representing the object, filter, and image planes, with the Fourier transformation between them carried out using an FFT routine. Objects were generated in the form of letter 'A' in a number of different sizes, and a letter 'E' just less than the critical scale(figure 4.8). The regular filter and the input objects are represented graphically in figures 4.7 and 4.8.

Figure 4.9 shows the computational simulation of $|t_2(x,y)|^2$, along with simulations of $|u_3^-(x,y)|^2$, the output of the optical processor, in which letters of a set size have been used as the input object. The image display package xv [88] was

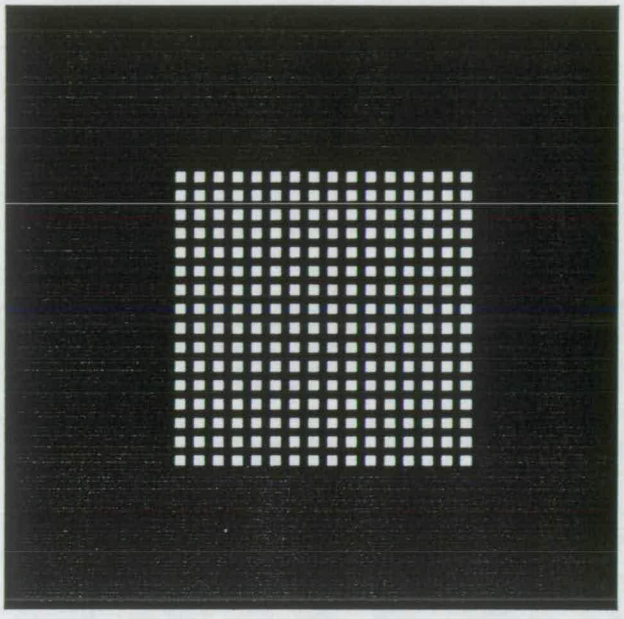

Figure 4.7: 16x16 regularly pixellated square array: $\alpha = 2, \beta = 1$.

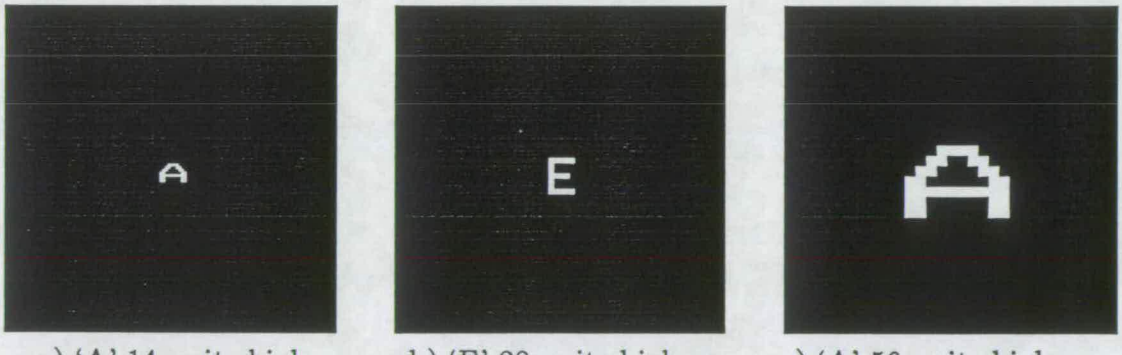

a) 'A' 14 units high b) 'E' 28 units high c) 'A' 56 units high

Figure 4.8: Form of letters 'A' and 'E' used as objects.

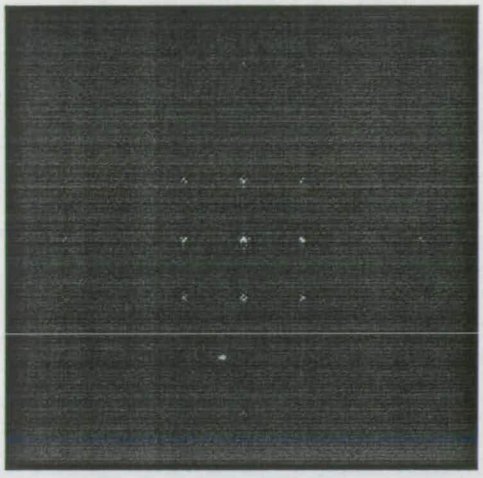

a) Power spectrum of regular array used as filter

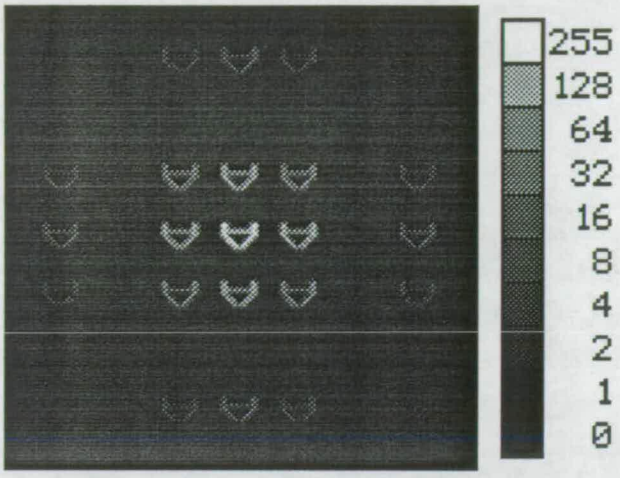

b) Output image illustrating over-sampling

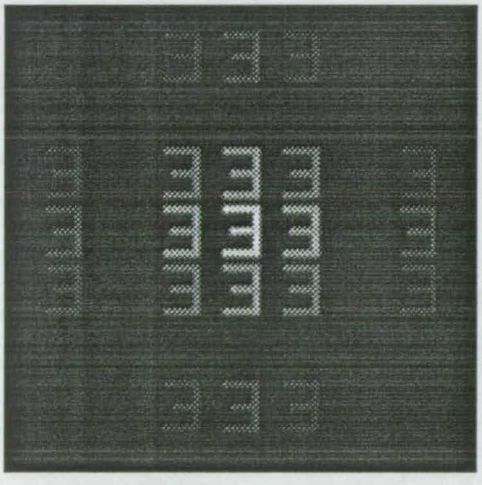

c) Output image illustrating critical sampling

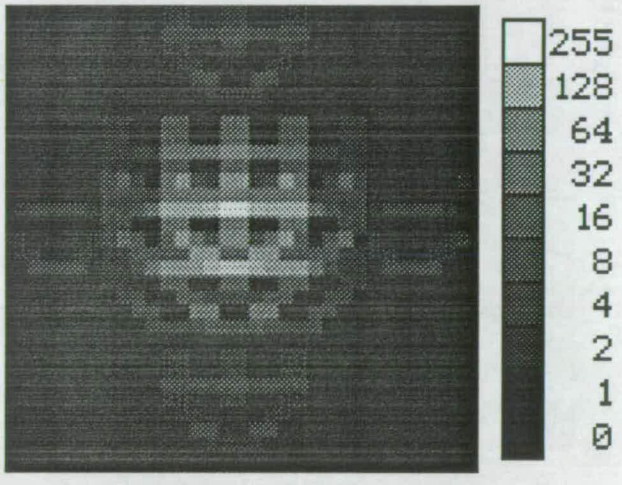

d) Output image illustrating under-sampling and aliasing

Figure 4.9: Computational simulations for regularly pixellated filter.

used to generate the figures, with the gamma value of each image set at 1.5. The effect of this is to emphasise the lower intensities and prevent loss of detail when the figures are reproduced.

The simulations illustrate over-, under-, and critical sampling. Figure 4.10 also shows $|t_2(x,y)|^2$, and takes a cross section along the x-axis which can be compared with figure 4.2 from the earlier 1-d analysis. Note the relative intensities of the

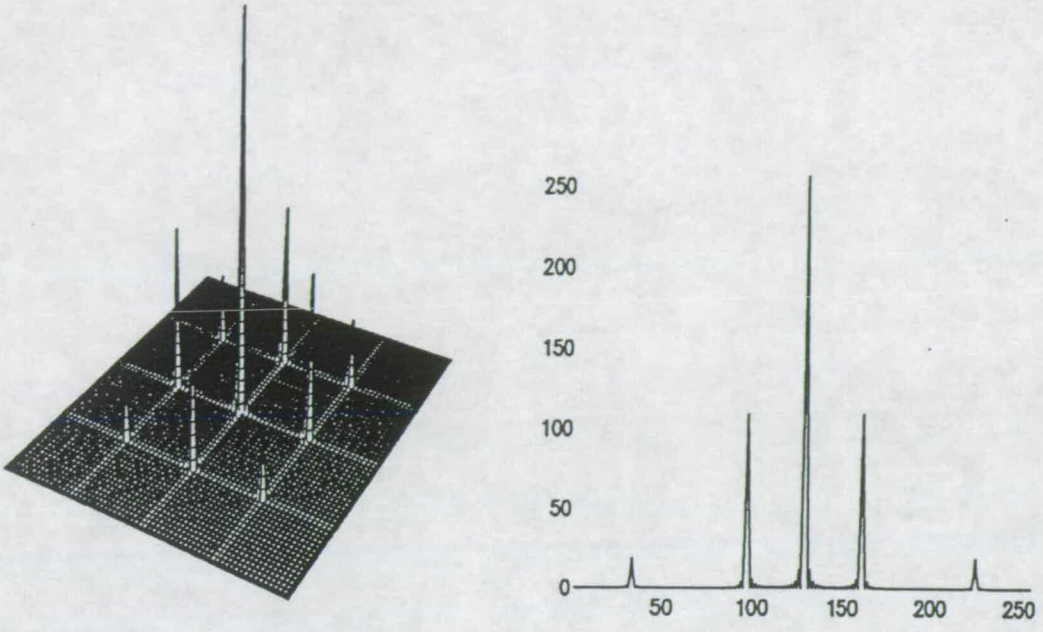

a) Power spectrum, $|t_2(x,y)|^2$.

b) Cross section along the x-axis of a)

Figure 4.10: Further detail of the power spectrum,$|t_2(x,y)|^2$ of the regular array

spectral orders, and in particular the absence of the second spectral orders.

In the simulation the filter was scaled such that it was confined to the central 128x128 points of the numerical array. This is to prevent the filter from appearing infinite in extent due to the cyclical and sampling nature of the FFT– this would lead to the spectral orders appearing as δ-functions instead of being broadened by convolution with the transform of the filter's window function, as described in Treatment 2. Thinking of this in another way, if the filter filled the entire 256x256 points of the numerical array, i.e. $Q\alpha = 256$, then the sinc function from the Fourier transform of this window would have its zeroes at the sampling points of the array representing the transform, giving apparent δ-function behaviour. However, if the window is half the size, the zeroes will be spaced every two sampling points and the effect of the convolution will be evident.

In these simulations the critical size of the object is 32x32 points: the over-, under- and critical sampling cases used letters 14, 56 and 28 points high respectively. Thus the case decribed as critical sampling should perhaps be more rigourously described as 'near-critical sampling'.

Optical simulations

High quality hardcopies of the graphical representation of the filter were obtained from a laser printer, and reduced photographically onto lithographic film at the appropriate scale, to provide high contrast filters for the optical simulations. This was repeated for the input letters.

The optical simulations using the transparencies are shown in figure 4.11, made on a Fourier optical processing bench with 0.31m focal length transforming lenses, of diameter 5cm.

The transparencies were mounted, and placed in a liquid gate containing Dow Corning DC704 silicone ferrofluid (refractive index at room temperature = 1.557) in an attempt to reduce the effect of surface roughness of the transparency by index matching. The liquid gate was constructed for the experiment using Ealing Beck 'utility-grade' optical flats. A photograph of the liquid gate is shown in figure 4.12.

The photographs in figure 4.11 were taken on Ilford HP5 film, and overexposed by approximately two stops with development reduced by a third. This reduces contrast and makes the detail in the low intensity regions easier to see.

4.1.4 Discussion of results

The computational simulations are seen to be in agreement with the theory by comparison of the cross-section along the x-axis of the power spectra and figure 4.6.

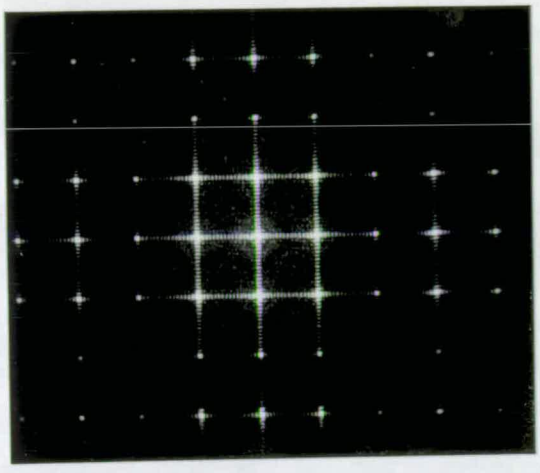

a) Power spectrum of regular array used as filter

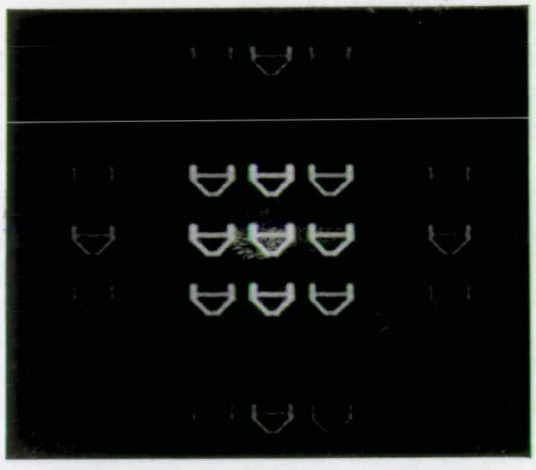

b) Output image illustrating oversampling

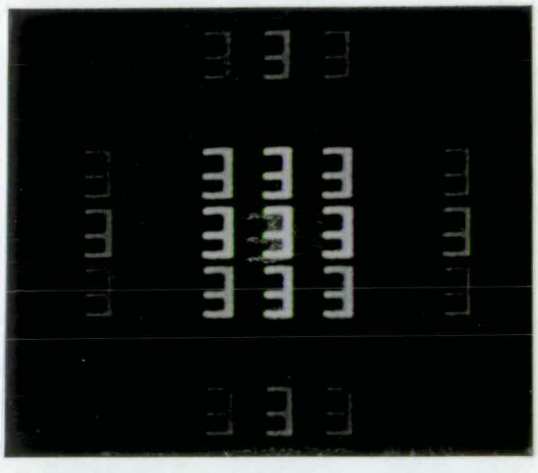

c) Output image illustrating critical sampling

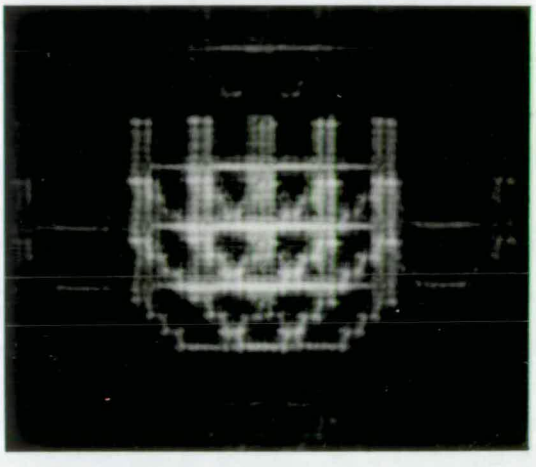

d) Output image illustrating undersampling and aliasing

Figure 4.11: Optical simulations for regularly pixellated filter.

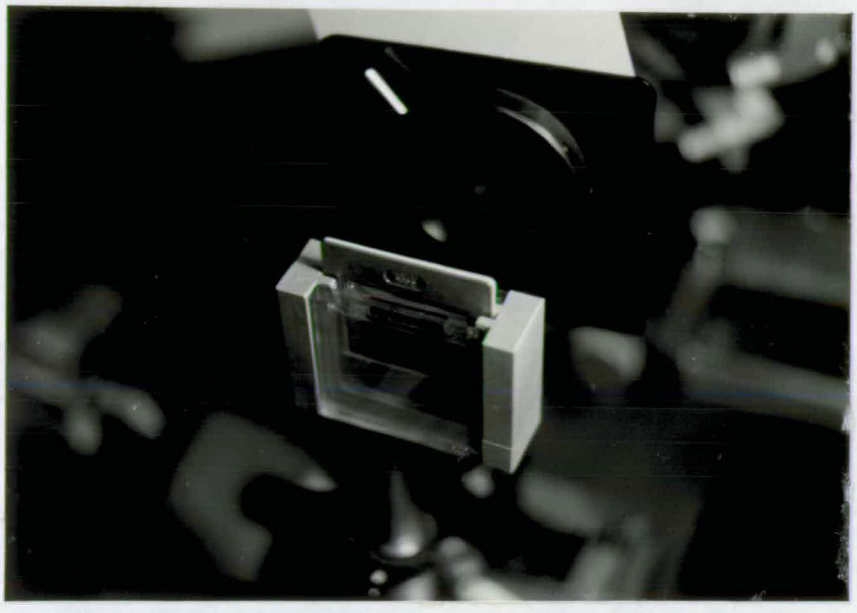

Figure 4.12: Liquid gate mounted on 4-f bench

The relative heights of the spectral orders are also reflected in the relative intensities of the replicated images in figure 4.9.

The optical simulations are much more important, as the theory must be seen to model the real life situation effectively. Qualitatively this seems to be the case, simply from comparing figures 4.9 and 4.11, though the optical simulations differ from the theory in that the higher even orders are present, albeit extremely weakly. This may be attributed to a small discrepancy in the size of the pixels, which leads to the zeroes of the transform of the pixel function falling close to, but not exactly on, the centres of the higher even orders, spatial non-uniformity of the laser illumination, or finite transmittance through the black areas of the transparency. The effect is extremely small, and only shows up through overexposure of the zero and first orders.

A more quantitative relation of the theory and simulation may be derived from measuring the size of the objects which lead to critical sampling, and by measuring the intensities of the various spectral orders in the filters diffraction pattern.

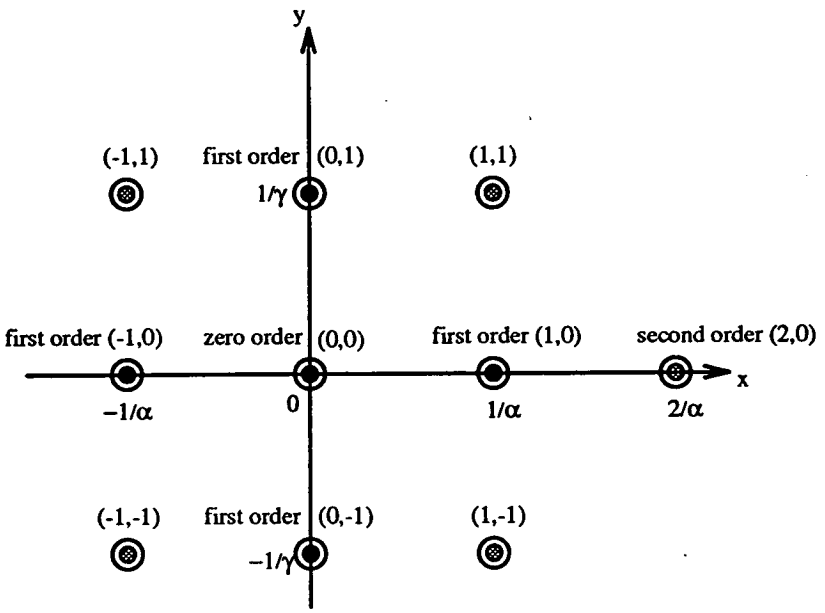

Figure 4.13: Nomenclature for labelling orders

Object sizes

The optical simulations were carried out on a processing bench employing (0.31±0.01)m focal length lenses, using a 16x16 array with pixels on a 200μm pitch, illuminated by a HeNe laser operating at 632.8nm. As described earlier, this gives a critical size for the object of $\frac{\lambda f}{\alpha} = \frac{632.8\times10^{-9}\times0.31}{200\times10^{-6}}$m $= (0.98 \pm 0.03)$mm, in both the x and y directions. This tallies with the results as the oversampled A is less than this, being (0.62 ± 0.02)mm high. the near-critically sampled E slightly smaller at (0.88 ± 0.02)mm, and the undersampled A 1.26 times larger at (1.24 ± 0.02)mm.

Magnitudes of spectral orders

Before comparing the magnitudes of the spectral orders it is useful to establish a simple nomenclature by which they can be labelled in 2-d. This is described in figure 4.13 and can be seen to be related to scaled x-y coordinates, with the zero order labelled as the $(0,0)^{th}$ order at (0,0), and the first order lying on the x-axis at $(1/\alpha, 0)$ labelled as the $(1,0)^{th}$ order. Substituting the position of the $(n,m)^{th}$

Order	Calculated relative $I_{n,m}$	Measured $I_{n,m}$	Measured relative $I_{n,m}$
(0,0)	1.00	144± 4	1.00
(1,0)(0,1)(-1,0)(0,-1)	$(\frac{2}{\pi})^2 = 0.41$	66±4	0.46±0.04
(1,1)(1,-1)(-1,-1)(-1,1)	$(\frac{2}{\pi})^4 = 0.16$	29±4	0.20± 0.04

Table 4.1: Calculated and measured relative intensities of spectral orders in the power spectrum of a regular array

order (i.e. $(n/\alpha, m/\gamma)$) into $|P(x,y)|^2$ gives its intensity directly. Labelling this intensity as $I_{n,m}$ and using $\beta = \alpha/2$, $P(x,y) = \beta^2 \mathrm{sinc}^2(\beta x, \beta y)$ gives

$$I_{n,m} = (Q\beta)^4 \left(\frac{\sin(\frac{n\pi}{2})}{\frac{n\pi}{2}} \right)^2 \left(\frac{\sin(\frac{m\pi}{2})}{\frac{m\pi}{2}} \right)^2 . \tag{4.39}$$

Some values for $I_{n,m}$ are calculated and shown in table 4.1, with the $(Q\beta)^4$ multiplier omitted to normalise the intensities relative to the zero order.

The third and fourth columns in the table were completed using data gathered from a 2-d image, collected by an Avcam 405 CCD array camera placed in the transform plane of the uniformly illuminated transparency. This image was grabbed and digitised to 256 grey levels using a PIP 512 framestore mounted on an IBM PS/2 personal computer. The camera control unit had its gamma set at unity to help ensure linearity, and care was taken to attenuate the illumination to prevent saturation occuring at the zero order. The measured intensity values have had the background intensity of 3 subtracted from them, and the second and third entries are the averages over the orders which have the same calculated intensity. The uncertainty estimates are taken from the averaging process. The background intensity is the mean of the intensity values from a 47x96 pixel area of the 512x512 pixel digitised image, well away from the spectral orders. The faint even orders noticed in the photographs were indistinguishable from the background, confirming overexposure. The linearity of the camera was established experimentally using Malus' Law, and the calibration graph is shown in Appendix A.

The results are in good agreement as should be expected for this simple case. The

main objective here was to establish the reliability of the intensity measurements, as the same measurement techniques are used to establish the validity of the randomised pixel cases discussed later.

4.2 Conclusions

To summarise the analysis of the effect of a regularly pixellated amplitude filter in the Fourier plane we note the following points:

- Spectral orders arise in the image plane due to the regularly spaced pixels on the spatial light modulator.
- The function which describes the pixel size and shape helps determine the magnitude of the spectral orders, and can be used to eliminate some of them.
- The spectral orders have a width inversely proportional to the number of pixels across the SLM.
- The presence of the spectral orders leads to replication in the image plane, and aliasing if the Fourier transform is undersampled.
- The product of the passband and the sampling rate is a constant, equal to the number of pixels across the device, and shows the benefit of devices with a large number of pixels.
- The optimum sampling rate is obtained when the replicated images are as large as possible without overlapping, and this fixes the scale of the transform.
- The transform scale is determined by the object size and the focal length of the transforming lens– a transforming lens of focal length f_1 sets an object size of $\lambda f_1/\alpha$ by $\lambda f_1/\gamma$ for optimum sampling of the 2-d case above.
- The nth spectral order can be removed if the ratio of pixel size:pixel pitch is k/n, where k is an integer $\leq n$.

- Making the pixels as large as possible minimises the size of the regular array's first order replicas and maximises the diffraction efficiency of the SLM.

The spectral orders arise due to the regular spacing of the pixels. In the next chapter a random element is introduced to the positioning of the pixels in an attempt to suppress the spectral orders.

Chapter 5

Use of a randomly pixellated SLM in a 4-f optical processor

The idea that randomly positioning the pixels will reduce the size of the higher spectral orders is intuitively appealling, as the spacing of the spectral orders is directly related to the spacing of the pixels. If the spacing of the pixels could be made completely random, it would seem reasonable to expect that a continuous spectrum of spatial frequencies would be necessary in the Fourier transform plane to represent this, thus removing the distinct spectral orders.

The following analysis shows that it is indeed possible to reduce the size of the spectral orders by randomising the pixel positions. However, for SLMs the matter is complicated by restrictions placed on the positioning of the pixels to prevent pixels overlapping and allow practical fabrication of devices. These are included in the analysis, and when randomisation schemes are considered a special case is described which gives results which belie its underlying regularity. First, however, we note the practical motivation in removing the higher spectral orders.

5.1 Motivation to remove the higher spectral orders

We now consider the consequences of the replication on optical correlation experiments, where the objective is to detect the presence of a target object in an input scene. Whether or not the replicas constitute a problem will depend on the spectral content of the input scene, the nature of the target contained within, and the resolution of the filter SLM. In many applications the input to the processor will itself be implemented using an SLM. This is assumed to be the case here, with the input SLM labelled as SLM1 and the filter SLM as SLM2. It is convenient to characterise each SLM by its space-bandwidth product (SBP), equal to the number of pixels across the SLM, Q.

As noted earlier, to prevent the replicas overlapping (i.e. aliasing), the Fourier transform must be sampled at frequency greater than or equal to the Nyquist rate. This condition imposes a maximum size on the object– if the pixel pitch of SLM2 is α, the maximum size of the object will be $\lambda f/\alpha$. If this object then occupies Q_1xQ_1 pixels on SLM1 then the lowest spatial period that can be represented is $(2/Q_1)(\lambda f/\alpha)$. This in turn will transform to a point $Q_1\alpha/2$ distant from the position of the zero spatial frequency component of the transform, which passes through the centre of the central pixel of SLM2, requiring that SLM2's SBP be at least as large as that of the target in order to represent accurately its transform or complex conjugate. In the case where the SBPs of SLM1 and SLM2 are equal (necessary if the target occupies all of SLM1), then the replication of the output will not be a problem as the replicas will fall outwith the region of interest in plane 3, which will be restricted to the area defined by the geometrical image of the perimeter of SLM1. If however the target or targets each occupy a 16×16 pixel region in a 256x256 input SLM for example, then a 16×16 pixel SLM in plane 2, though adequate to represent the target's Fourier transform, will produce replicas lying within the region of interest in the output plane [89, 90] which could give rise to peaks which may be misinterpreted as signifying correlation. Increasing the SBP of SLM2 so that it was a 256x256 pixel device would solve the problem, but the advantages to be gained from using a low-resolution SLM in the Fourier

plane (fast frame rates, compact size, modest memory size for the filter bank [90]) suggest that a worthwhile alternative approach to the problem is to try to eliminate or at least reduce the magnitudes of the spectral orders in $t_2(x, y)$ which lead to the replicas.

5.2 Randomly Pixellated Amplitude Filter

5.2.1 Analysis in 1-d

The method used here is consistent with that of Treatment 1 for the regularly pixellated amplitude filter described previously. Again, the diffraction pattern of the filter itself with all pixels 'on' is calculated. The general analysis is partly derived from work by Stark [91, 92] and Aime [93, 94] on diffraction patterns of nonoverlapping circular grains, and the power spectrum of solar granulation, respectively.

Using the same coordinate system as before we postulate a 1-d filter function whose amplitude transmittance is

$$T_2(\xi) = p(\xi) * \sum_{n=0}^{Q-1} \delta(\xi - n\alpha - \phi_n) \tag{5.40}$$

where as before $p(\xi)$ is the pixel function and Q is the number of pixels. ϕ_n is a random displacement of the pixel centre from the position it would have occupied were it regularly positioned with pitch α. Q and n are integers.

Fourier transforming and using the convolution and shift theorems gives

$$t_2(x) = P(x) \times \sum_{n=0}^{Q-1} e^{-i2\pi n\alpha x} e^{-i2\pi\phi_n x} \tag{5.41}$$

using the same notation as before.

To solve this equation analytically requires knowledge of all the ϕ_n values, and will not result in a useful general expression. However, noting that the equation represents just one of an ensemble of possible pixel arrays, dependent on the choice of the ϕ_n, we can use an ensemble average to replace the part of the equation involving the random values, which will yield an expression showing the form of the diffraction pattern. Since we detect intensities, the ensemble averaging is performed over $|t_2(x)|^2$.

Thus we wish to evaluate

$$\langle |t_2(x)|^2 \rangle = \langle |P(x)|^2 \times \sum_{n=0}^{Q-1} \sum_{n'=0}^{Q-1} e^{-i2\pi n\alpha x} e^{i2\pi n'\alpha x} e^{-i2\pi\phi_n x} e^{i2\pi\phi_{n'} x} \rangle \tag{5.42}$$

where the angle brackets denote the ensemble averaging operation.

By removing the $\phi_n, \phi_{n'}$ terms which are dependent on each other (i.e. when $n = n'$) from the summation

$$\langle |t_2(x)|^2 \rangle = |P(x)|^2 \times \left[Q + \sum_{\substack{n=0 \\ n \neq n'}}^{Q-1} \sum_{n'=0}^{Q-1} e^{-i2\pi n\alpha x} e^{i2\pi n'\alpha x} \langle e^{-i2\pi\phi_n x} \rangle \langle e^{i2\pi\phi_{n'} x} \rangle \right]. \tag{5.43}$$

where the Q arises from the $n = n'$ terms.

This can be readily evaluated in general, as

$$\langle e^{-i2\pi\phi_n x} \rangle = \int_{-\infty}^{\infty} P_{\phi_n}(\phi) e^{-i2\pi\phi x} \, d\phi \tag{5.44}$$

where $P_{\phi_n}(\phi)$ is the real, positive probability of finding ϕ_n with value ϕ, and is equal to $P_{\phi_{n'}}(\phi)$. Using this, and taking $P_{\phi_n}(\phi)$ to be even, allows us to write

$$\langle e^{i2\pi\phi_{n'}x}\rangle = \int_{-\infty}^{\infty} P_{\phi_n}(\phi)e^{i2\pi\phi x}\,d\phi \tag{5.45}$$

$$= \left[\int_{-\infty}^{\infty} P_{\phi_n}(\phi)e^{-i2\pi\phi x}\,d\phi\right]^* \tag{5.46}$$

giving

$$\langle e^{-i2\pi\phi_n x}\rangle\langle e^{i2\pi\phi_{n'}x}\rangle = |\mathcal{F}[P_{\phi_n}(\phi)]|^2 \tag{5.47}$$

$$= |\rho(x)|^2 \tag{5.48}$$

where $\rho(x)$ is the Fourier transform of the probability distribution function, $P_{\phi_n}(\phi)$.

The remaining geometrical series terms work out to be

$$\underbrace{\sum_{n=0}^{Q-1} e^{-i2\pi n\alpha x} \sum_{n'=0}^{Q-1}}_{n\neq n'} e^{i2\pi n'\alpha x} = \frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)} - Q \tag{5.49}$$

Substituting these results into equation 5.42 gives the following important result

$$\langle|t_2(x)|^2\rangle = |P(x)|^2 \times \left[Q + \left(\frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)} - Q\right)|\rho(x)|^2\right] \tag{5.50}$$

$$= |P(x)|^2 \times Q\left[1 - |\rho(x)|^2 + \frac{|\rho(x)|^2}{Q}\frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)}\right] \tag{5.51}$$

This expression is clearly similar to that derived for the regularly pixellated case, and indeed is the same if $\rho(x)$ is identically 1. This corresponds to the probability

distribution function $P_{\phi_n}(\phi)$ being equal to $\delta(0)$, or the ϕ_n being identically zero, as would be expected. Other specific choices of probability distribution function will be considered when the expression is generalised to 2 dimensions, but it is useful at this stage to make some general observations regarding $\rho(x)$.

1. Applying the theorem quoted as equation 4.37 shows that if $P_{\phi_n}(\phi)$ is properly normalised (i.e. $\int_{-\infty}^{\infty} P_{\phi_n}(\phi)\ d\phi = 1$), then $\rho(0)$ will be equal to 1, regardless of the form of $P_{\phi_n}(\phi)$.

2. As $0 \leq P_{\phi_n}(\phi) \leq 1$, an upper bound of 1 can be established for $\rho(x)$ using the following general result for Fourier transform pairs [8, p157]:

$$|\rho(x)| \leq \int_{-\infty}^{\infty} |P_{\phi_n}(\phi)|\ d\phi \tag{5.52}$$

3. Combining the previous two points shows that $|\rho(x)|^2$ is a maximum at $x = 0$ and must either decrease or remain equal to 1 as as we move away from the origin. An instance of the latter case is that of the regular array described above. The special cases of probability distribution function considered later will give rise to $\rho(x)$'s which will decrease away from the origin.

These points help show quantitatively how the randomisation will affect the form of $\langle |t_2(x)|^2 \rangle$, the power spectrum of $T_2(\xi)$, and how the randomisation may be used to suppress the higher order spectral order peaks which give rise to well defined replicas.

It is apparent that, with the exception of the zero order, the spectral order peaks resulting from the $\frac{\sin^2(\pi Q \alpha x)}{\sin^2(\pi \alpha x)}$ term are attenuated by the $|\rho(x)|^2$ multiplier. This suggests that with a suitable choice of probability distribution function, $|\rho(x)|^2$ may be tailored to have zeroes at the centres of some or all of the spectral orders, thus eliminating them. As the spectral orders have finite width this is an approximation, but even for a small number of pixels it is still a good one, as will be demonstrated later. Thus some of the spectral orders may be eliminated, and the remaining orders will be no larger than for the corresponding regularly pixellated case.

The magnitude of the zero order is proportional to the total transmissive area of the array, and is thus independent of the chosen randomisation scheme so long as the pixels do not overlap and reduce the total transmissive area[1]. As overlapping pixels are neither practical nor particularly desirable in an SLM, all the cases considered will constrain the pixels from overlapping.

The $(1 - |\rho(x)|^2)$ part of the equation gives rise to a *light depletion region* [91]. This refers to a dark patch surrounding the zero order (where $1 - |\rho(0)|^2 = 0$) in the midst of illumination arising from where $1 - |\rho(x)|^2$ has non-zero values. This background illumination is a redistribution of the energy removed from the higher spectral orders by $|\rho(x)|^2$, and must be conserved (expressed here by Rayleigh's theorem: $\int_{-\infty}^{\infty} |\rho(x)|^2 \, dx = \int_{-\infty}^{\infty} |P_{\phi_n}(\phi)|^2 \, d\phi$).

Little more can be said about the light depletion zone without choosing specific examples. However, if it is in the form of diffuse, low-level background illumination it will not contribute to the formation of distinct replicas when $t_2(x)$ is convolved with $t_1(-x)$ to give the output of the optical processor. This is beneficial when correlations are carried out, as distinct replicas can, through aliasing, lead to incorrect results. The extremely blurred replicas arising from a diffuse background can hopefully be easily separated from the actual correlation peak(s) by thresholding the output of the processor. This will be demonstrated in chapter 7.

It is also important to note that the ratio of the background and zero order intensities is

$$\frac{\langle \text{background} \rangle}{\langle |t_2(0,0)|^2 \rangle} = \frac{Q\,[1 - |\rho(x)|^2]}{Q\,[\frac{|\rho(0)|^2}{Q} \frac{\sin^2(\pi Q \alpha . 0)}{\sin^2(\pi \alpha . 0)}]} \leq \frac{1}{Q} \tag{5.53}$$

emphasising the advantage of having a large number of pixels, Q.

With the basic ideas that arise from the 1-d treatment established, it is now

[1]this follows from the general relation for Fourier transform pairs that $f(0) = \int_{-\infty}^{\infty} F(\xi) \, d\xi$, and from $\rho(0) = 1$

possible to transfer them to the more useful 2 dimensional case. The generalisation to 2-d is less straightforward for randomised pixels than it is for regularly spaced pixels.

5.2.2 Generalisation to 2-d

For the regular case the generalisation was straightforward, due to the separability of the Fourier transforms in the x and y directions. This is not the case with a randomly pixellated array, as can be seen when the transmission function of the 2-d filter is written down.

$$T_2(\xi,\eta) = p(\xi,\eta) * \sum_{n=0}^{Q-1}\sum_{m=0}^{Q-1} \delta(\xi - n\alpha - \phi_{n,m}, \eta - m\gamma - \psi_{n,m}) \tag{5.54}$$

The coordinate system and variables are as defined earlier except for the introduction of the random displacements $\phi_{n,m}$ and $\psi_{n,m}$. Here $\phi_{n,m}$ and $\psi_{n,m}$ are $\xi-$ and $\eta-$ components of the displacement of the centre of the pixel labelled by n and m from the position it would have occupied were it regularly positioned with pitch α in the ξ direction and pitch γ in the η direction. Q, n and m are integers.

The problem here is that the random variables' displacements are specified for each pixel in the QxQ array by both labels n and m, which are counters labelling the pixels in the ξ and η directions respectively. Thus neither the transmission function of the filter nor its Fourier transform is easily separable. Therefore $T_2(\xi,\eta)$ must be Fourier transformed directly to give

$$t_2(x,y) = P(x,y) \times \sum_{n=0}^{Q-1}\sum_{m=0}^{Q-1} e^{-i2\pi n\alpha x} e^{-i2\pi\phi_{n,m}x} e^{-i2\pi m\gamma y} e^{-i2\pi\psi_{n,m}y} \tag{5.55}$$

When intensities are considered and ensemble averaging takes place to allow an

analytic solution to the equation, we arrive at

$$\begin{aligned}\langle|t_2(x,y)|^2\rangle &= \langle|P(x,y)|^2 \\ &\quad\times\sum_{n=0}^{Q-1}\sum_{m=0}^{Q-1}\sum_{n'=0}^{Q-1}\sum_{m'=0}^{Q-1}\Big[e^{-i2\pi(n-n')\alpha x}e^{-i2\pi(m-m')\gamma y} \\ &\quad\times e^{-i2\pi(\phi_{n,m}x+\psi_{n,m}y)}e^{i2\pi(\phi_{n',m'}x+\psi_{n',m'}y)}\Big]\rangle \qquad (5.56)\\ &= |P(x,y)|^2 \\ &\quad\times\langle\sum_{n=0}^{Q-1}\sum_{m=0}^{Q-1}\sum_{n'=0}^{Q-1}\sum_{m'=0}^{Q-1}\Big[e^{-i2\pi(n-n')\alpha x}e^{-i2\pi(m-m')\gamma x} \\ &\quad\times e^{-i2\pi\phi_{n,m}x}e^{i2\pi\phi_{n',m'}x}e^{-i2\pi\psi_{n,m}y}e^{i2\pi\psi_{n',m'}y}\Big]\rangle \qquad (5.57)\end{aligned}$$

This can be simplified greatly if the values of the random variables are mutually independent. This is the case except when $n = n'$ and $m = m'$, fixing the values of $\phi_{n',m'}$ and $\psi_{n',m'}$ if $\phi_{n,m}$ and $\psi_{n,m}$ are known. If these cases are explicitly removed from the summation then the averaging can be performed over each function of the random variables, in a manner similar to the 1-d analysis, to give

$$\begin{aligned}\langle|t_2(x,y)|^2\rangle &= |P(x,y)|^2\times\Big[Q^2 \\ &\quad+\underbrace{\sum_{n=0}^{Q-1}\sum_{n'=0}^{Q-1}e^{-i2\pi(n-n')\alpha x}\sum_{m=0}^{Q-1}\sum_{m'=0}^{Q-1}e^{-i2\pi(m-m')\gamma y}}_{n,m\neq n',m'} \\ &\quad\times\langle e^{-i2\pi\phi_{n,m}x}\rangle\langle e^{i2\pi\phi_{n',m'}x}\rangle\langle e^{-i2\pi\psi_{n,m}y}\rangle\langle e^{i2\pi\psi_{n',m'}y}\rangle\Big] \qquad (5.58)\end{aligned}$$

Substituting for the $\langle ...\rangle$ terms, and rearranging using

$$\underbrace{\sum_{n=0}^{Q-1}\sum_{n'=0}^{Q-1}e^{-i2\pi(n-n')\alpha x}\sum_{m=0}^{Q-1}\sum_{m'=0}^{Q-1}e^{-i2\pi(m-m')\gamma y}}_{n,m\neq n',m'} = \frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)}\frac{\sin^2(\pi Q\gamma y)}{\sin^2(\pi\gamma y)} - Q^2 \qquad (5.59)$$

gives the general 2-d equation for a $Q \times Q$ randomised pixel array:

$$\begin{aligned}\langle|t_2(x,y)|^2\rangle &= |P(x,y)|^2 \times Q^2 \left[1 - |\rho(x)|^2|\rho(y)|^2 + \right. \\ &\left. \frac{1}{Q^2}\frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)}\frac{\sin^2(\pi Q\gamma y)}{\sin^2(\pi\gamma y)}|\rho(x)|^2|\rho(y)|^2\right] \qquad (5.60)\end{aligned}$$

where $\rho(y)$ is the Fourier transform of the probability distribution function $P_{\psi_{n,m}}(\psi)$.

This equation is similar to that of the 1-d case, and noting that observations made for $\rho(x)$ carry over and also apply to $\rho(y)$, is indeed exactly the same along the axes except for a multiplying constant. More important, however, is that the general form is the same, with the $|\rho(..)|^2$ terms attenuating the spectral orders, the redistributed energy giving rise to a background illumination with a light depletion zone surrounding the zero order, and the zero order being unchanged in magnitude. Further, the ratio of background and zero order intensities is still inversely related to the number of pixels i.e.

$$\frac{\langle\text{background}\rangle}{\langle|t_2(0,0)|^2\rangle} \leq \frac{1}{Q^2} \qquad (5.61)$$

where this time the mean background contribution is proportional to $[1-|\rho(x)|^2|\rho(y)|^2]$.

Thus in the real 2-d case it is possible that some or all of the replicas may be suppressed, depending on the randomisation scheme used.

5.3 Specific examples of randomised arrays

Two different randomisation schemes have been tried, both subject to the constraint that pixels do not overlap,which is necessary for practical implementation on an SLM. Choosing the pixels to be square with side β means that

$$-\frac{\alpha-\beta}{2} \leq \phi_{n,m}, \psi_{n,m} \leq \frac{\alpha-\beta}{2}, \tag{5.62}$$

and makes $p(\xi,\eta) = \text{rect}(\xi/\beta, \eta/\beta)$, and $P(x,y) = \beta^2 \,\text{sinc}(\beta x, \beta y)$ as before. This is equivalent to considering the pixel to be confined within a square cell of sides equal to the regular pixel pitch, with no part of the pixel allowed to lie outside the cell.

The pixels were chosen to be square for consistency with silicon backplane SLMs developed at the University of Edinburgh [67]. For the same reason the array size was set to be 16x16 and the pixel pitch made the same in both the ξ and η directions.

The first example has the offsets uniformly distributed within the allowed range, and is perhaps intuitively the best candidate to reduce the size of the spectral orders. In practice this is not the case, though the example is worth studying as means of qualitatively describing the effect of phase variation induced by surface roughness on transparencies (by analogy with detour-phase coding in holography), assuming the variations are uniformly distributed over a certain range of sizes.

The second basic example arises from simple consideration of what is required from $\rho(x)$ or $\rho(y)$, and results in selecting one of four discrete positions for the pixel within each cell. This surprisingly simple randomisation scheme makes it possible to eliminate all the odd spectral orders, and allows all the higher orders to be eliminated if α:β is 2. A problem with the continuity of the ITO electrodes used in the matrix addressing arises which can prevent this scheme being implemented exactly, but an effective means of surmounting this problem is presented.

The analysis has thus far concentrated on the ensemble average of all possible sets of choices of $\phi_{n,m}, \psi_{n,m}$, while any implementation of the randomisation schemes in an SLM will use the values of random variable specific to one member of the ensemble. Simulations generating the Fourier power spectrum $|t_2(x)|^2$ of one member of each of the randomisation schemes' ensembles are included for comparison with those from the average of a number of members, and shown to be in good

agreement.

As $\phi_{n,m}$ and $\psi_{n,m}$ have the same form of probability distribution function, we need consider only $\phi_{n,m}$ in detail.

Note on presentation of results

In all the simulations the intensity values are referred to a maximum value of 256 units at the zero order, which in the computational simulations is offset to the point (129,129). It was necessary to use the **unigraph** [95]software package to display grey scale plots of the low-level details of the power spectra, as the details were lost in the image files used by xv which are quantised to 256 grey levels. Otherwise, the 2-d grey scale pictures are displayed using xv with gamma=1.5.

5.3.1 Example 1: Uniformly random distribution

As stated $\phi_{n,m}$ is uniformly distributed in the allowed range $-\frac{\alpha-\beta}{2} \leq \phi_{n,m} < \frac{\alpha-\beta}{2}$, which gives a normalised probability distribution function

$$P_{\phi_{n,m}}(\phi) = \frac{1}{\alpha-\beta}\text{rect}\left(\frac{\phi}{\alpha-\beta}\right) \tag{5.63}$$

shown in figure 5.1.

Using this distribution function, and using NAG routines to generate the random numbers, a FORTRAN program was written to generate a 2-d representation of one member of the ensemble. An example of the resulting 16x16 array is illustrated in figure 5.2 for the case $\alpha = 2, \beta = 1$.

The probability distribution function is Fourier transformed giving

$$\rho(x) = \text{sinc}\left((\alpha-\beta)x\right). \tag{5.64}$$

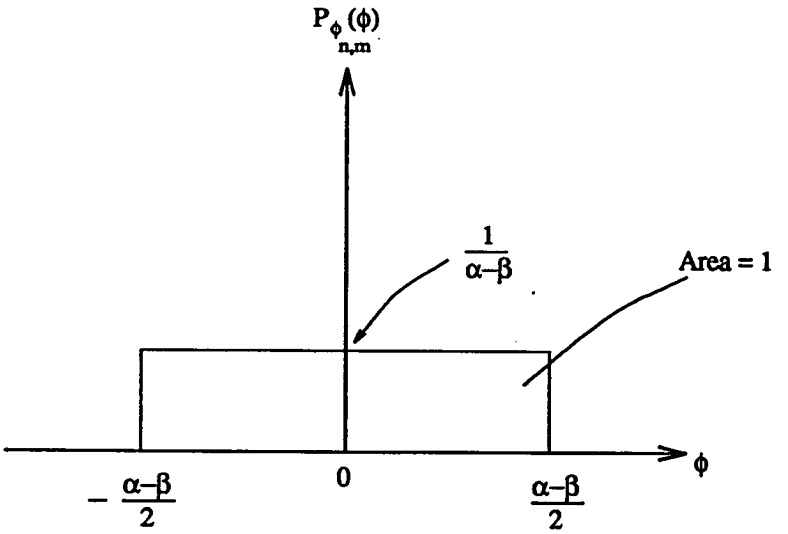

Figure 5.1: Normalised probability distribution function for $\phi_{n,m}$, uniformly distributed random distribution.

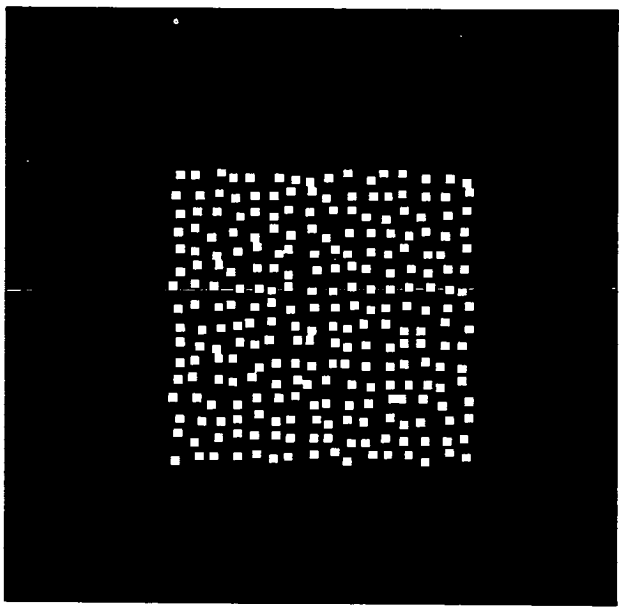

Figure 5.2: 16x16 array with uniformly distributed, non-overlapping pixels

Similarly $\rho(y) = \mathrm{sinc}((\alpha-\beta)y)$, establishing the form of the diffraction pattern as

$$\begin{aligned}\langle|t_2(x,y)|^2\rangle &= Q^2\beta^4\mathrm{sinc}^2(\beta x)\,\mathrm{sinc}^2(\beta y)\times\Big[1-\mathrm{sinc}^2\left((\alpha-\beta)x\right)\,\mathrm{sinc}^2\left((\alpha-\beta)y\right)\\ &\quad+\frac{1}{Q^2}\frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)}\frac{\sin^2(\pi Q\alpha y)}{\sin^2(\pi\alpha y)}\mathrm{sinc}^2\left((\alpha-\beta)x\right)\,\mathrm{sinc}^2\left((\alpha-\beta)y\right)\Big],\end{aligned} \tag{5.65}$$

substituting for $P(x,y)$.

This expression as it stands is quite unwieldy, until it is considered in terms of its smaller constituent parts as introduced earlier. When considered with respect to α and β, along with the practicalities of SLM construction it emerges that it is advantageous to set $\beta = \alpha/2$.

To understand this we first simplify equation 5.65 by taking a section along the x-axis, to give

$$\begin{aligned}\langle|t_2(x,0)|^2\rangle &= Q^2\beta^4\mathrm{sinc}^2(\beta x)\times\Big[1-\mathrm{sinc}^2\left((\alpha-\beta)x\right)\\ &\quad+\frac{1}{Q^2}\frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)}Q^2\mathrm{sinc}^2\left((\alpha-\beta)x\right)\Big]\end{aligned} \tag{5.66}$$

Given this equation the object now is to determine the relative sizes of α and β which will minimise the size of the spectral orders. This criterion must be most rigorously applied to the first and second orders, as we wish to suppress aliasing in the region of the zero order, with the ideal being to eliminate all the higher spectral orders and have the background illumination from the $1-\mathrm{sinc}^2((\alpha-\beta)x)$ term spread over as wide an area as possible. If we do not have to make an SLM this is trivial– simply set $\alpha = \beta$, and all the spectral orders disappear with no background illumination, approximating the finite width spectral orders as δ-functions. With the added constraint of non-overlapping pixels this corresponds to a completely regular array with a 100% fill factor, and cannot be fabricated.

In terms of equation 5.66 we note that when $\alpha = \beta$, $\text{sinc}^2((\alpha - \beta)x)$ is identically equal to 1, and the zeroes of $\text{sinc}^2(\beta x)$ occur at $x = n/\alpha$ (n an integer), as shown in figure 5.3a), coincident with the higher spectral orders. In this case all the higher orders have been suppressed by the part of the function due to the pixel dimensions, and the $\text{sinc}^2((\alpha - \beta)x)$ part equal to $|\rho(x)|^2$ has played no part in the reduction of the spectral orders.

When $\beta < \alpha$, as in all practical cases, $\text{sinc}^2((\alpha - \beta)x)$ can also contribute to suppressing the higher spectral orders. The effect of varying β with respect to α is considered for the constituent parts of equation 5.66:

- $\frac{\sin^2(\pi Q \alpha x)}{\sin^2(\pi \alpha x)}$ has had its behaviour described already, and gives rise to spectral orders at n/α of magnitude Q^2.

- $\text{sinc}^2(\beta x)$ has its zeroes at n/β. Hence its first zero will fall at increasing values of x as β decreases, allowing any one of the higher orders to be eliminated, plus any order which is an integral multiple of the lowest order eliminated– except the first order, as a consequence of $\beta < \alpha$. Hence the lowest order which can be eliminated as a consequence of the pixel function is the second order, corresponding to $\beta = \alpha/2$.

- $\text{sinc}^2((\alpha - \beta)x)$ has its zeroes at $n/(\alpha - \beta)$. Decreasing from $\beta = \alpha$, the position of the function's first zero falls at decreasing values of x, until $\beta = 0$, corresponding to the position of the first spectral order. However, the pixel has by now disappeared, so the lowest order which can be realistically eliminated as a consequence of this probability distribution function is the second order, corresponding to $(\alpha - \beta) = \alpha/2$: i.e. $\beta = \alpha/2$. This gives a fill factor of 25%, which coincidentally is the fill factor of the Applied Optics Group's 16x16 silicon backplane device.

- $1 - \text{sinc}^2((\alpha - \beta)x)$ shows the presence of the light depletion zone. Multiplying this by the $\text{sinc}^2(\beta x)$ term gives the shape of the generalised background illumination and allows an estimate of its maximium value. For the case $\beta = \alpha/2$, differentiation with respect to x shows that $\text{sinc}^2(\alpha/2x)\ (1 - \text{sinc}^2(\alpha/2x))$ has a maximum when $\text{sinc}^2(\alpha/2x) = 1/2$. This in turn implies

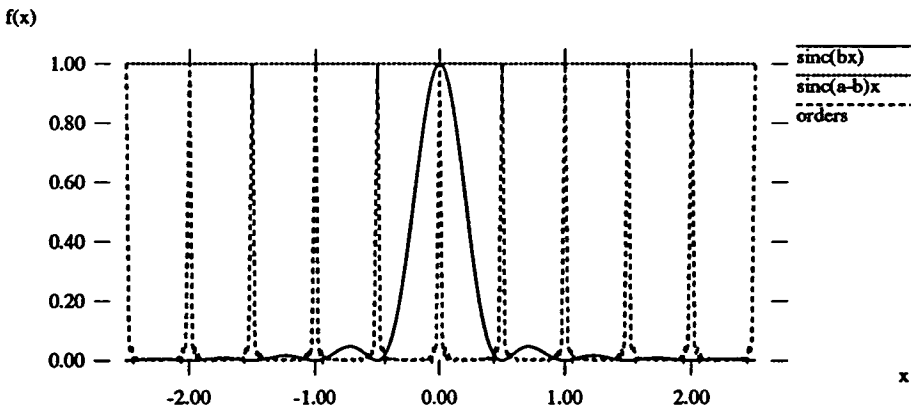

a) $\alpha = \beta$, $(\alpha = 2)$.

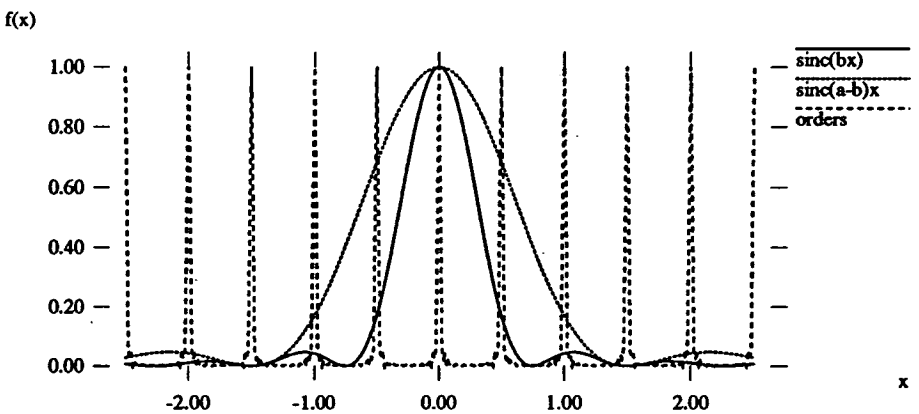

b) $\alpha = \frac{3}{2}\beta$, $(\alpha = 2)$.

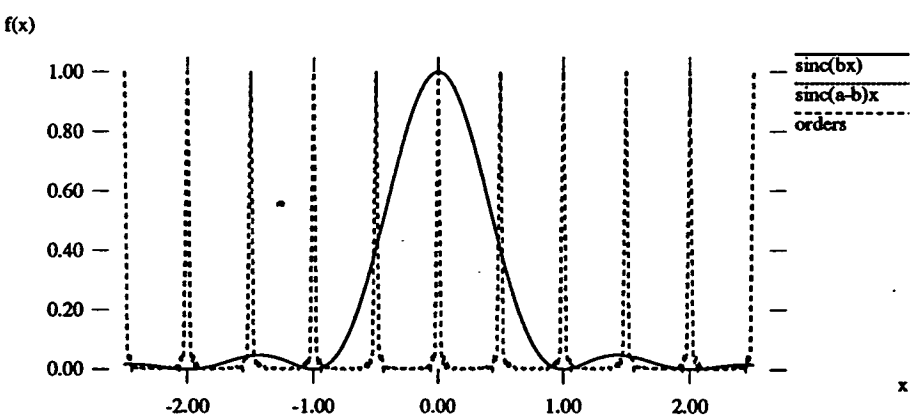

c) $\alpha = 2\beta$, $(\alpha = 2)$. Here $\mathrm{sinc}^2(\beta x)$ and $\mathrm{sinc}^2((\alpha - \beta)x)$ coincide.

Figure 5.3: Contributions to $\langle |t_2(x,0)|^2 \rangle$. The functions $\mathrm{sinc}^2(\beta x)$ and $\mathrm{sinc}^2((\alpha - \beta)x$ are plotted, and the positions of the spectral orders included for reference.

a maximum background intensity of $Q^2\beta^4/4$, when substituted back in the main equation. This can be compared with a zero order intensity of $Q^4\beta^4$.

The behaviour discussed above is shown in figures 5.3 b) and c).

Thus β is chosen to be $\alpha/2$ to eliminate the second spectral order, the lowest order which is practical to be eliminated, and all the higher even orders. This is also true for a regular array, but in the randomised case the remaining orders are attenuated further by the multiplying $\rho(...)$ terms.

Simulations for the uniformly randomised array

Simulations were carried out in the manner described earlier for the regular pixel array to verify the analysis. This time the computational simulations were carried out for an ensemble of randomised arrays as well as the single array shown in figure 5.2. The ensemble averaging was performed over 50 arrays with different sets of random displacements chosen according to the rules and probability distribution function described. These arrays' power spectra were calculated and averaged to give $\langle|t_2(x,y)|^2\rangle$, which is displayed in figure 5.4a). At this scale, the power spectrum of the single array of figure 5.2 looks identical. The low intensity details of the power spectra of the ensemble and the single array are shown in figure 5.4b) and c), which are thresholded at 2 units. These figures are comparable with the photograph from the optical simulation shown in figure 5.4d), where the low-level background detail in the power spectrum is shown clearly, at the expense of over-exposing the zero order. The optical simulation used an array photographically reduced from figure 5.2, mounted in a liquid gate, and with pixels on a pitch of 200μm.

Figure 5.5a) and b) shows computational simulations illustrating under- and near-critical sampling using the letters of figure 4.8 with the single random array used as a Fourier plane amplitude filter. Figure 5.5c) and d) show the corresponding optical simulations using the 4-f processor described for the regular array, with

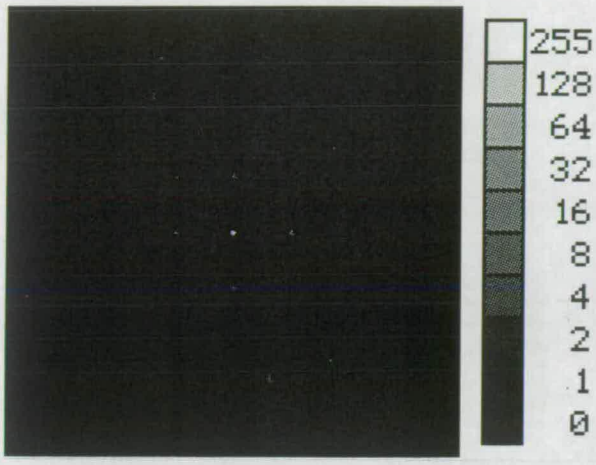

a) Average power spectrum of ensemble of 50 uniformly randomised arrays

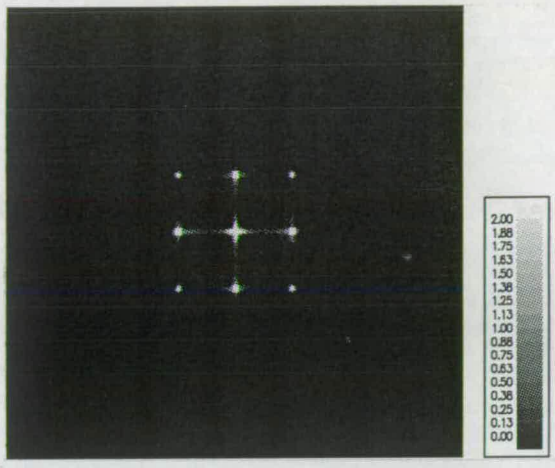

b) Low level detail in power spectrum of ensemble

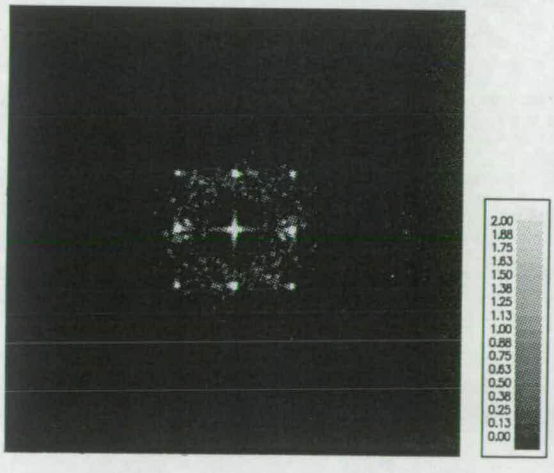

c) Low level detail in power spectrum of member of ensemble

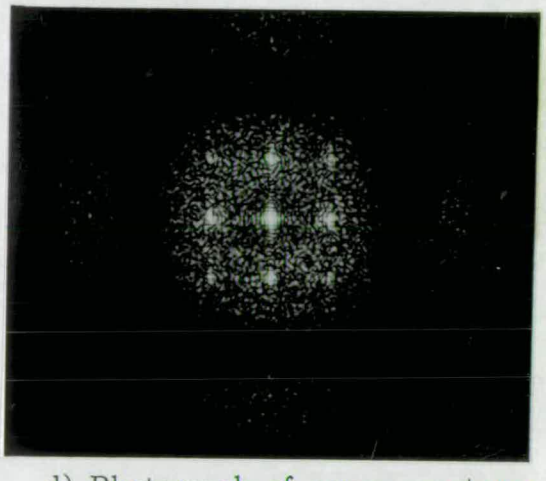

d) Photograph of power spectrum of single member from optical simulation

Figure 5.4: Power spectra simulations for uniformly randomised pixel arrays. Note light depletion regions in b), c), and d).

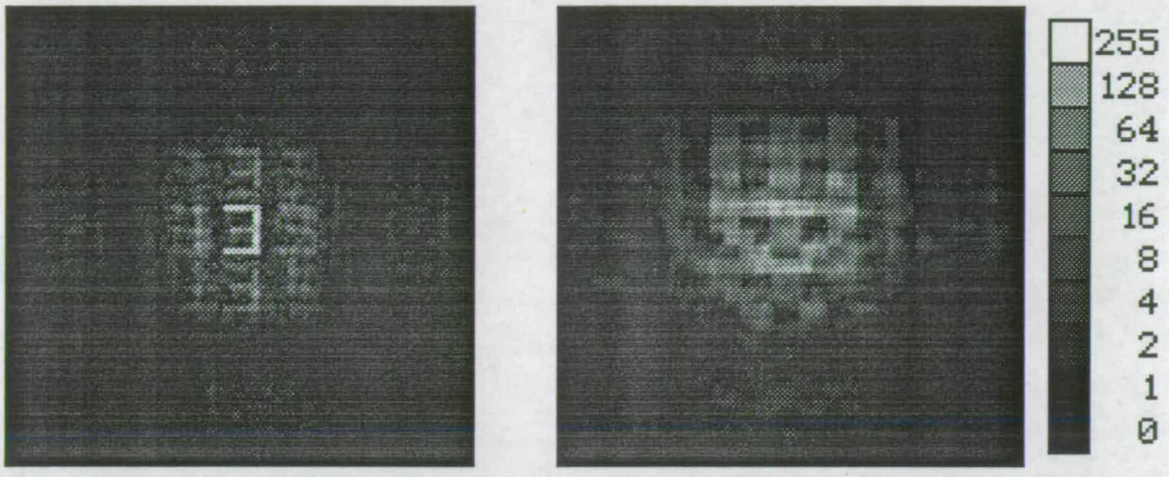

a) Computational simulation for near-critically sampled E

b) Computational simulation for undersampled A

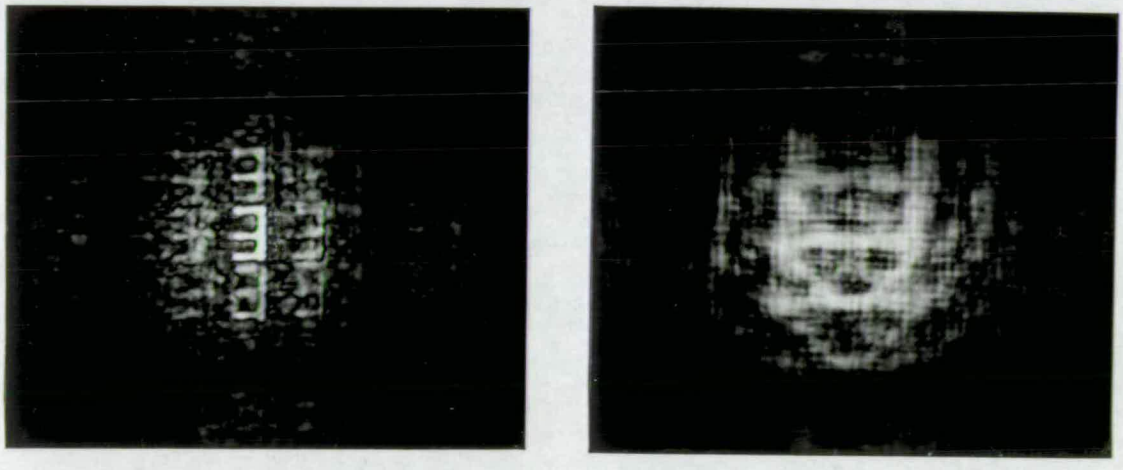

c) Optical simulation for near-critically sampled E

d) Optical simulation for undersampled A

Figure 5.5: Output images illustrating sampling for input letters filtered by uniformly randomised pixellated filter of figure 5.2.

the inputs and Fourier plane filter in the form of transparencies of the letters and random array.

Discussion of results

Qualitatively the ensemble averaged power spectrum and the power spectrum of the single member of the ensemble are in good agreement with each other, and with the analysis, showing the expected light depletion zone, generalised background away from the zero order, missing even orders and odd orders that are attenuated compared to the regular case. This is shown more quantitatively in figure 5.6 which shows the cross sections along the x-axis from the computational simulations. Figure 5.6c) shows that the low level background underlying the narrow $\mathrm{sinc}^2(..)$ functions of the spectral orders is of the $(1 - \mathrm{sinc}^2(\frac{\alpha}{2}x))$ form predicted by the analysis, and also of the correct magnitude. The figure also shows a 3-d representation of the power spectrum $|t_2(x,y)|^2$ of the single array shown in figure 5.2, which should be compared with figure 4.10a) from the regularly pixellated case.

The intensities of the spectral orders nearest the origin in the optical simulation were measured using an array camera and framestore as before, and are compared with the calculated values from the analytical expression in table 5.1. The optical results have the background value of 2 subtracted, and are averaged over the orders which should have equal intensities. The background value is the mean intensity over a 30x19 pixel area of the image well away from the spectral orders. The calculated values of the spectral order intensities $I_{n,m}$ are derived from

$$I_{n,m} = Q^2\beta^4 \,\mathrm{sinc}^2(\frac{n}{2})\,\mathrm{sinc}^2(\frac{m}{2})\left[1 - \mathrm{sinc}^2(\frac{n}{2})\,\mathrm{sinc}^2(\frac{m}{2}) + Q^2\,\mathrm{sinc}^2(\frac{n}{2})\,\mathrm{sinc}^2(\frac{m}{2})\right] \tag{5.67}$$

with the tabulated values divided by $(Q\beta)^4$ to normalise intensities relative to the zero order. The tabulated results show that the theory and optical simulations are in good agreement.

The simulations from the undersampled As, however, show that aliasing from the first orders is still significant, and places the same restriction on the size of

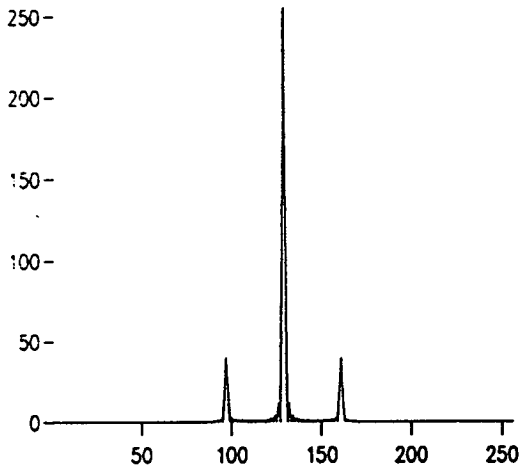

a) Cross-section along $x-$axis for ensemble average

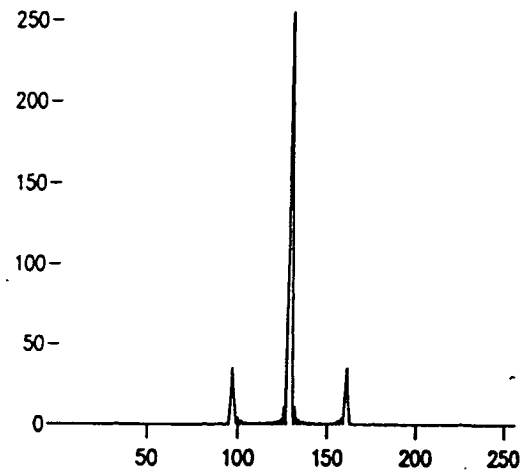

b) Cross-section along $x-$axis for individual member

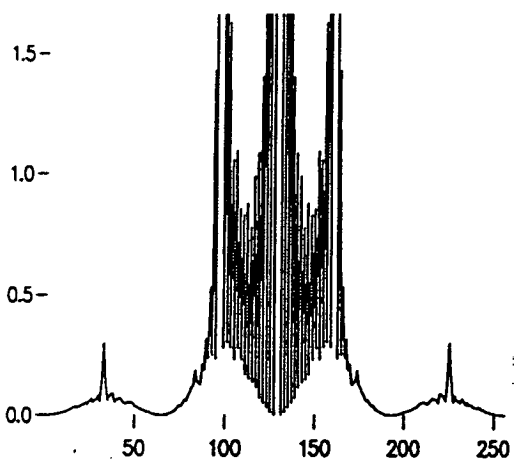

c) Low level detail of a) showing light depletion region

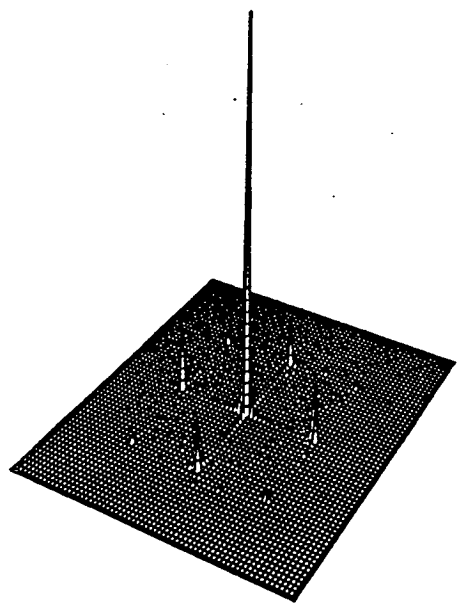

d) 3-d representation of power spectra of the single array

Figure 5.6: Power spectra cross-sections for the uniformly randomised array

Order	Calculated relative $I_{n,m}$	Measured $I_{n,m}$	Measured relative $I_{n,m}$
(0,0)	1.00	168±3	1.00
(1,0)(0,1) (-1,0)(0,-1)	$\frac{1}{Q^2}\frac{4}{\pi^2}[1-\frac{4}{\pi^2}+Q^2\frac{4}{\pi^2}]=0.17$	30±3	0.18±0.02
(1,1)(1,-1) (-1,-1)(-1,1)	$\frac{1}{Q^2}(\frac{4}{\pi^2})^2[1-(\frac{4}{\pi^2})^2+Q^2(\frac{4}{\pi^2})^2]=0.03$	10±3	0.06±0.02

Table 5.1: Spectral order intensities for uniformly randomised array

the input object as in the case of the regular filter with equivalent pixel pitch. To overcome this we consider a randomisation scheme which eliminates the odd orders in addition to eliminating the higher even orders.

5.3.2 Example 2: Discrete position random distribution

This randomisation scheme is specifically designed to eliminate the odd orders. With reference to equation 5.65, it has already been shown that choice of pixel function alone cannot practically achieve this, so it is necessary to select $\rho(x), \rho(y)$ such that they have zeroes at $x = \frac{2n+1}{\alpha}$ or $y = \frac{2m+1}{\alpha}$ respectively, with n, m both integers. A simple function which satisfies this is

$$\rho(x) = \cos(\frac{\pi\alpha}{2}x) \tag{5.68}$$

The resulting probability distribution function is found by Fourier transformation to be

$$P_{\phi_{n,m}}(\phi) = \frac{1}{2}\delta(\phi - \frac{\alpha}{4}) + \frac{1}{2}\delta(\phi + \frac{\alpha}{4}) \tag{5.69}$$

as illustrated in figure 5.7

This is clearly a useful result as it gives rise to possible pixel positions within $\pm\frac{\alpha}{2}$

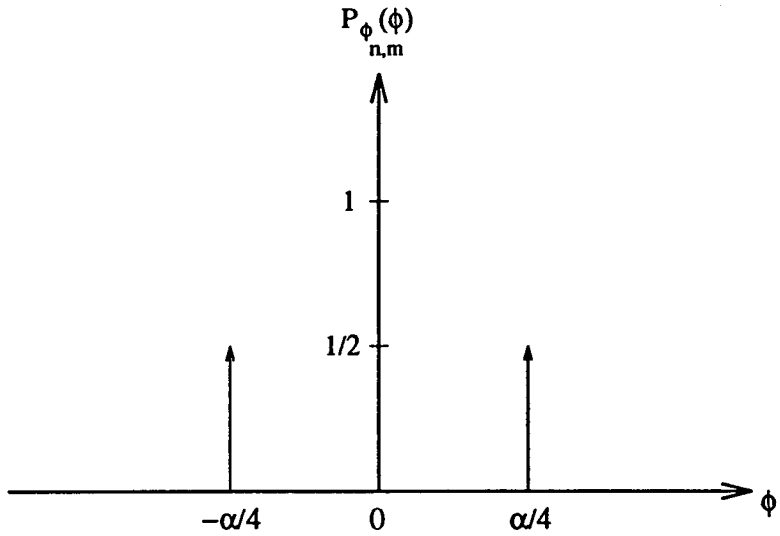

Figure 5.7: Normalised probability distribution function for $\phi_{n,m}$, discrete position random distribution.

of the undisplaced pixel centre, satisfying the earlier condition which requires the pixels not to overlap, equivalent to the pixel being confined in an $\alpha \times \alpha$ cell. This condition also constrains the pixel width $\beta \leq \frac{\alpha}{2}$. Obviously $\beta = \frac{\alpha}{2}$ is favoured both in terms of maximising the transmitting area (or fill factor, to borrow a phrase from the field of active backplane SLMs) and because such a pixel width will lead to the elimination of the higher even orders. Any attempt to use a discrete position pixel scheme with a wider pixel width will result in the δ-functions marking the pixel centres lying less than $\alpha/4$ away from the undisplaced pixel centre, due to the non-overlapping condition. This in turn will broaden the cos(...) function of $\rho(x)$ due to the reciprocal nature of the Fourier transform, resulting in $\rho(x)$'s first zero being too far from the origin to eliminate the first orders.

With $\rho(y)$ specified similarly to $\rho(x)$, i.e. $\rho(y) = \cos(\frac{\pi\alpha}{2}y)$, each pixel can occupy one of four positions in the $\alpha \times \alpha$ cell with equal probability for each position. These positions are shown in figure 5.8.

Again, a FORTRAN program was written to implement this randomisation scheme and an example of the resulting 16x16 array is shown in figure 5.9. The form of the diffraction pattern of this array is then given by

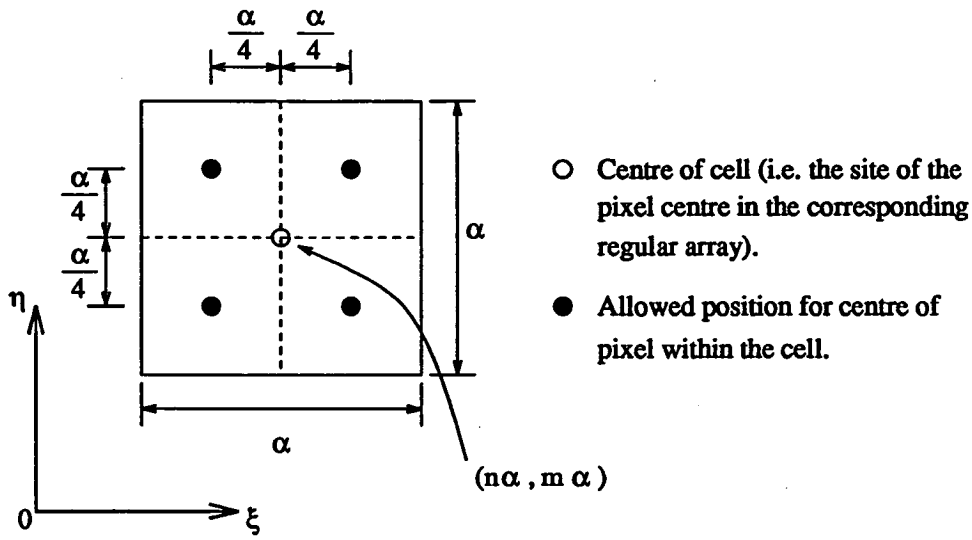

Figure 5.8: Possible pixel positions within $\alpha \times \alpha$ cell.

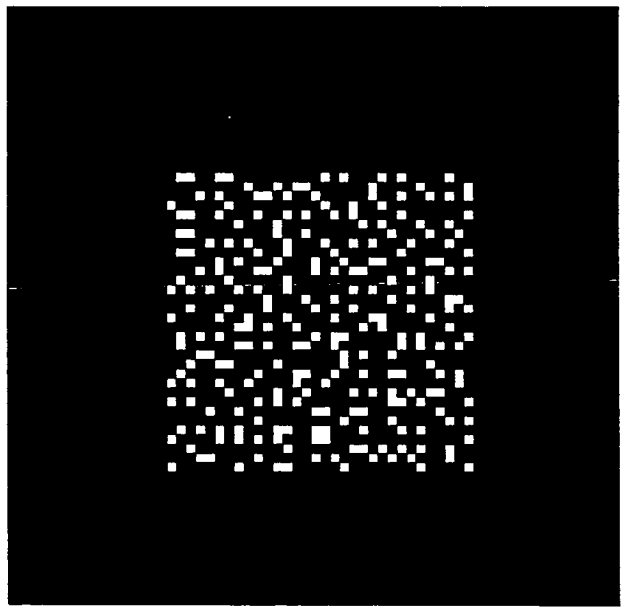

Figure 5.9: 16x16 array with discretely positioned non-overlapping pixels

$$\langle|t_2(x,y)|^2\rangle = Q^2\beta^4\mathrm{sinc}^2(\tfrac{\alpha}{2}x)\,\mathrm{sinc}^2(\tfrac{\alpha}{2}y) \times \Big[1-\cos^2(\tfrac{\pi\alpha}{2}x)\,\cos^2(\tfrac{\pi\alpha}{2}y) + \frac{1}{Q^2}\frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)}\frac{\sin^2(\pi Q\alpha y)}{\sin^2(\pi\alpha y)}\cos^2(\tfrac{\pi\alpha}{2}x)\,\cos^2(\tfrac{\pi\alpha}{2}y)\Big] \tag{5.70}$$

Rearranging using the trigonometric identity $(\cos(A)\sin(A))^2 = \frac{1}{4}\sin^2(2A)$ yields

$$\langle|t_2(x,y)|^2\rangle = \underbrace{Q^4\beta^4\,\mathrm{sinc}^2(Q\alpha x)\,\mathrm{sinc}^2(Q\alpha y)}_{\text{zero order}} + \underbrace{Q^2\beta^4\,\mathrm{sinc}^2(\tfrac{\alpha}{2}x)\,\mathrm{sinc}^2(\tfrac{\alpha}{2}y)\,\Big[1-\cos^2(\tfrac{\pi\alpha}{2}x)\,\cos^2(\tfrac{\pi\alpha}{2}y)\Big]}_{\text{everything else: generalised background}} \tag{5.71}$$

This explicitly separates the zero order from the background and immediately shows that the background to zero order intensity ratio is $\leq \frac{1}{Q^2}$.

By setting $y = 0$ the equation reduces further to

$$\langle|t_2(x,0)|^2\rangle = Q^4\beta^4\mathrm{sinc}^2(Q\alpha x) + Q^2\beta^4\frac{\sin^4(\frac{\pi\alpha}{2}x)}{(\frac{\pi\alpha}{2}x)^2} \tag{5.72}$$

and by differentiating the background contribution with respect to x, it can be shown that the background has a maximum when $tan(\frac{\pi\alpha}{2}x) = \pi\alpha x$. This implies a maximum background intensity of $0.525\,Q^2\beta^4$, occuring at $x = 0.742\,\alpha$.

Simulations for discrete position randomised array

Computational and photographic simulations were again carried out in the manner described earlier. As in the previous example, an ensemble averaged power

spectrum, averaged over 50 different arrays with different sets of choices for the pixel positions, was calculated. This is shown in figure 5.10. Again, the power spectrum of the single array is indistinguishable from that of the ensemble at this scale. Figure 5.10 b) and c) shows the low intensity detail of the respective power spectra following thresholding at 2 units, and figure 5.10d) shows the photograph from the corresponding optical simulation using an array photographically reduced from figure 5.9, and mounted in a liquid gate.

Figure 5.11 shows computational simulations illustrating under- and near-critical sampling for the single array used as a Fourier plane amplitude filter, along with the corresponding optical simulations.

Discussion of results

The computational simulations show that indeed all the diffraction orders other than the zero order have been eliminated, and replaced by a low level background, the form of which is as expected from the analysis. The single array compares well with the ensemble in terms of similarity, though the background is less uniform. This is evident when cross-sections along the $x-$axis are taken, shown in figure 5.12a) and b). The computational simulation of the ensemble averaged spectrum is in particularly good agreement with the analysis, with reference to figure 5.12c). As can be seen, the simulation shows a damped oscillatory $\mathrm{sinc}^2(Q\alpha x)$ function added to the background, with a central peak 4 points wide and oscillations every 2 points. As α is equivalent to a width of 8 points in the 256x256 array used in the computational simulations, representing a size of 8/256 in the object array, separation of the spectral orders will be $\frac{1}{\alpha} = 32$ points in the transformed array, and the zeroes of the $\mathrm{sinc}^2(Q\alpha x)$ function should indeed be spaced by $\frac{1}{Q\alpha} = 2$ points. The zero order in this representation, as in all the 2-d computational simulations, is equal to 256, so using the earlier result the maximum value of the background should be $\frac{0.525}{Q^2} \times 256 = 0.525$, occurring at $\pm 0.742\alpha = \pm 24$ points from the zero order, i.e. at 129 ± 24, allowing for the offset of the zero order in the computational simulations. The figure shows this to be the case. The zeroes of the background function occur at the positions at which the even

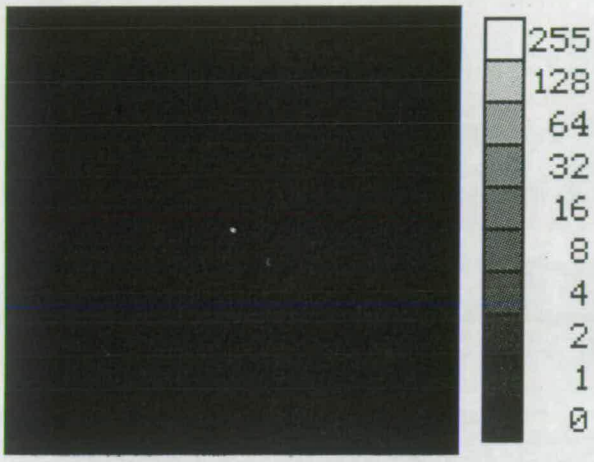

a) Average power spectrum of ensemble of 50 discrete position randomised arrays

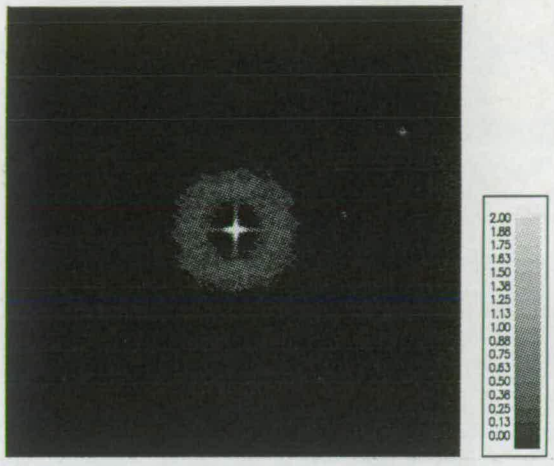

b) Low level detail in power spectrum of ensemble

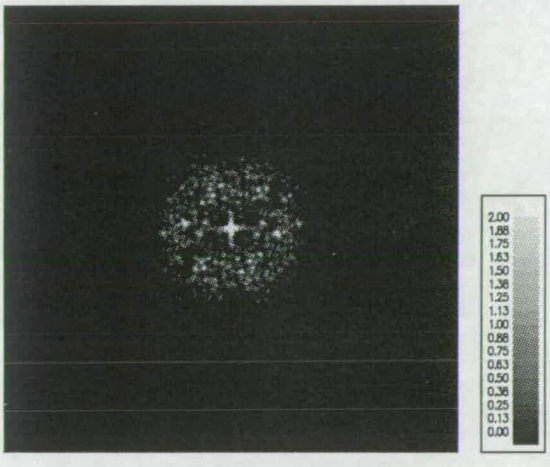

c) Low level detail in power spectrum of member of ensemble

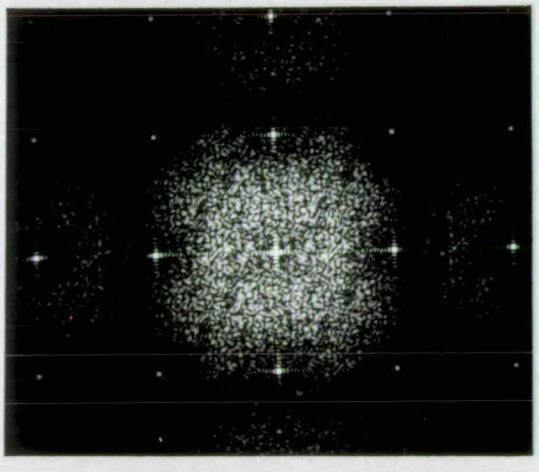

d) Photograph of power spectrum of single member from optical simulation

Figure 5.10: Power spectra simulations for discrete position randomised pixel arrays.

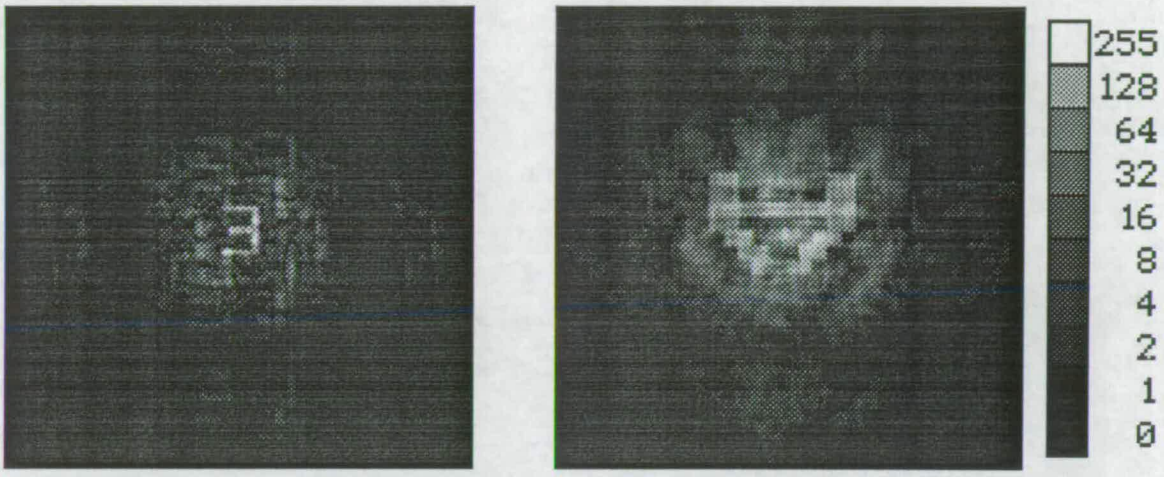

a) Computational simulation for near-critically sampled E

b) Computational simulation for undersampled A

c) Optical simulation for near-critically sampled E

d) Optical simulation for undersampled A

Figure 5.11: Output images illustrating sampling for input letters filtered by discrete position randomisation pixellated filter of figure 5.9.

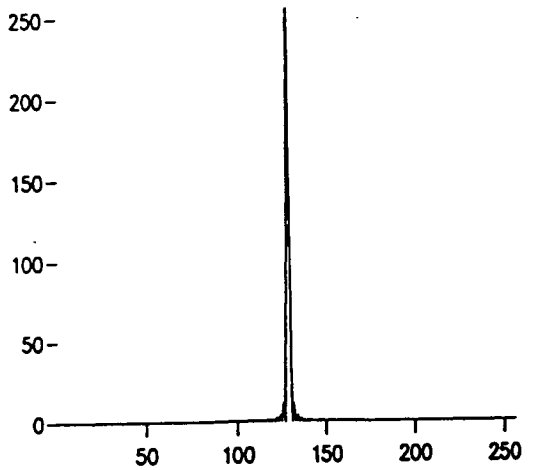

a) Cross-section along $x-$axis for ensemble average

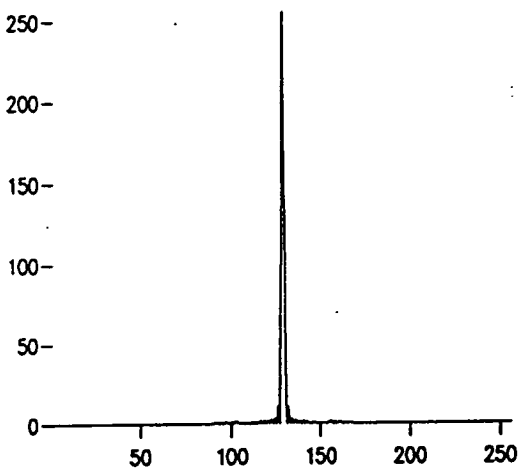

b) Cross-section along $x-$axis for individual member

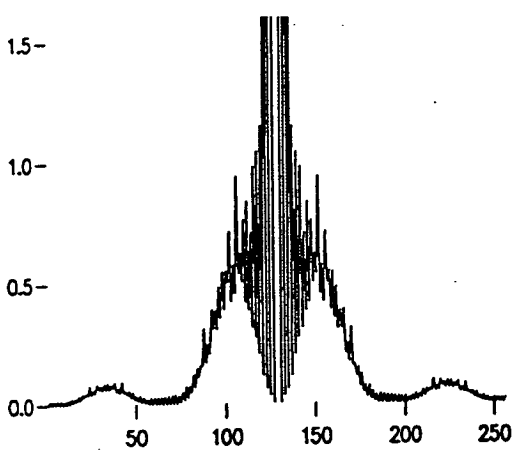

c) Low level detail of a) showing light depletion region

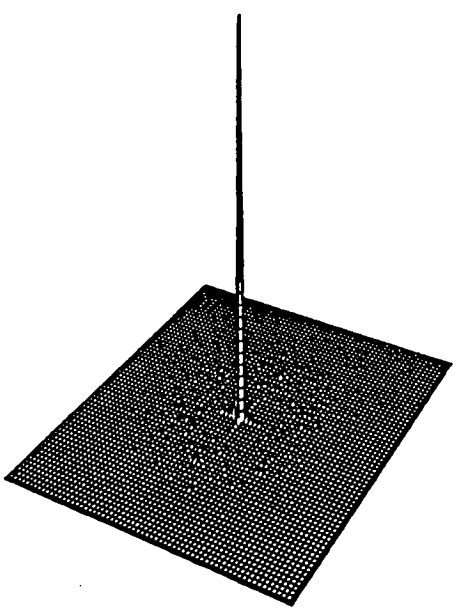

d) 3-d representation of power spectra of the single array

Figure 5.12: Power spectra cross-sections for the discrete position randomised array

spectral orders would occur.

The optical simulations show faint signs of the spectral orders, but were too low in intensity to distinguish from the dark background when intensities were measured using the framestore and array camera. The light depletion region is also not evident when the intensity of the zero order is at a level that does not saturate the array camera– this is as expected, given that the framestore can resolve 255 grey values, and that the light depletion region should be $\leq 1/256$ of the zero order. The faint peaks can also just be detected on the computational simulations. Small deviations from the ideal such as this are to be expected when only one member of the ensemble is considered. Nevertheless, this randomisation scheme gives excellent results, with the sharp peaks from the higher spectral orders being almost completely eliminated and replaced by a smooth low level background, leading to much greater tolerance of aliasing, as shown in the examples of undersampling in figure 5.11.

Unfortunately, before making use of this result in optical processing applications with an SLM, one serious problem must be overcome– a matrix addressed SLM using the array pattern shown in figure 5.9 cannot be fabricated, nor others like it generated from the same rules. The reason for this is apparent with reference to the mask showing the electrodes which would overlap to define the pixels of the array (figure 5.13). As can be seen, the relative positions of some of the neighbouring pixels cause the connecting electrodes to pinch off, leading to unwanted open circuits and failure of many of the pixels. Fortunately it is possible to reproduce the preceding results using an array which can be fabricated as an SLM, as shown in the following example.

5.3.3 Example 3: Special case of discrete position randomisation

As is implied from the heading, this randomisation scheme is closely related to the previous example. Here the possible pixel positions are the same as in Example 2

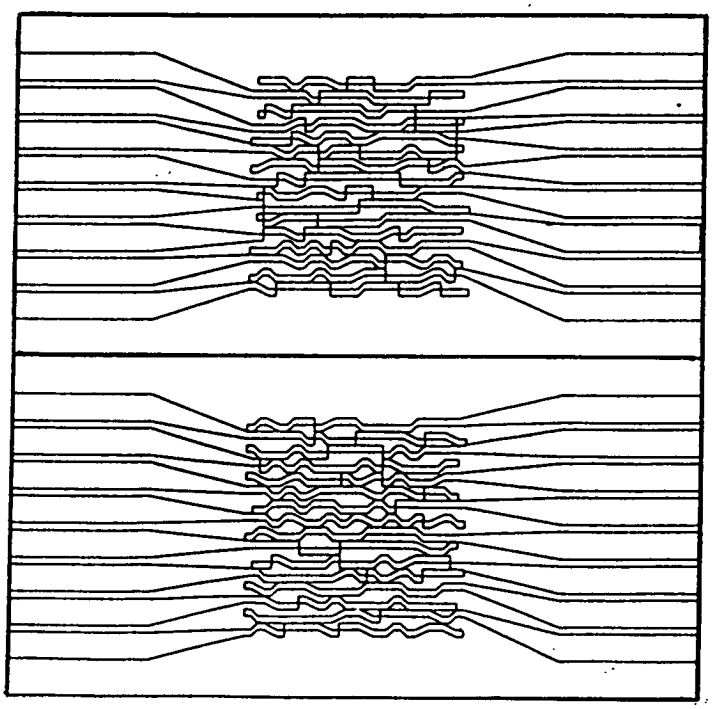

Figure 5.13: Mask to fabricate SLM using discrete position randomisation.

except that a 'nearest neighbour' condition has been added to the choice of pixel position which prevents the pinch off situation occurring.

The nearest neighbour condition chosen was to restrict the selection of pixel position to possible positions in the cell which would not lead to pinch off occuring between the pixel under consideration and the ones in the cells above and to its left, these cells being labelled A and B respectively. As the pixel can take one of four positions in each of A and B, this leads to a set of 16 groups of possible positions, as shown in figure 5.14.

Thus prior knowledge of the pixel positions in the adjoining cells A and B is necessary before any new choice of pixel position can be made. This is accomplished by defining a set of initial conditions in the form of an imaginary row of cells above the first row of the actual array, and an imaginary column of cells to the left of the first column of the actual array, with all the pixels in these imaginary entries in position 0 according to the labelling scheme of figure 5.14. With this set of initial conditions in place a choice of pixel position can be made from the allowed positions for each cell in the actual array, starting in the top left corner,

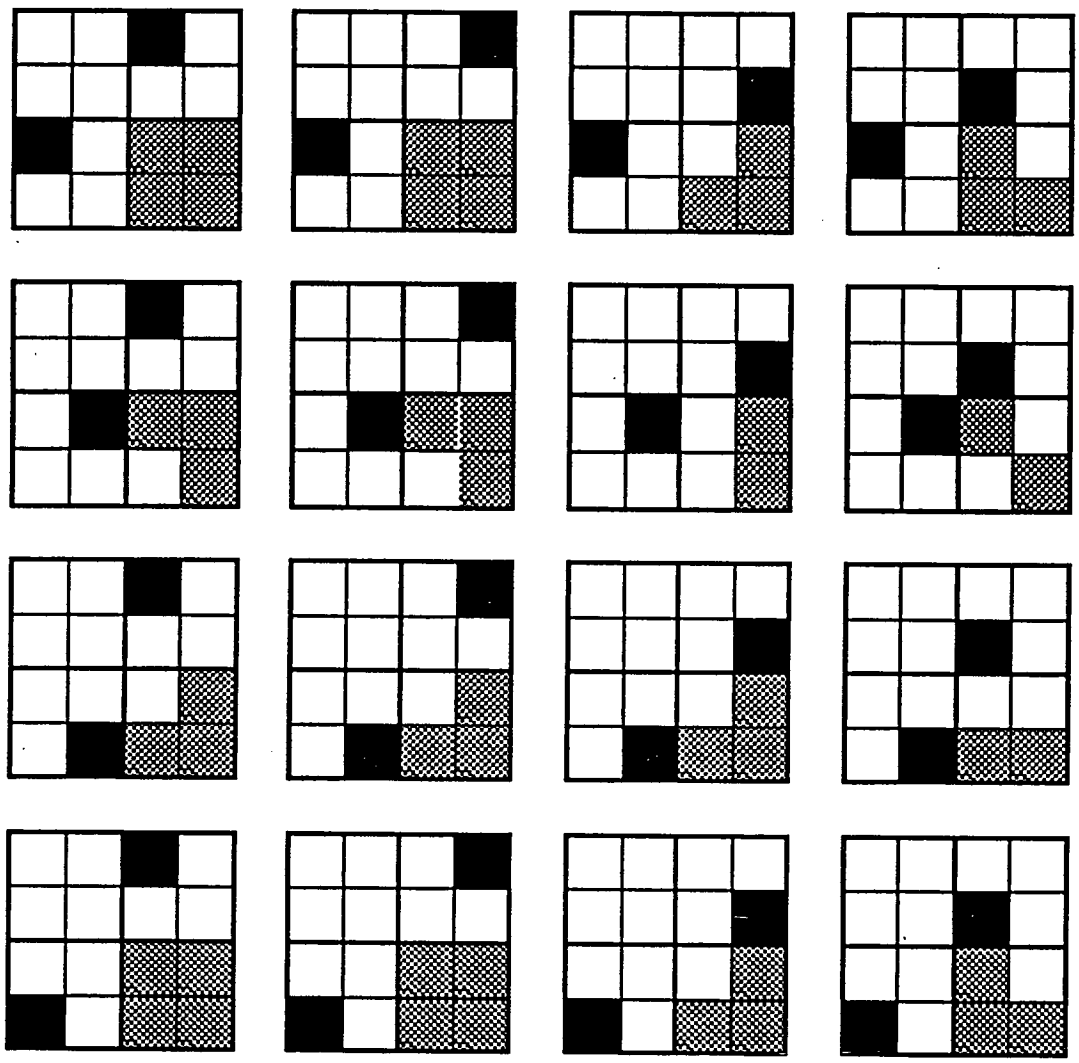

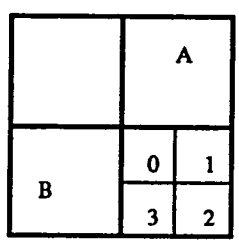

The shaded areas are the possible positions that the pixel can take (corresponding to the labels 0, 1, 2, 3) which prevent breaks in the connecting electrodes, given the positions of the pixels in A and B

Figure 5.14: Set of groups of allowed pixel positions.

working along the top row, and left to right along each row in turn until the final cell in the bottom right corner of the array. Position 0 was chosen for the initial conditions to give the widest possible choice of allowed positions for the entries in the top row and first column of the actual array.

No mention has been made yet of the relative probabilities of the selections from the allowed positions. The first step was to write a program which repeatedly picked one of the four possible positions at random each with probability $\frac{1}{4}$, until the position chosen corresponds to one of those allowed, given the pixel positions in A and B. This initial attempt yielded an array showing obvious periodicity in the η direction, suggesting that some weighting of the probabilities of choice of pixel position was necessary. This is apparent from consideration of figure 5.14 which shows that while position 2 is possible in all 16 of the nearest neighbour configurations, position 0 is possible only 9 times, and positions 1 and 3 are each possible 12 times out of 16. This would suggest that weighting in favour of position 0 and against position 2 be carried out. In fact the situation is much more complex than this, as the selection procedure for the pixel positions is explicitly dependent on the neighbouring pixels, and a rigourous determination of an optimum weighting scheme would require knowledge of the relative probabilities of the nearest neighbour pixel configurations– which depends on the weighting scheme. Hence an assumption was made that each of the 16 nearest neighbour pixel configurations was equally probable and weighting carried out on the basis of the occurrences of allowed pixel positions.

The weighting is simply proportional to the inverse of a position's occurrence in figure 5.14, so if P(n) is the weighted probability of picking a pixel to be in position n, we have

Position, n	Normalised Probability, P(n)	Weighting
0	$P(0)=\omega \times \frac{1}{9} = 0.326$	1.304
1	$P(1)=\omega \times \frac{1}{12} = 0.245$	0.981
2	$P(2)=\omega \times \frac{1}{16} = 0.184$	0.736
3	$P(3)=\omega \times \frac{1}{12} = 0.245$	0.981
Total	1.000	4.000

where ω is a normalisation constant.

In the FORTRAN program written to generate the resulting array the weighting and choice of pixel is carried out in a subroutine which has the possible positions passed to it as arguments, and which returns a choice to the main program along with a flag saying whether it is allowed or not. If the choice is flagged as invalid the process is repeated until an allowed selection is made. The weighting is implemented by altering the decisive limits that the random variable xx lies between i.e.

if $0.00 < xx \leq 1.31$, choose position 0
if $1.31 < xx \leq 2.29$, choose position 1
if $2.29 < xx \leq 3.02$, choose position 2
if $3.02 < xx \leq 4.00$, choose position 3

where xx is uniformly distributed over the range 0 to 4. In practice the limit values were chosen to be 1.33, 2.33, and 3.00 (instead of 1.31, 2.29, and 3.02), after a process of iteration to obtain an array showing no obvious periodicity in either the ξ– or η-direction..

The subroutine described is part of the program **pixel.f** used to generate all the 16x16 arrays used in this thesis, and is called **choose**(p0, p1, p2, p3, pout, x, pflag). pixel.f is listed in Appendix B. The resulting array is plotted in figure 5.15.

This special case of the previous discrete position randomisation scheme gives rise to a subset of the previous ensemble, hereafter referred to as 'special case' arrays for brevity. The special case arrays can be fabricated as matrix-addressed SLMs without pinch-off of the connecting electrodes (see figure 6.4b) for the electrode patterns for the array shown), but to be of use must perform in the simulations in a similar manner to the earlier discrete position randomised arrays.

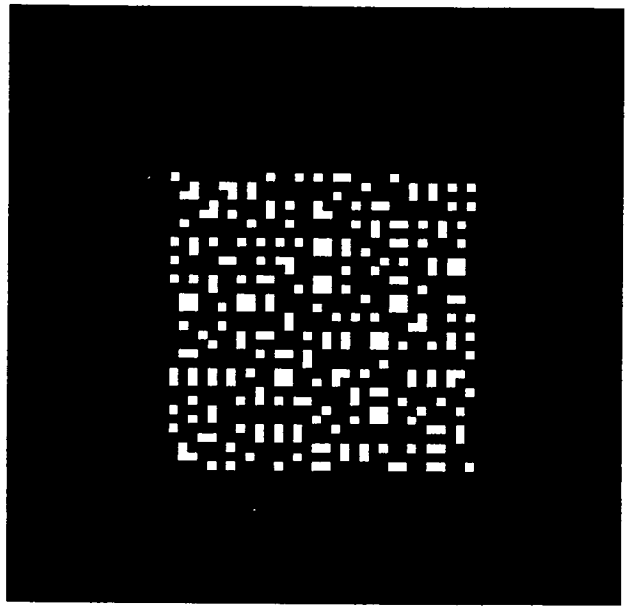

Figure 5.15: 'Special case' version of discrete position randomisation

Simulations for 'special case' randomised arrays

The simulations were carried out as before. 50 different members of the ensemble of 'special case' arrays were generated, and their power spectra calculated and averaged to give the spectrum displayed in figure 5.16a). At this scale, the power spectrum of the single array of figure 5.15 is indistinguishable from that of the ensemble of 50 arrays. The low intensity details of the power spectra of the ensemble and the single array are shown in figure 5.16b) and c), for comparison with the photograph, d), of the power spectrum from the optical simulation. The optical simulation used a transparency photographically reduced from figure 5.15, mounted in a liquid gate..

Figure 5.17a) and b) shows the results of computational simulations of near-critical and undersampling using the single special case array as a Fourier plane filter, with the letter E and large A as input to a 4-f processor. The corresponding optical simulations are shown in figure 5.17c) and d).

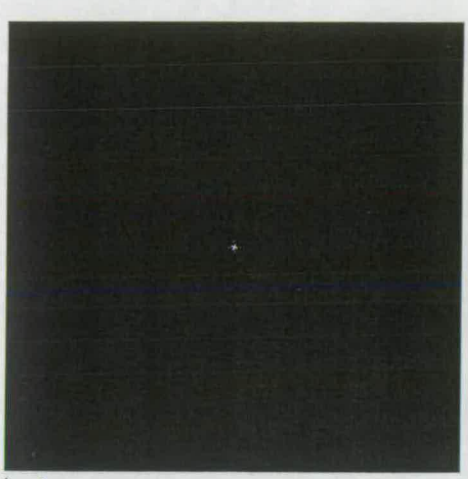

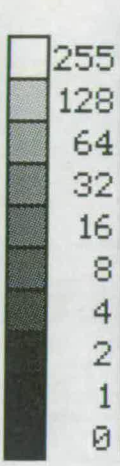

a) Average power spectrum of ensemble of 50 'special case' randomised arrays

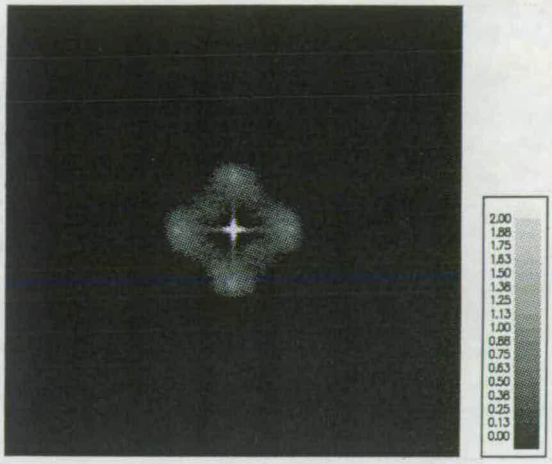

b) Low level detail in power spectrum of ensemble

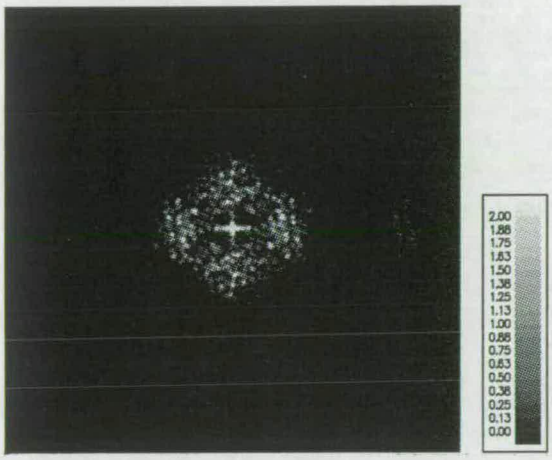

c) Low level detail in power spectrum of single member of ensemble

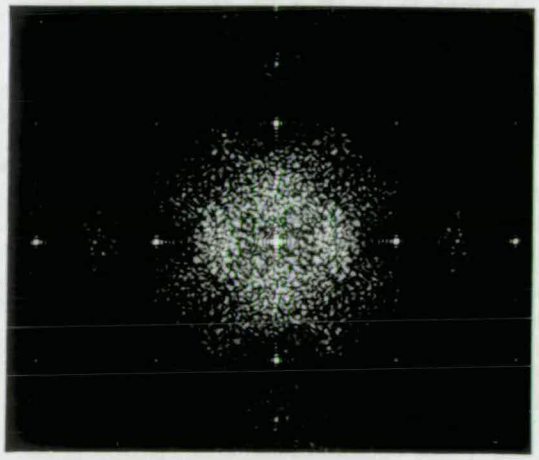

d) Photograph of power spectrum of single member from optical simulation

Figure 5.16: Power spectra simulations for special case randomised pixel array

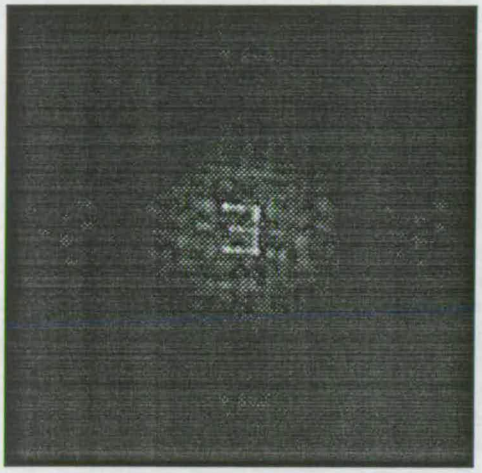

a) Computational simulation for near-critically sampled E

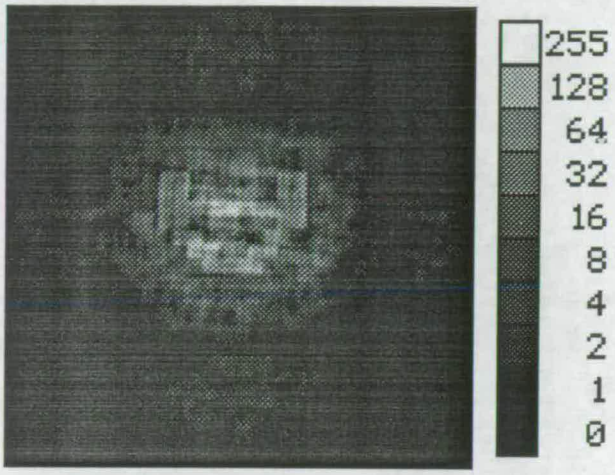

b) Computational simulation for undersampled A

c) Optical simulation for near-critically sampled E

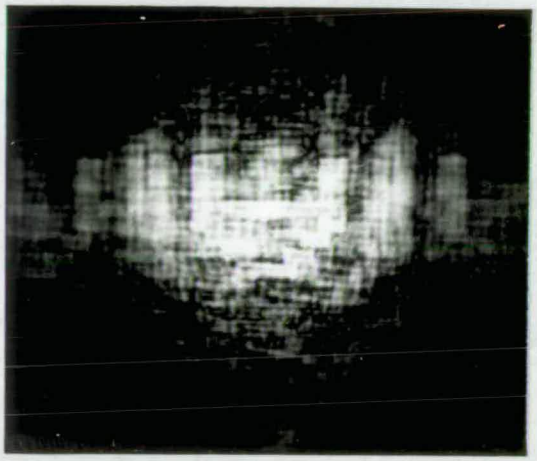

d) Optical simulation for undersampled A

Figure 5.17: Output images illustrating sampling for input letters filtered by special case of discrete position randomised pixellated filter of figure 5.15.

Discussion of results

The power spectra results for the single special case array and the ensemble average of the 50 different arrays are very similar to each other, and as required, to the power spectra of the previous discretely randomised arrays. Studying the background detail of the single array's spectrum it is evident that the background distribution is of a similar level to before, showing the same gross light depletion region structure, but less diffuse. In particular some structure is evident at the positions of the first orders, indicating that the constraints on the pixel positioning have modified the pixel position probability distribution function. This deviation from the model is shown to be nominal when cross-sections along the $x-$axis are taken, shown in figure 5.18– it is apparent that the first order structure is only of the order of 0.5% of the height of the zero order peak (cf 17% for the uniformly randomised array). This is extremely low and indeed could not be distinguished from the background when the array camera was used to establish the relative heights of the spectral orders. Again, tolerance of aliasing was demonstrated by the imaging of the undersampled A.

5.3.4 Additional comments

Other pixel sizes: $\beta \neq \alpha/2$

We consider the discrete position randomisation scheme for pixels with $\beta \geq \alpha/2$. The non-overlapping condition then requires that the random displacements are correspondingly smaller, with the joint probability distribution defined as

$$\begin{aligned} P_{\phi_{n,m},\psi_{n,m}}(\phi,\psi) \;=\; & \frac{1}{4}\,\delta(\phi-\frac{\alpha-\beta}{2},\psi-\frac{\alpha-\beta}{2})+\frac{1}{4}\,\delta(\phi-\frac{\alpha-\beta}{2},\psi+\frac{\alpha-\beta}{2}) \\ & +\frac{1}{4}\,\delta(\phi+\frac{\alpha-\beta}{2},\psi-\frac{\alpha-\beta}{2})+\frac{1}{4}\,\delta(\phi+\frac{\alpha-\beta}{2},\psi+\frac{\alpha-\beta}{2}) \end{aligned} \tag{5.73}$$

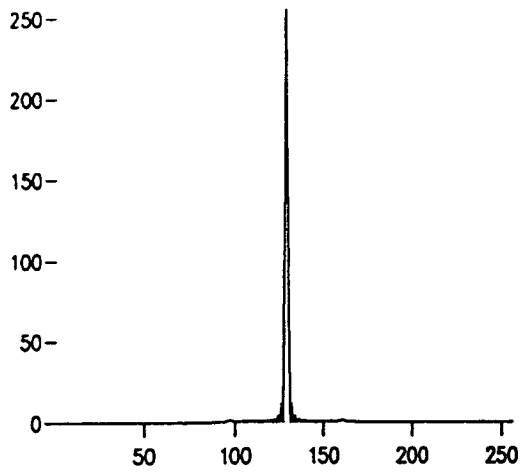

a) Cross-section along $x-$axis for ensemble average

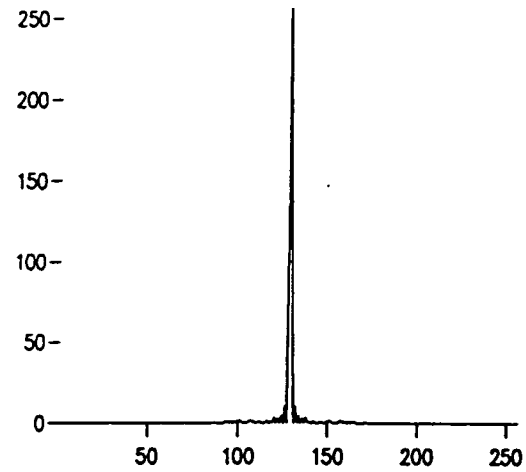

b) Cross-section along $x-$axis for individual member

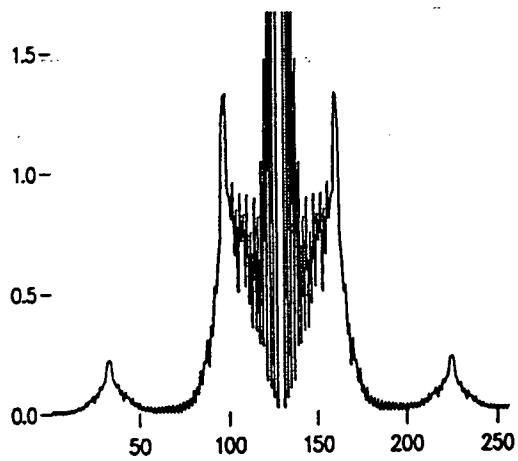

c) Low level detail of a) showing light depletion region

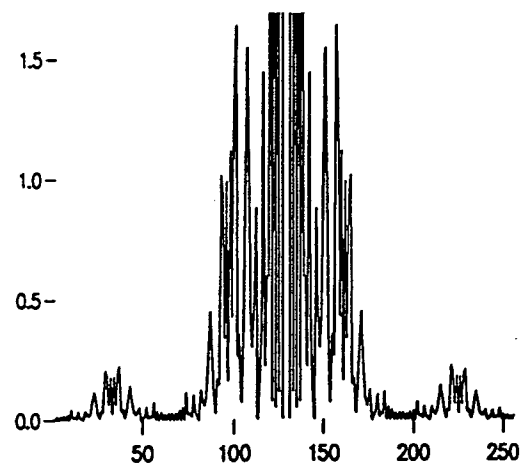

d) Low level detail of b)

Figure 5.18: Power spectra cross-sections for the special case randomised array

β/α	Relative intensities of first orders		
	Regular case[1]	Uniform random distbn.[2]	Discrete random distbn.[3]
0.95	0.0082	0.0081	0.0079
0.90	0.0119	0.0115	0.0108
0.80	0.0547	0.0479	0.0358
0.70	0.1353	0.0997	0.0467
0.60	0.2546	0.1458	0.0243
0.50	0.4052	0.1642	0.0000

[1]Function attenuating spectral orders is $\text{sinc}^2(\beta x)$.

[2]Function attenuating spectral orders is $\text{sinc}^2(\beta x)\text{sinc}^2((\alpha-\beta)x)$.

[3]Function attenuating spectral orders is $\text{sinc}^2(\beta x)\cos^2(\pi(\alpha-\beta)x)$.

Table 5.2: Comparison of first order intensities for regular and randomised arrays, relative to normalised zero order.

and $\rho(x) = \cos^2(\pi(\alpha-\beta)x)$, $\rho(y) = \cos^2(\pi(\alpha-\beta)y)$. While this will not eliminate the first order unless $\beta = \alpha/2$, the first order will still be attenuated compared to that of the regularly spaced array with the same shape and size of pixel. This is shown in Table 5.2.

Table 5.2 shows a direct comparison between the power spectra of the regular and randomised arrays described, by tabulating the intensity of the spectral order nearest the origin for various ratios of β:α. The table shows how both the randomisation schemes described reduce the magnitude of the first spectral order compared to that of the regular array. As intimated earlier, the choice of $\beta = \alpha/2$ to eliminate the higher even spectral orders is not the optimum choice for attenuating the odd orders when using the uniformly randomised array of section 5.4.1, and in actual fact gives the least possible attenuation for that randomisation scheme.

In every case the discrete case randomisation scheme attenuates the first order the most, with all the higher orders being eliminated for $\beta = \alpha/2$. If $\beta < \alpha/2$ the odd orders can still be eliminated if the allowed position for the pixel centres is as for $\beta = \alpha/2$, i.e. $\phi_{n,m}, \psi_{n,m} = \alpha/4$ or $-\alpha/4$. This is important if the area that the pixel is allowed to lie in is slightly less than $\alpha \times \alpha$ (as is necessary to prevent electrical short-circuits between adjacent pixels when using a matrix-addressed SLM) thus preventing the size of the pixel from being as large as $\alpha/2$. If ϵ is the separation necessary to prevent short-circuits between adjacent pixels

then consequently the transformed pixel function will become $P(x) = (\alpha/2 - \epsilon)\text{sinc}(\alpha/2 - \epsilon)x$ and the higher even orders will be eliminated– however, so long as ϵ is small compared to α, the effect of any spectral order higher than the first should be small due to the decreasing $P(x)$ being small-valued at their positions.

Effect of light depletion region on the zero order

Even for the extremely low resolution case of a 16x16 array, the effect of the light depletion zones which arise in the random array's spectrum is not intrusive, becoming apparent only on closer inspection of the spectra. As noted the magnitudes of the zero orders are the same, for the same total transmissive area. The analysis also shows that as the magnitude of the light depletion zone is zero at the origin, and rises slowly with respect to the width of the zero order in the random examples described, the zero order will effectively be unchanged, leading to negligible loss of sharpness in the central image.

5.4 Conclusions

Replication due to the underlying structure of a pixellated SLM has been identified as a problem in optical correlation, in the form of false correlations arising due to aliasing when the SBP of the Fourier plane filter is less than that of the input. This situation will arise if the SLM has a SBP matched to that of a relatively small target in the overall input scene, which is desirable in terms of SLM frame rate and economy. Randomisation of the pixel positions has been proposed as a solution, and has been shown to be effective in removing unwanted spectral orders, and redistributing their energy into a diffuse low-level background illumination which falls to zero around the zero order; the zero order itself is unchanged in magnitude for a given pixel size. The power spectra of randomised arrays has been analysed and used to predict a 'discrete position' randomisation scheme, which results in the removal of the first spectral orders for pixel sizes half the pixel pitch or less. For the case where the pixel size is exactly half of the pixel pitch all distinct spectral

orders other than the zero order may be eliminated. A special case of the discrete position randomisation has been identified which is suitable for fabrication as a matrix addressed SLM. This is described in the following chapter, and used in the subsequent correlation experiments of chapter 7.

Chapter 6

Construction of a transmissive SLM

Spatial light modulators were constructed to demonstrate and implement optical applications of Fourier plane filters with randomly positioned pixels. The spatial light modulators which were constructed were transmissive, matrix addressed liquid crystal devices with an individual pixel's 'off' or 'on' state selected using a multiplexed drive scheme. Both phase only and amplitude only binary SLMs were constructed, for use in the correlation experiments described in the next chapter. This chapter explains the reasons for choosing this kind of device and details their construction, and the choice and operation of the multiplexed drive.

6.1 Basic design of SLM

The basic design of the SLM is as shown in figure 6.1, and will be described in more detail in the rest of the chapter.

The light modulating layer is a 12 μm layer of nematic liquid crystal (LC) sandwiched between two glass substrates (1) separated by polyester spacers (2). The liquid crystal is aligned between the surfaces with either a 90^0 twist (for an amplitude modulator) or a 0^0 twist (for a phase-only modulator) by alignment layers of rubbed PVA (3) on the in-facing sides of the glass. In the amplitude modulat-

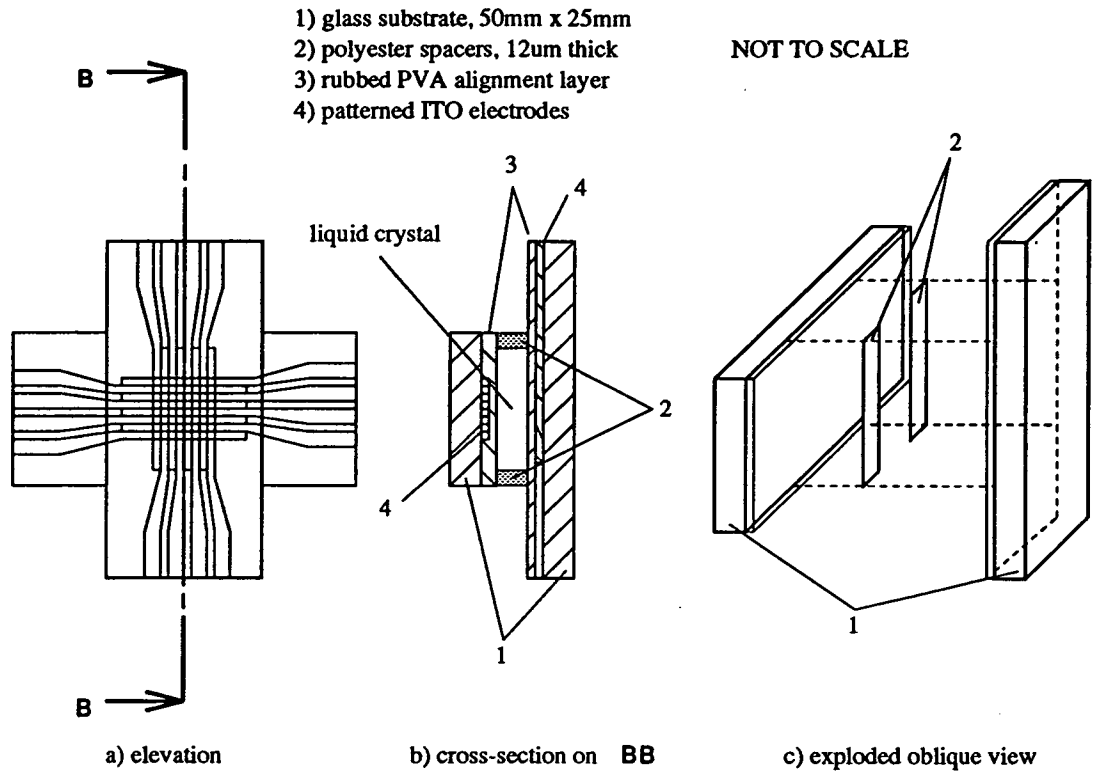

Figure 6.1: Basic design of SLM

ing configuration polarisers are mounted either side of the SLM, with their axes aligned with the rubbing direction of the PVA on the glass plate nearest the first polariser. In the phase-only configuration only the polariser on the input side of the SLM is necessary. Transparent conducting indium tin oxide (ITO) electrodes (4) are patterned onto the glass surfaces using lithography and sputtering before application of the PVA, with the pixels defined by the overlap of the two sets of electrodes when the device is assembled.

Addressing individual pixels makes use of the non-linear optical response of the liquid crystal to applied voltage described in Chapter 3, whereby the liquid crystal does not respond until the voltage across it reaches some threshold value V_{th}. The electrical signals of the multiplexed drive scheme are supplied to the SLM by an interface/waveform generator which is connected to the SLM by printed circuit board (PCB) edge connectors. The interface receives logic control signals and pattern information from a BBC microcomputer running a control program

	Nematic	FLC
Cell thickness	$\sim 12\mu$m	$\sim 2\mu$m
Alignment treatment	Rubbed PVA or evaporated SiO_2	Rubbed polymer or evaporated SiO_2
Drive scheme	Multiplexed, 3 or 4 waveforms	Multiplexed, 4 waveforms
Cell filling	Slowly, under vacuum at room temperature	Slowly, under vacuum, cooling from elevated temperature

Table 6.1: Typical differences in constuction of FLC and nematic LC devices.

written in BASIC.

6.2 Choice of matrix addressed transmissive configuration

The configuration described above was chosen for several reasons, but primarily because of the wide range of devices that could be made using it, once the necessary technology had been mastered. Devices using widely differing types of liquid crystal with correspondingly different electrooptic responses could be made using the same prepared substrates: typical differences in the construction of ferroelectric and nematic liquid crystal devices are described in Table 6.1, which details differences of procedure rather than technology. In both cases the devices' fill factors, pixel dimensions, spacings, shapes and positions are determined by the electrode patterns on the glass substrates which can be varied by changing a single photomask.

The technology used to make the devices is well proven [65], and sufficiently simple that the devices could be made in-house without recourse to outside specialists such as a silicon foundry, which is necessary with other SLMs designed at the University of Edinburgh. These two points meant that the progress of the devices could be monitored throughout their construction and engendered confidence at the outset they could be made to work within a reasonable time scale, with the time for developing successive generations of device determined mainly by the

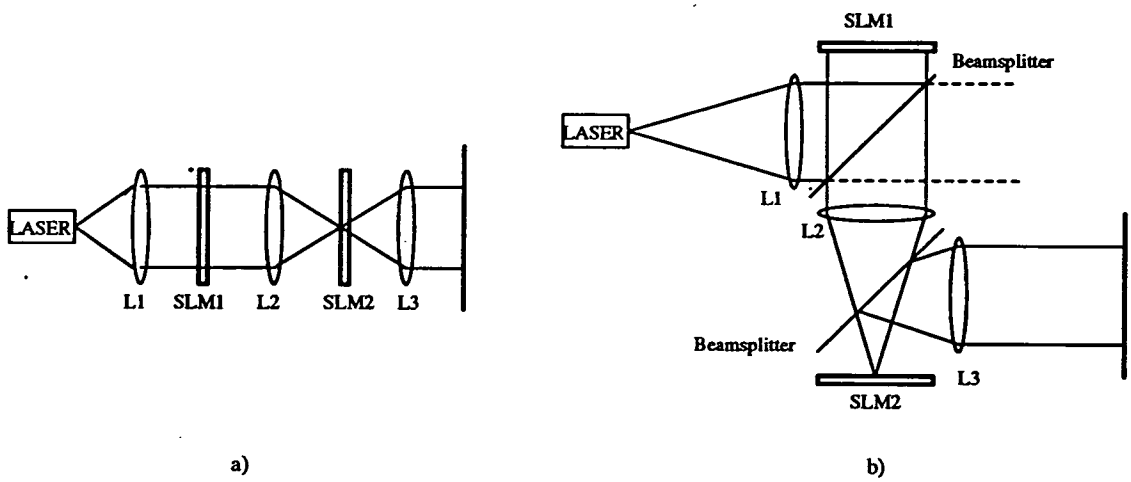

Figure 6.2: Optical processor architecture, using a) transmissive SLMs and b) reflective SLMs

time taken to design and make a new photomask.

Transmission was chosen as the favoured mode of operation due to the simpler processor architectures possible compared to those using reflective devices. Figure 6.2 shows optical processors with SLMs in the input and Fourier planes, a) using transmissive SLMs and b) using reflective liquid crystal based SLMs. The difference in complexity is manifest, even with the omission of polarisers and analysers. The reflective system also suffers from unwanted reflections from the front and back surfaces of the LC cover glass/counter electrode which can swamp the desired 'signal' reflection from the mirrored surface(s) of the reflective SLM. In principle these reflections could be removed by anti-reflection coating the surfaces of the cover glass, but in the two-SLM system described unwanted reflections between the SLMs would remain, originating from the pixellated mirrors of the SLMs.

The transmissive SLM also has the advantage over the reflective SLM of greater tolerance to thickness variation of the LC layer. When the SLM is configured as a phase modulator the thickness and uniformity of thickness of the LC layer are particularly crucial, with a deviation from the desired thickness leading to a phase error proportional to the deviation: using a reflective SLM the effect of the deviation is doubled due to the LC layer being traversed twice.

Matrix addressing is necessary to reduce the number of connections between the substrate and drive electronics to a manageable level. Matrix addressing of a QxQ array of pixels reqires $2Q$ external connections and necessitates the use of a multiplexed drive scheme and LC with a sharp threshold characteristic; directly addressing each pixel relaxes the LC threshold condition but requires Q^2 external connections. The previous chapter showed that often it is desirable to have a filter SLM with a relatively low space bandwidth product, and so the devices constructed were arrays of only 16×16 pixels.

6.3 Construction of the device

Once the basic design of the SLM has been decided the process of detailed design and construction proceeds in a fairly linear fashion, as shown in the flowchart in figure 6.3, with the various stages described in the subsections indicated.

6.3.1 Design and fabrication of the photomask

The overall dimensions and conducting areas for electrical connection at the edges were common to all the photomasks fabricated. The substrate size was chosen to be 50mm by 25mm due to the availability of optical flats in dimensions which are multiples of this, and of microscope slides in these dimensions which were used for test cells and prototypes. This allowed a pixel array up to approximately 20mm square to be designed, and edge connection to be performed using standard 0.1" pitch PCB edge connectors, taking off 8 electrode connections at each end of the substrate. The electrodes were interdigitised to keep them as straight as possible and minimise resistance differences arising from variations in length.

Typical mask designs are shown in figure 6.4 with each design composed of patterns for the top and bottom sets of electrodes placed side by side. Plotfiles for the designs were generated by a specially written FORTRAN program (**ranpix**), and plotted at a magnification of x5 on a high quality pen plotter. The program was

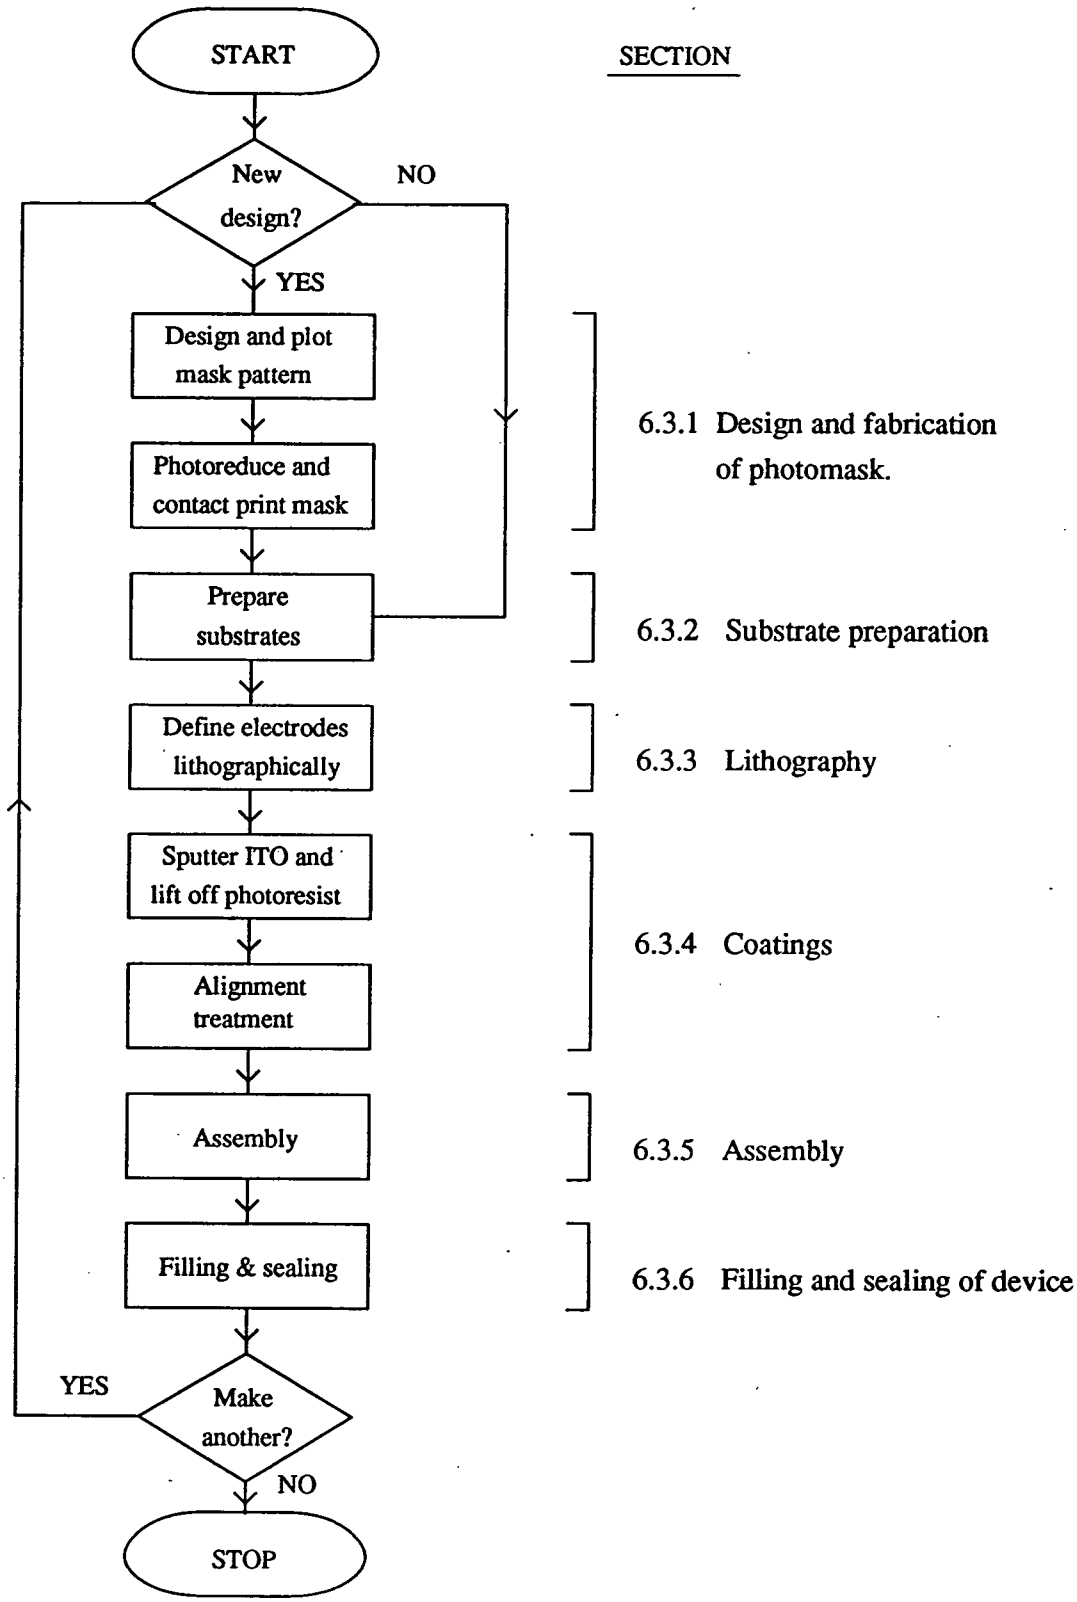

Figure 6.3: Construction process for a transmissive SLM

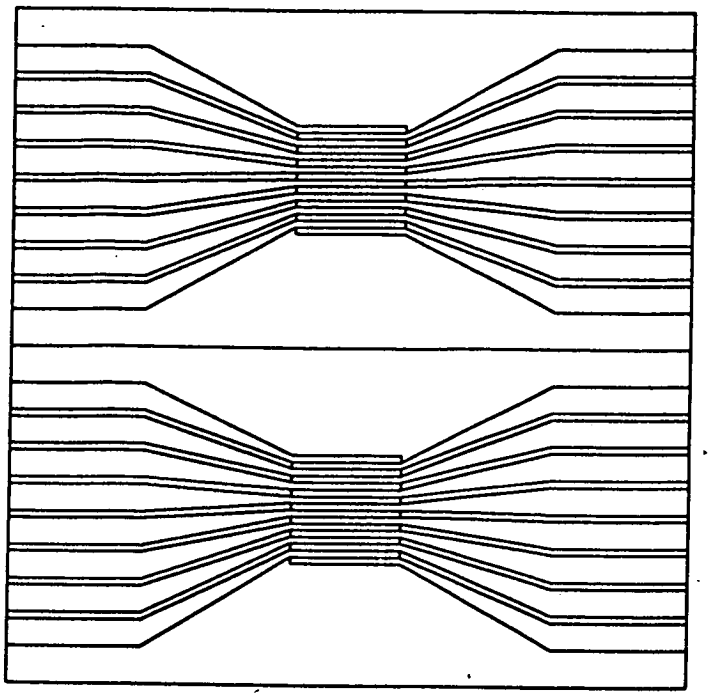

a) simple regular array

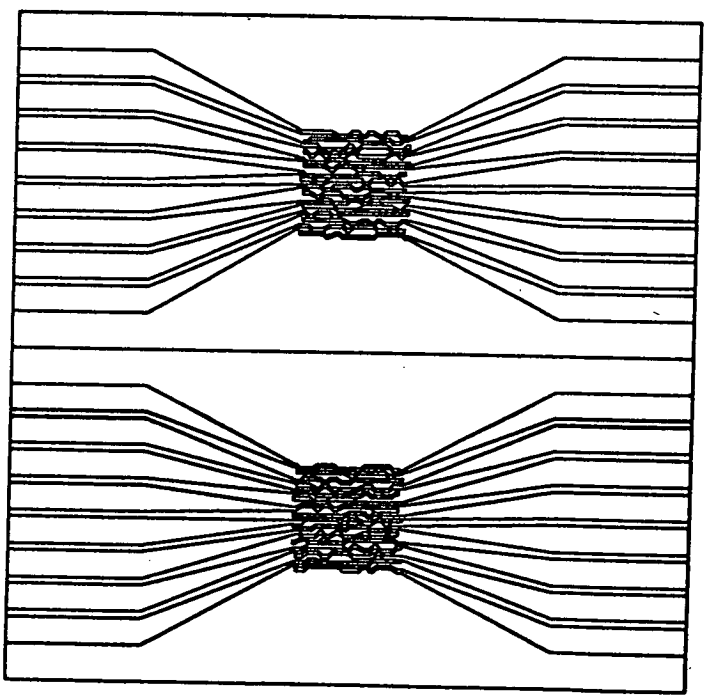

b) Random pixel array used in later correlation experiments

Figure 6.4: Typical mask designs

written to be flexible, generating 16x16 arrays of varying fill-factor, size and pixel-position randomisation scheme depending on the responses to a series of menus. The program also accepts pixel position information from the program **pixel** used to generate the arrays in Chapters 4 and 5, requiring only coordinates of the top right hand corner of each pixel. Scaling and interconnection of the pixels and electrodes is carried out by **ranpix**, along with the 90^0 rotation + reflection + translation of coordinates necessary to plot the top electrode pattern beside the bottom electrode pattern.

The pair of top and bottom electrode patterns were plotted side by side to reduce the possibility of scaling differences betweeen top and bottom in the subsequent stage when the design was photoreduced to $\times 1$. The designs were plotted oversize to retain detail of the electrode pattern, and to reduce the thickness of the plotted lines relative to the size of the pixel area to a minimum, while maintaining them thick enough to be photoreduced successfully.

Photoreduction was done onto Kodak Kodalith film to yield a high contrast negative on a celluloid backing. This negative was then contact printed onto a 4"x4" high resolution plate coated with Agfa HR 1A emulsion, and developed in Kodak HRP developer to yield a finished mask, after washing and drying. Contact printing onto the photographic plate is necessary as a rigid mask compatible with a Carl Süss mask aligner was required. The design was not photoreduced directly onto a high resolution plate due to the extremely low sensitivity of the emulsion, which resulted in underexposure of the plates due to low intensity reciprocity failure. Contact printing onto the plates initially also proved difficult until a high intensity flashgun was used for the exposure.

6.3.2 Substrate preparation

The substrates must first be cut to size, either resulting in two 50mm x 25mm glass plates, or more preferably a single 50mm x 50mm plate, which is easier to handle and register with the photomask. It is a simple task to split the 50mm x 50mm plate into separate top and bottom electrode patterns at a later stage, following

lift-off of the photoresist. Microscope slides were used for the prototype SLMs, but a better quality substrate (in terms of surface defects, flatness) was desired for the SLMs to be used in correlation experiments. Glass wafers corresponding to the dimensions of 4" diameter silicon wafers showed promise due to their good optical quality and compatibility with lithography equipment used for silicon. In practice the glass wafers proved almost impossible to cut to size, were too flexible to maintain a uniform layer of liquid crystal in a finished device, and caused problems when using automated lithography equipment which relied on opaque (to visible light) substrates interrupting beams of light to control progress along the wafer track. Utility grade optical flats with flatness guaranteed to $\lambda/4$ were also tried, but their 6mm thickness led to difficulties in making external electrical connections to the electrodes. High resolution plates with the emulsion stripped off were used as a compromise - the plates used were Ultra Flat grade from Precision Photoglass guaranteed to have a bow of less than 5μm over the diagonal of a 3" square plate. Approximating the bow to an arc of a circle, this would scale down to approximately $\lambda/10$ over the diagonal of the 8mm x 8mm arrays used in the later correlation experiments– although this seems optimistic, interferometric measurements suggest that the flatness is very good. Also, at 0.060" ($\sim$ 1.5mm) thick the plates are rigid enough to ensure a uniform LC layer, and are compatible with standard PCB edge connectors for connections to the electrodes.

Before further processing the glass is cleaned rigorously according to the following procedure:

1. Washed under a tap with a nailbrush and Teepol.
2. Rinsed in tap water and allowed to dry.
3. Placed in a plastic carrier and left in hot chromic acid at approx 80 ^{o}C for 3-4 hours.
4. Removed from the acid and rinsed in deionised water for $>$ 2 hours.
5. Immersed in a beaker of unused Analar acetone to remove the water.
6. Immersed in a fresh beaker of unused Analar acetone for the same reason.

7. Allowed to dry in clean air in a laminar flow cabinet.

The glass substrate is now ready for lithography, if making an SLM, or for sputtering with ITO, if making a simple LC test cell.

6.3.3 Lithography

The electrode pattern is defined lithographically [96] in such a manner as to leave raised areas of photoresist corresponding to the opaque areas on the photomask. This process is shown in parts a) to c) of figure 6.5. Photoresist must first be spun onto the substrate. After the liquid photoresist is dropped on, the substrate is spun at 4000rpm for 20 seconds to leave a uniform coating of resist over the surface, approximately 2μm thick. The photoresist used was a positive resist, Shipley Microposit 1880 [97]. The coated substrate is then baked at 100^{o}C for 1 hour to remove excess solvent from the photoresist. Exposure follows, using contact printing. This is carried out using a Carl Süss mask aligner/contact printer with the substrate held in place on a vacuum chuck, then pressed firmly against the photomask. The photoresist is then exposed to parallel UV light through the photomask and developed in Shipley Microposit 351 developer [98], diluted to the strength of 4 parts deionised water to 1 part developer, for one minute. The substrate with developed photoresist is then washed in deionised water and left to dry in preparation for sputtering with ITO.

6.3.4 Coatings

Several coatings must be applied to the substrates before assembly into a spatial light modulator. In all the devices made, it is necessary to apply a sputtered coating of indium tin oxide (ITO) to form transparent conducting electrodes, and an alignment layer for the liquid crystal. This is shown in stages d) to h) of figure 6.5.

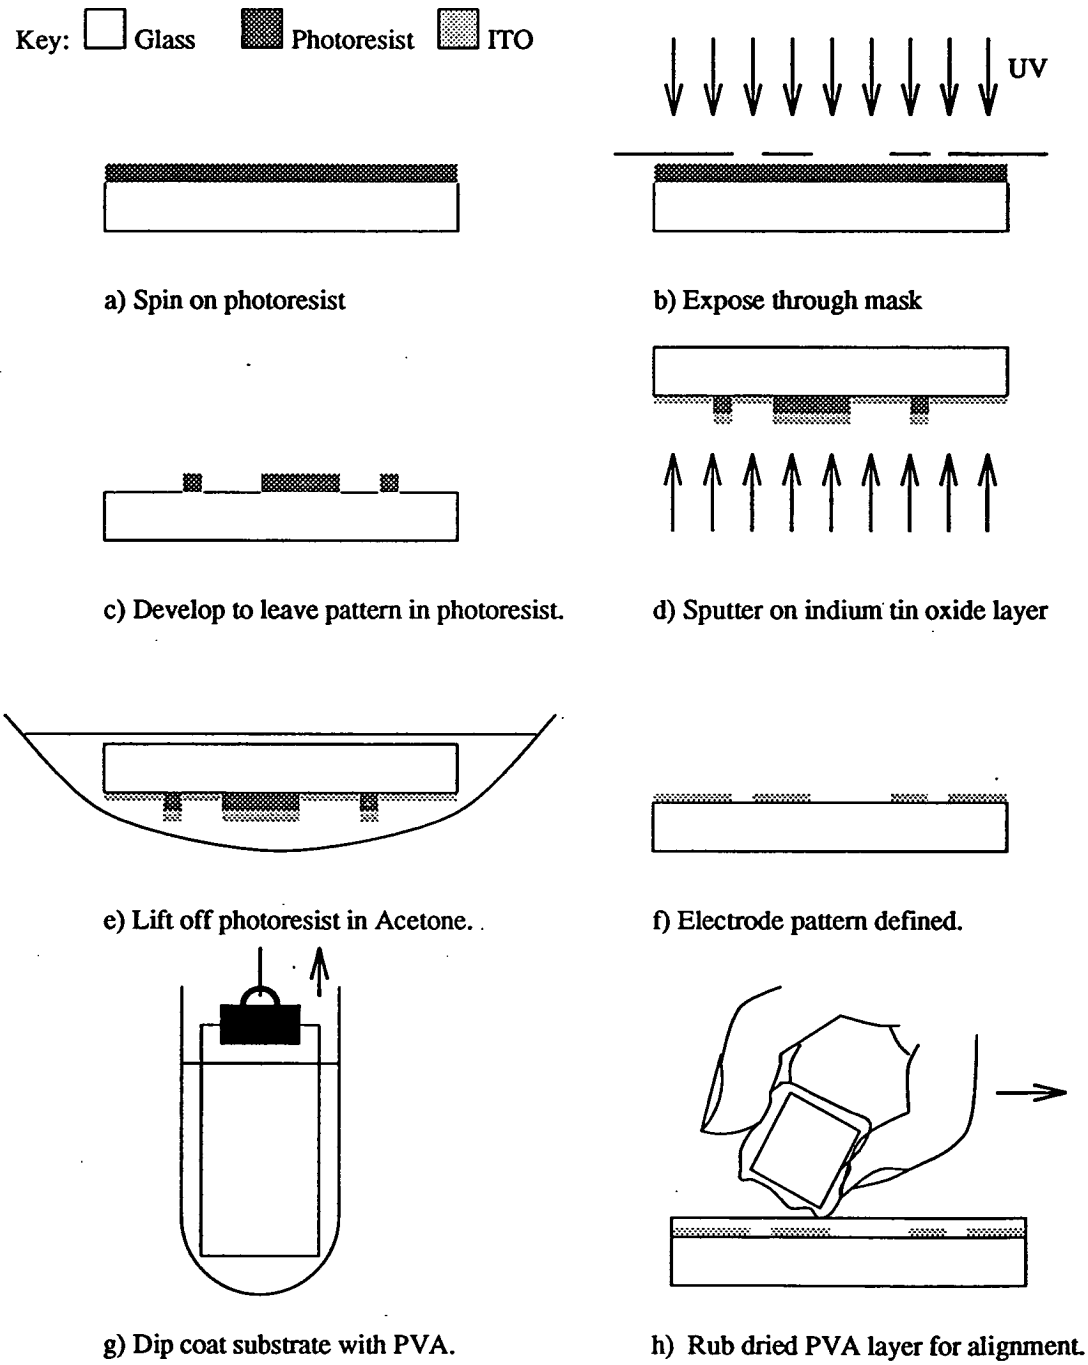

Figure 6.5: Illustration of lithography and coating stages

Sputtering of ITO

The ITO layers are deposited using the dc magnetron reactive sputtering technique in the department's Techniques Laboratory [99]. For this the substrate is mounted face down in a vacuum system, 132mm above a target of 90% indium-10% tin alloy. Sputtering takes place in a tenuous argon/oxygen atmosphere at a total pressure of $\leq 8.5 \times 10^{-3}$ torr with the target covered by a shutter until conditions in the vacuum chamber have stabilised, with the oxygen and argon gas flow rates at the desired values. After a few minutes the shutter can be opened and sputtering onto the substrate commences, continuing for a further ten minutes. This should give a film approximately 400 Å thick, nearly optically transparent, with a measured resistance of 200-300 Ω, as measured between the probes of an AVO meter 1 cm apart.

Lift-off of photoresist

The photoresist used to define the boundaries of the electrodes must be removed before further coating of the substrates. This constitutes a 'lift-off' process [100, p304, 364] whereby the photoresist is removed along with the sputtered coating to leave uncoated regions defining the electrodes. When removing the photoresist the substrate is placed face down in a watchglass and immersed in the resist stripper to prevent the lifted-off ITO slumping back onto the substrate and potentially drying out, causing short circuits. The resist stripper used was Analar acetone, and the stripping operation was carried out in an ultrasonic bath for greater speed and efficiency. Following drying, the substrates are ready to be prepared for the next coating, that of the alignment layer.

Alignment treatment

A rubbed layer of PVA was used for alignment purposes in preference to evaporated silicon monoxide. Earlier experiments had shown little visible difference in

maximum contrast, and rubbed PVA alignment could be effected more quickly and conveniently.

Before the alignment layer is applied, the substrate is split into its two halves, and cleaned to remove any organic residues of the acetone. Cleaning is done by scrubbing with a cotton bud following immersion in a dilute solution of Decon 90, rinsing in deionised water and leaving to dry.

The actual coating of the surfaces is carried out by dipping the substrates into a large test tube filled with a solution of PVA (1g of PVA, molecular wt. 125,000, in 400ml of water). The substrate is then drawn slowly out of the test tube, leaving a uniform layer of PVA over the surface, which is then left to dry in air. If any organic residues have been left on the surface of the substrate the PVA will have difficulty in adhering to the surface and will not leave a uniform layer. After the PVA has dried, the surface is rubbed by a lens tissue doubled over and wrapped around a smooth flat metal block. Applying pressure similar to that applied to a ballpoint pen when writing, the block and lens tissue is drawn slowly over the surface in a straight line. This is repeated another 4 times, with the rubbing direction the same for all 5 strokes. The rubbing direction is noted. The procedure is carried out for both halves, with the relative rubbing directions such that, when assembled, the rubbing directions on the top and bottom halves will be perpendicular for an amplitude modulator, or parallel for a phase modulator.

Further consideration must be given to the rubbing direction for an amplitude modulator if the liquid crystal used contains a chiral dopant. The chiral dopant is added to remove the degeneracy which exists between the left and right handed 90^0 twist states, ensuring that only one twist state is present in the aligned liquid crystal cell, and removing the optical defects that occur at the interface of regions of opposite twist. However, the rubbed PVA alignment layer also favours one handedness of twist dependent on the relevant orientations of the rubbing direction, due to a tilt induced in the liquid crystal molecules at its surface. For the liquid crystal mixture used here, BDH E300 + 0.05% CB15, this means that when viewing the assembled device, the rubbing direction of the top plate must be from north to south and the bottom plate from west to east (in an obvious

notation), for conflict between preferred twist directions to be avoided [101].

Opaque mask layer for phase modulator

In the devices configured as phase modulators it was also necessary to apply an opaque layer of aluminium to the back side of one of the substrates to mask off the areas of liquid crystal which could not be switched. The switchable pixel regions were defined as clear regions in this layer using the lithography techniques described above i.e. generating a mask matching the pixel areas, coating the back of one of the glass substrates with photoresist, registering the mask relative to the pixels using the mask aligner, exposing and developing the photoresist . The aluminium layer was evaporated onto the surface in a vacuum, to a thickness of approximately 2000Å. The areas corresponding to the pixels were then removed using the lift-off process described earlier.

6.3.5 Assembly

The two halves of the SLM must now be brought together with spacers of suitable thickness separating them. The spacers are cut from 12μm thick polyester sheeting into strips approximately 3mm wide, which are placed onto the bottom half as shown in figure 6.6a). The top half is brought carefully into position, and adjusted until the overlap of the electrodes define the pixels as required.

With the top and bottom plates registered correctly, using a microscope for viewing if necessary, the two halves are clamped lightly together. Four crocodile clips are used for clamping, with their points of contact with the glass over the spacers to minimise stress-induced bowing of the plates. The two halves are now glued together using UV curing adhesive which is spread thinly along all four of the overlapping sides, except for a small open area at a corner as shown in figure 6.6. Liquid crystal enters through this gap to fill the 12μm thick cell between the plates in the subsequent vacuum filling stage. Note that the gap is orientated

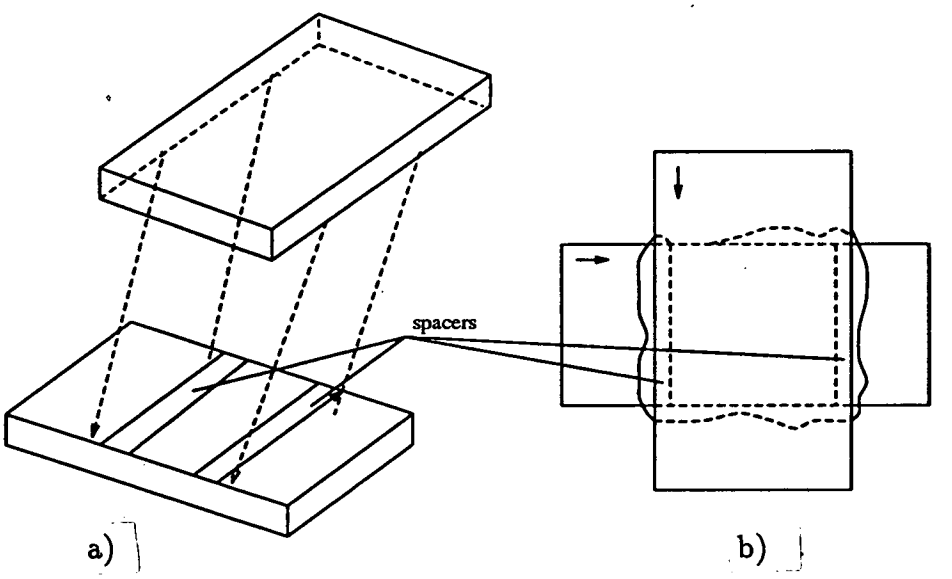

Figure 6.6: Assembly of SLM, showing spacers and opening for vacuum filling

with respect to the rubbing directions. The liquid crystal approximately follows the rubbing direction as it fills the cell, starting from the open corner.

6.3.6 Filling and sealing of device

The cell is filled in a vacuum chamber by a drop of liquid crystal which is initially at the end of a probe, attached to a rotating mechanical feed-through, as depicted in figure 6.7. The vacuum chamber is pumped down for several hours to outgas both the liquid crystal and the cell, getting down to a pressure $\leq 10^{-4}$ torr. After outgassing, the drop of liquid crystal is placed at the open corner of the cell, using the rotating feed through, while maintaining the vacuum. The cell is allowed to fill by capillary action until it is approximately half full. At this point the chamber is slowly let back up to atmospheric pressure, forcing the liquid crystal to fill the vacuum remaining in the cell. The entire filling process is carried out at room temperature and results in a well-aligned, defect-free, nematic liquid crystal cell. After filling is complete, the excess liquid crystal is removed, and the open corner sealed using 'Araldite Rapid' epoxy adhesive to complete the construction of the liquid crystal device.

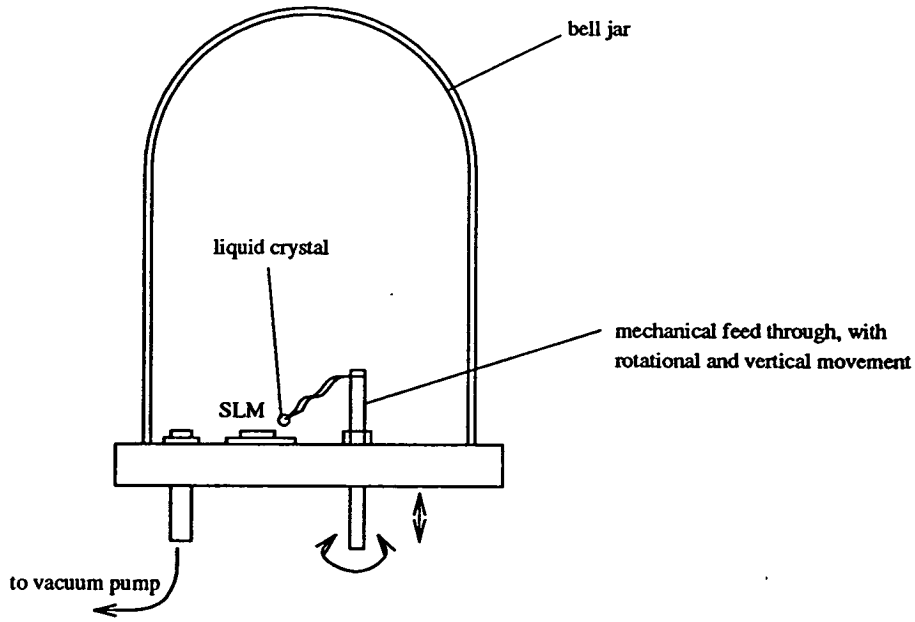

Figure 6.7: System for vacuum filling of liquid crystal cell

6.4 Interfacing

An interface is necessary to implement the chosen multiplexed drive scheme and transfer the frame information (which pixels are off/on) to the spatial light modulator from a microcomputer, which controls the operation of the interface.

The basic principle of the drive scheme is that row and column waveforms can be specified which maintain all the pixels below threshold except the pixels which are required to be switched 'on' in a particular column, with each column being addressed in turn. This scheme is clearly of a dynamic nature, relying on each column being addressed at a high enough frequency to prevent the 'on' pixels from relaxing to the 'off' state by the time they are next addressed. A further consideration is that the net d.c. voltage across the liquid crystal integrated over a longish time interval (e.g. the time taken to write all the pixels) should come to zero, to prevent the liquid crystal degrading electrolytically over time.

6.4.1 Interface for prototype devices

The first prototype devices built during this project were addressed by an interface built using simple digital ICs and buffers. This interface was designed to implement the drive scheme of Alt and Pleshko described in chapter 3, with the frame information and timing being supplied from a BBC microcomputer via its user port and 1MHz bus.

An interface for driving a 16x16 silicon backplane device existed already but was considered to be too slow for dynamically addressing an array, as it required 80 bytes of data per frame. Most of these bytes were concerned with the interface's timing and control of data routing through it, so in the new interface most of these functions were incorporated on it. This resulted in a frame being addressable using 33 bytes: 32 bytes of frame data (1 bit per pixel), and 1 control byte to reset a counter after each frame. With the data output routine on the BBC micro written in assembly language (by Andrew Garrie) it takes approximately 25ms to write a frame, i.e. a frame refresh rate of about 4kHz. The block diagram of the interface shown in figure 6.8 helps describe its operation.

The fundamental timing of the circuit is carried out by the strobe pulse (STB) from the 1MHz bus which accompanies each data byte as it appears on the 8 lines of the user port. These bytes appear at R1 and R2 and the first is clocked through R1, the second through R2, the third through R1, the fourth through R2 etc. After a pair of bytes has been toggled through R1 and R2 to be set up at R3 and R4, R3 and R4 are clocked simultaneously, with the circuit's timing ensuring that R1 and R2 are not transparent when this occurs. The TTL logic levels appearing at the outputs of R3 and R4 are scaled using CMOS NAND's (type 4011B) which drive the rows. The requirement that the net d.c. voltage be zero is satisfied by making the output from the NAND's bipolar. Therefore a logic one on the input of the NAND appears as $-V_d$ on output, and a logic zero as $+V_d$. This is accomplished using separate power rails for the NAND's, and should not strictly be possible, given the standard noise immunity of CMOS. However, it works.

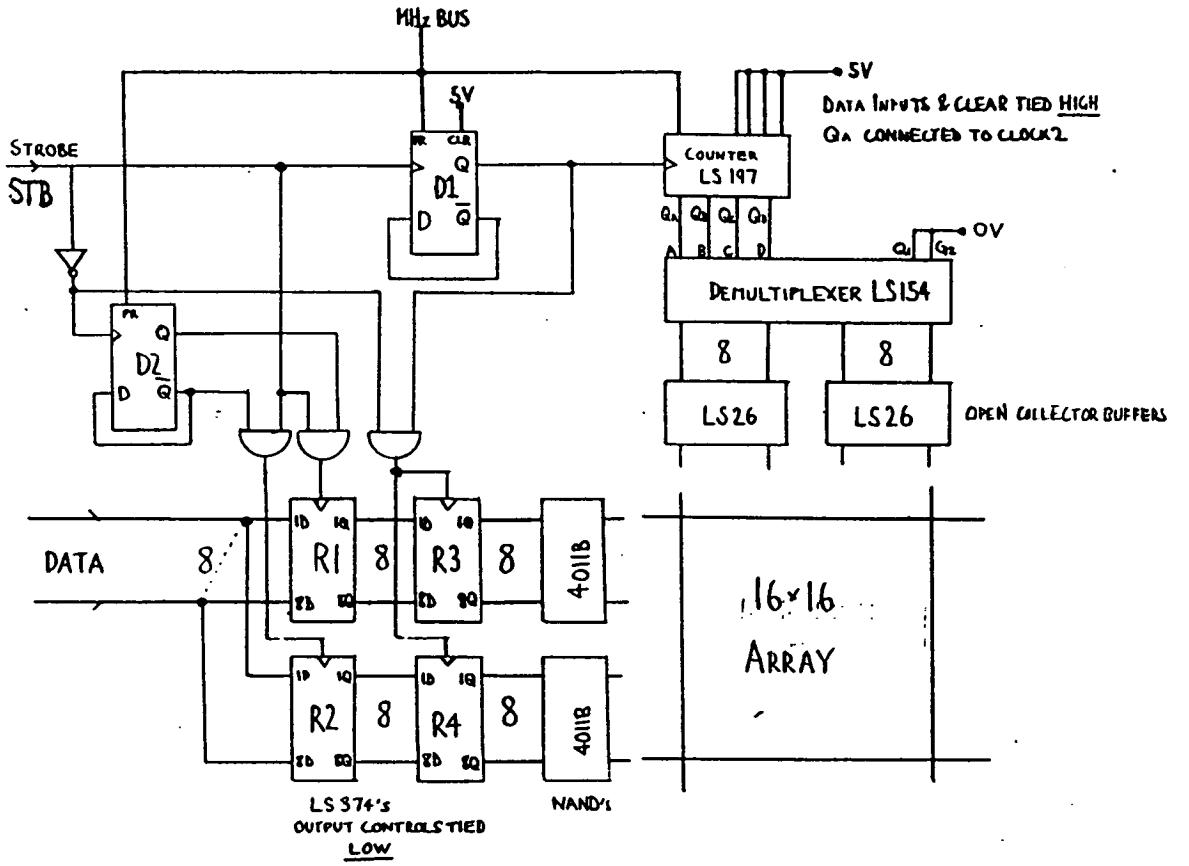

Figure 6.8: Block diagram of prototype interface

Addressing the rows is much simpler. The control byte which precedes every frame resets the counter (type LS197) and flip-flop D1. D1 halves the clock frequency so that the counter advances after every two bytes of data as required. The demultiplexer takes the counter output and steps through the columns as the clock advances, setting the selected column to logic zero and the other fifteen to logic one. This output is scaled and inverted using open collector buffers, so that logic one on input becomes $0V$ and logic zero becomes $+V_s$. This is made explicit in the timing diagram, figure 6.9.

This interface was adequate for testing the prototype devices, but performance suffered from the strobe data being unipolar, resulting in a net d.c. voltage over the liquid crystal. This net d.c. voltage could be eliminated by adding a bias voltage to the columns to optimise the prototype SLM's contrast ratio. However,

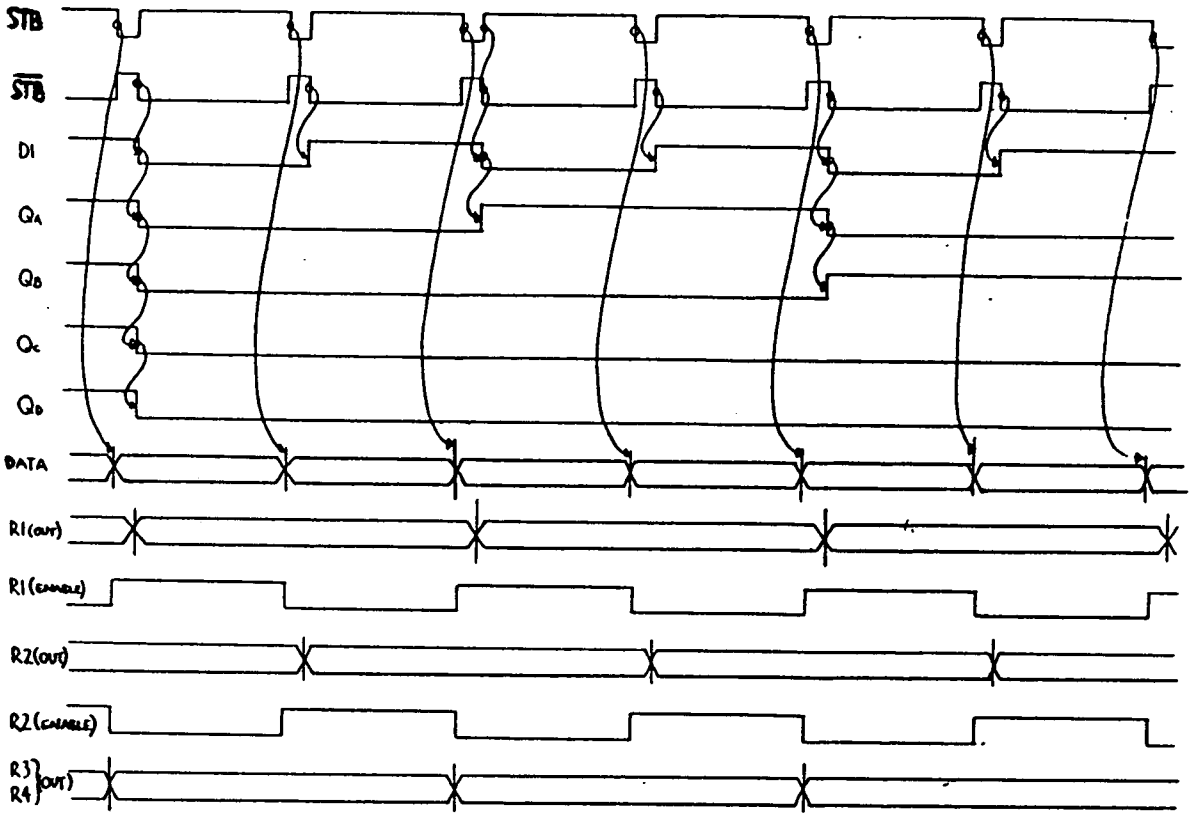

Figure 6.9: Timing diagram for prototype interface. The labels refer to the block diagram of figure 6.8.

the bias voltage was not a constant, but depended strongly on the pattern being written to the SLM. Thus, although any pattern could be displayed on the SLM with good contrast, use of the SLM as an easily reconfigurable input or Fourier plane filter was prevented by the variable bias requirement of this interface.

6.4.2 Interface for devices used in an optical processor

The devices used in the optical processor were interfaced to the computer using a proprietary Hitachi LCD Controller/Driver LSI, the HD44780 [102]. This CMOS integrated circuit is mounted in an 80 pin package and is significantly more sophisticated than the previous interface.

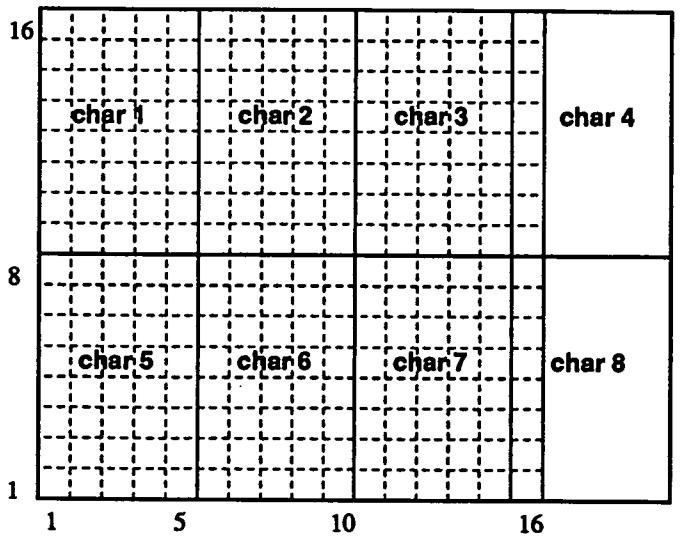

Figure 6.10: Mapping of pixels to user defined characters

The IC is designed as a single-chip driver/controller for an alphanumeric dot matrix display of 2 lines of 8 characters, each character occupying 5x7 pixels, with an optional 5x1 pixel cursor underneath. For the purpose of driving a 16x16 pixel SLM displaying arbitrary non-alphanumeric patterns it is necessary to consider the SLM as two lines of a display, each displaying 4 characters (see figure 6.10), and to use the chip's character generator RAM to build up eight 8x5 pixel user-defined characters which will yield the desired pattern. Software was written by Ian Redmond and Douglas McKnight [103] which allowed a 16x16 pattern to be input from the keyboard of a BBC microcomputer while controlling the operation of the HD44780 chip, including the definition of the user-defined characters. No further intervention from the computer is necessary once these characters have been defined and written to the chip due to the provision on the IC of an internal clock and the logic to continually refresh the frame being written to the SLM. Hence it is possible for a series of SLMs to be run from the same program by read-enabling a particular driver/controller IC individually whenever a new frame needs to be written to its SLM, leaving the others to continue operating by themselves.

The actual drive waveforms generated are shown in figure 6.11 along with the resultant superpositions seen by selected and non-selected pixels. Despite the apparent complexity of these waveforms, only one additional power supply con-

nection of -5V is necessary in addition to the ground and V_{CC} (+5V) already made to the driver/controller IC.

The origins of the various voltage levels are shown in figure 6.12, with the variable resistor allowing an adjustable negative bias. In this case the bias voltage allows contrast to be maximised by adjusting the overall amplitudes of the waveforms, rather than adding an offset for d.c. balancing, as in the prototype interface. The pattern dependent offset is unnecessary in this case as each waveform sums to 0 over one frame period, as required. The frequency of the chip's internal clock is set by an external resistance of 92KΩ between two of the pins, giving a frame frequency of 78.1Hz for the SLM. This interface proved convenient and reliable, and, when used in conjunction with the amplitude modulators, resulted in excellent contrast ratios irrespective of the pattern on the SLM.

6.5 Characterisation of the SLMs

The three SLMs fabricated for use in the correlation experiments of the next chapter were first characterised, so that the results could be sensibly related to the performance of the SLMs. All the SLMs were examined interferometrically to determine thickness variation across the pixel array, and to examine the phase modulating properties of the phase-only devices. In addition contrast ratio measurements were carried out for the amplitude modulator. Switching times were not examined, as the factor limiting the frame rate in the experimental set-up was the speed of the microcomputer running the interface.

6.5.1 Interferometric examination of the SLMs

As the SLMs are transmissive a Mach-Zehnder interferometer (as shown in figure 6.13) is appropriate for their examination [104]. With no sample in either arm of the interferometer, a small tilt is introduced between the outgoing beams to

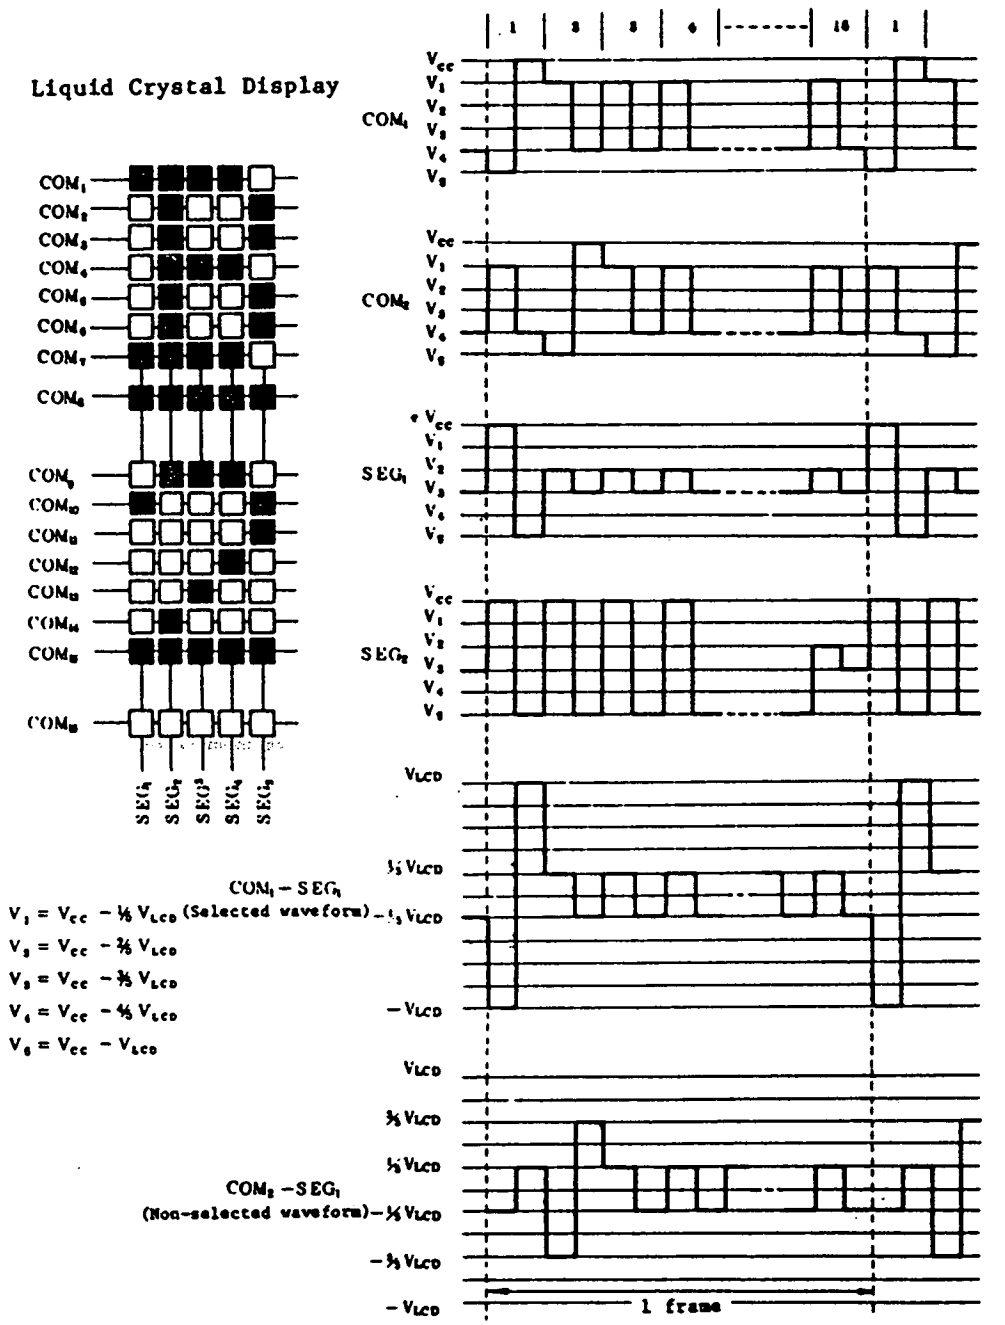

Figure 6.11: Drive waveforms from HD44780

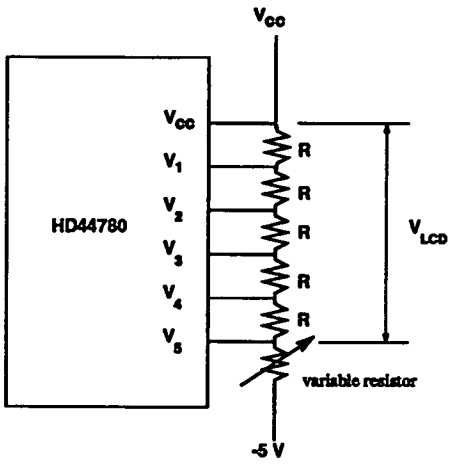

Figure 6.12: Voltage inputs to HD44780

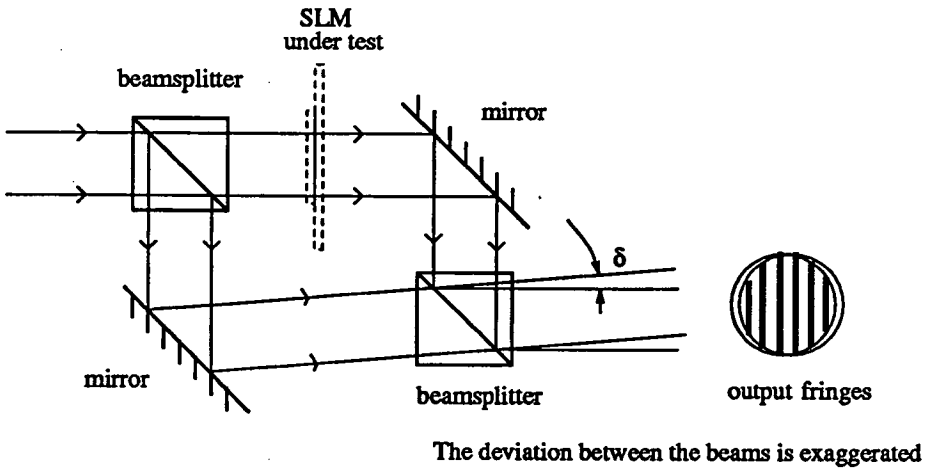

Figure 6.13: Mach-Zehnder interferometer

give a series of equispaced straight reference fringes, with which the interferograms of the SLMs can be compared.

In practice the modified implementation of the Mach-Zehnder interferometer shown in figure 6.14 was used, as its 'in-line' configuration meant that it could be easily mounted on the optical bench used in the processing experiments. There is a significant difference in path lengths between the two beams, but distinct fringes were obtained as the coherence length of the laser is much longer ($\sim$ 10m [105]). The plane of polarisation of the incident collimated light was aligned with the rubbing direction at the front surface of each SLM.

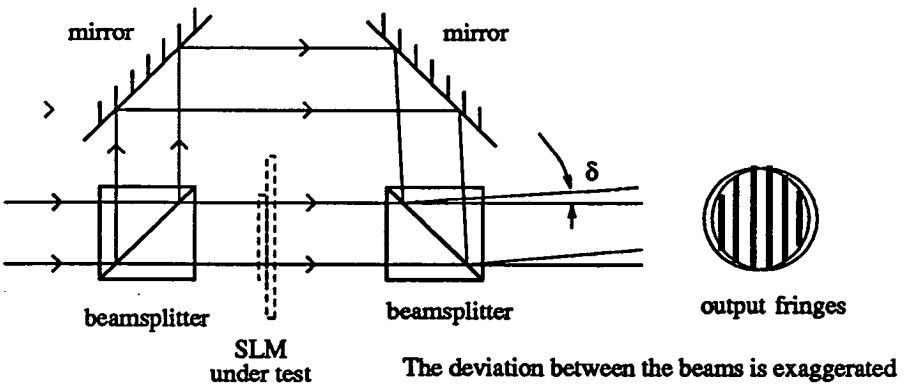

Figure 6.14: Modified experimental apparatus

Interferograms of the SLMs

The reference interference pattern and the interferograms from the three SLMs are shown in figure 6.15. All four pictures were captured with the adjustment of the interferometer and position of the array camera held constant. The pictures were analysed by photocopying the interferograms onto overhead transparencies at a magnification of 1, so that they could be compared directly with the reference when the borders of the two pictures were aligned.

It is not possible to make absolute measurements of the thicknesses of the SLMs, as the displacements of the reference fringes could not be followed. However, relative variations in thickness manifest themselves as differences in the fringe spacing, so this can be estimated qualitatively. With respect to the reference fringes shown, a variation in the fringe spacing implies thickness variation in the x-direction, and a change in the fringe angle implies thickness variation in the y-direction. Curvature of the fringes implies a combination of the two, and bowing of the liquid crystal cell. Estimates of the phase/thickness variations over the SLMs were made by counting fringes on the interferograms relative to a reference point, chosen here as the bottom right hand corner of the pixellated array, and comparing with the reference fringes. The size of the arrays was 8mm $\times$ 8mm.

The regular pixel phase modulator shows marked curvature of the fringes, and increasing spacing between fringes, indicating that the liquid crystal cell is bowed,

a) Reference fringes, no SLM in beam.

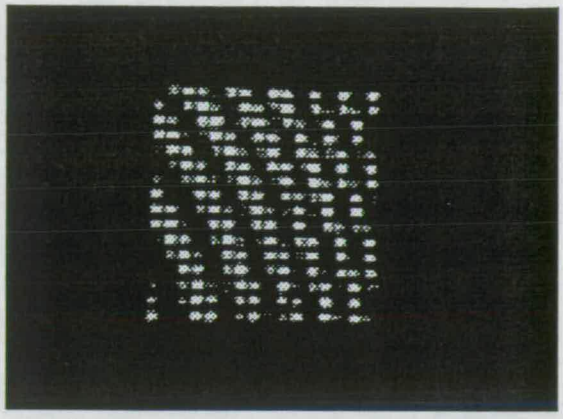

b) Regular pixel phase modulator in beam

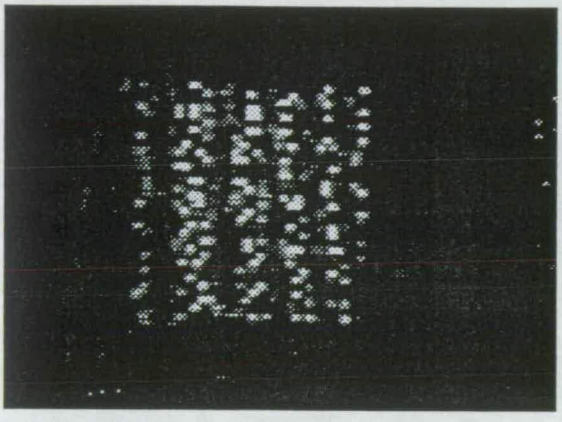

c) Random pixel phase modulator in beam

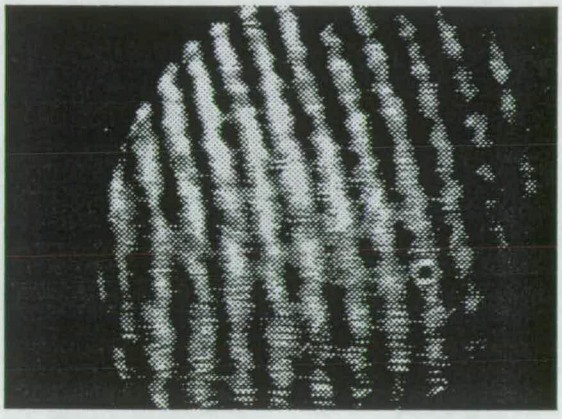

d) Random pixel amplitude modulator in beam

Figure 6.15: Interferograms from the modified Mach-Zehnder interferometer

and will thus behave like a low powered lens in addition to its specified task. It should in principle be possible to estimate the effective power of the lens effect, but this has not been done here. The main effect is wedging in the y-direction, as evidenced by a difference of 3 fringes between the top edge and the bottom edge, with respect to the references, with a wedge of half a fringe in the x-direction. This is obviously not ideal, but not necessarily disastrous– if the SLM can modulate uniformly over its area, the wedging will simply lead to a shift of the output in the output plane, with only the slight lensing degrading the form of the output.

The random pixel phase modulator was much more uniform, showing less than

half a fringe of wedging in the x-direction, and approximately half a fringe of wedging in the y-direction. Because of the random pattern of the transmissive area it is difficult to count fringes with an accuracy of better than $\pm\frac{1}{2}$ a fringe.

The random pixel amplitude modulator was also fairly uniform in thickness, with half a fringe of wedging in each of the $x-$ and $y-$directions. This amounts to a total difference in thickness between the bottom right corner and top left corners of 1 fringe. This is approximately 1.3μm, taking the refractive index of E300 as ~ 2.

Uniformity of phase modulation

The Mach-Zehnder interferometer was also used to estimate the uniformity of the phase modulation over the area of each of the phase modulators. Interferograms were obtained for each SLM in its 'off' state, and in a state with all pixels 'on' and the applied voltage tuned to maximise the number of pixels showing a reversal of contrast (i.e. a phase shift of π). The tilt between the beams of the interferometer was adjusted to give wider fringes, so that a change in contrast was easier to note.

The results are shown in figure 6.16. The random pixel phase modulator is quite uniform, although the figure shows a modulation depth slightly less than $\pi-$ the fringes have moved uniformly, but slightly less than the half fringe necessary to give contrast reversal everywhere. The regular pixel phase modulator is less successful, with no contrast reversal at the right hand side, or bottom left hand corner. This may be attributed to the extreme wedging and curvature of the cell, and to the third and fourth columns from the right having no applied voltage across them, due to open circuits in the electrodes.

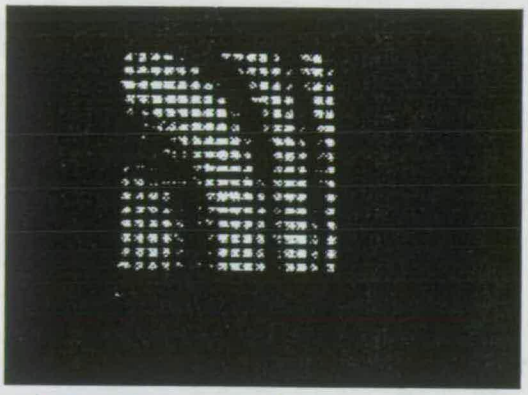

a) Regular pixel phase modulator, all pixels 'off'

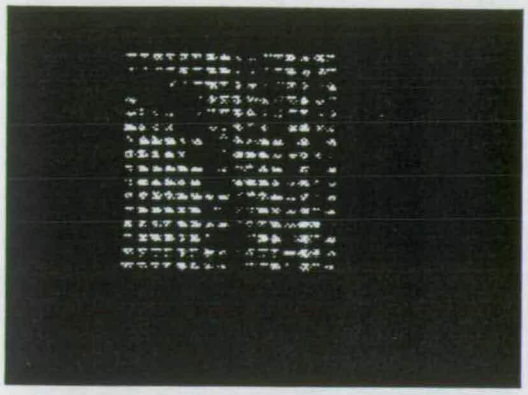

b) Regular pixel phase modulator, all pixels 'on'

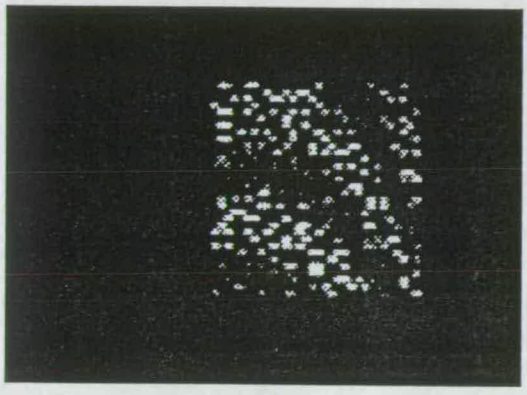

c) Random pixel phase modulator, all pixels 'off'

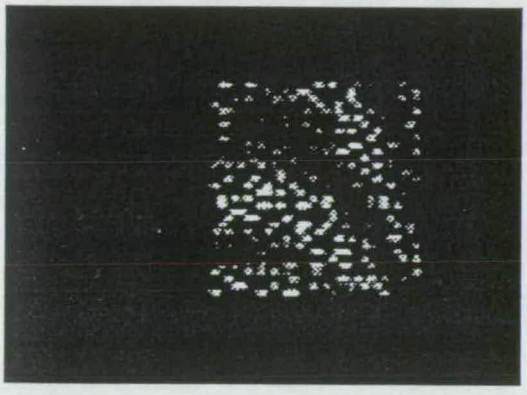

d) Random pixel phase modulator, all pixels 'on'

Figure 6.16: Interferograms from the modified Mach-Zehnder interferometer to determine the uniformity of modulation for the phase modulators

6.5.2 Contrast measurement for the amplitude modulator

The random pixel amplitude modulator was placed between parallel aligned polarisers and illuminated with collimated coherent light from a helium-neon laser. The SLM was oriented such that the first polariser was aligned with the PVA rubbing direction on the side of the liquid crystal cell nearest the polariser. The SLM was imaged onto the calibrated array camera, and frames captured of the two areas shown in figure 6.17a), with all pixels 'off' and all pixels 'on'. The frames for all pixels 'on' are shown in figure 6.17b) and c).

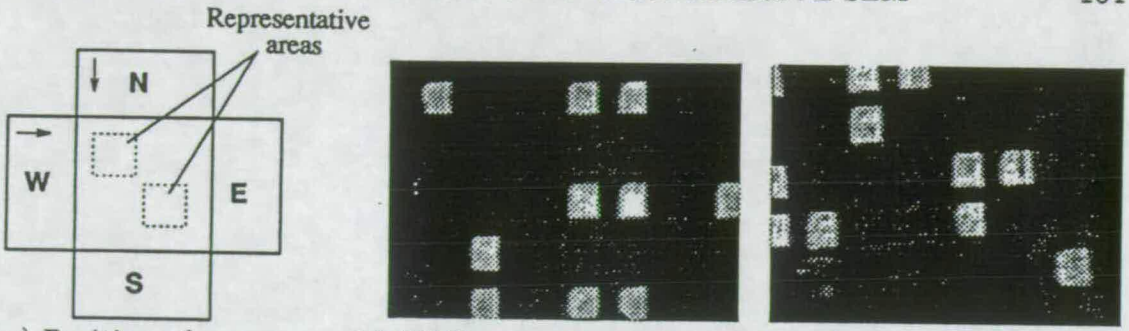

a) Position of areas on the SLM

b) 'NW' area with pixels 'on'

c) 'SE' area with pixels 'on'

Figure 6.17: Areas on the random pixel amplitude modulator selected as representative for contrast measurements

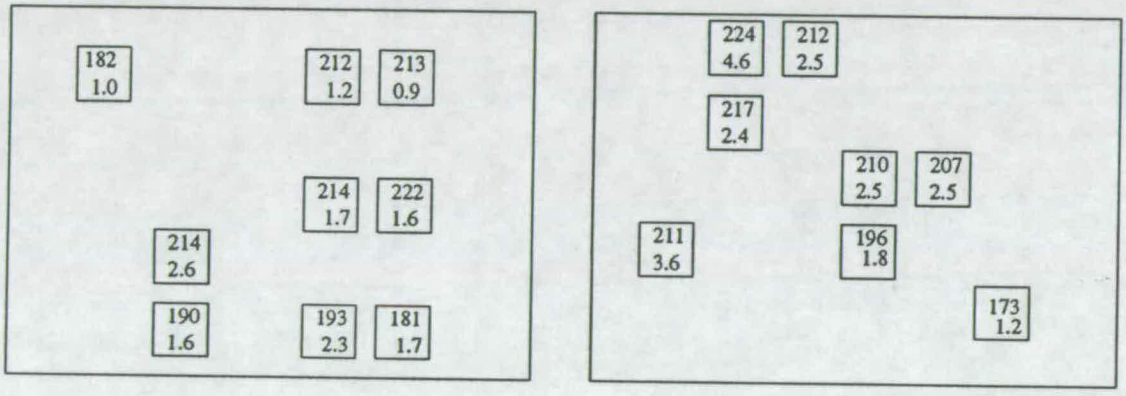

Figure 6.18: Results of contrast measurements

The average intensity over the area of each pixel was calculated for the 'off' and 'on' states, and the 'on' value for each pixel divided by its corresponding 'off' value to give a contrast ratio for each individual pixel. The mean of these contrast ratios was calculated to give an average figure for the entire SLM, with the uncertainty in this figure taken as the standard error on the mean for the contrast ratios for each pixel.

The experimental results are presented in figure 6.18 with each 'pixel' containing the 'off' and 'on' intensity values. The experimentally determined average value for the contrast ratio of the SLM is (115 ± 12):1. This high value is an excellent result and compares favourably with other available SLMs. The working SLM with all pixels 'on' is shown in figure 6.19.

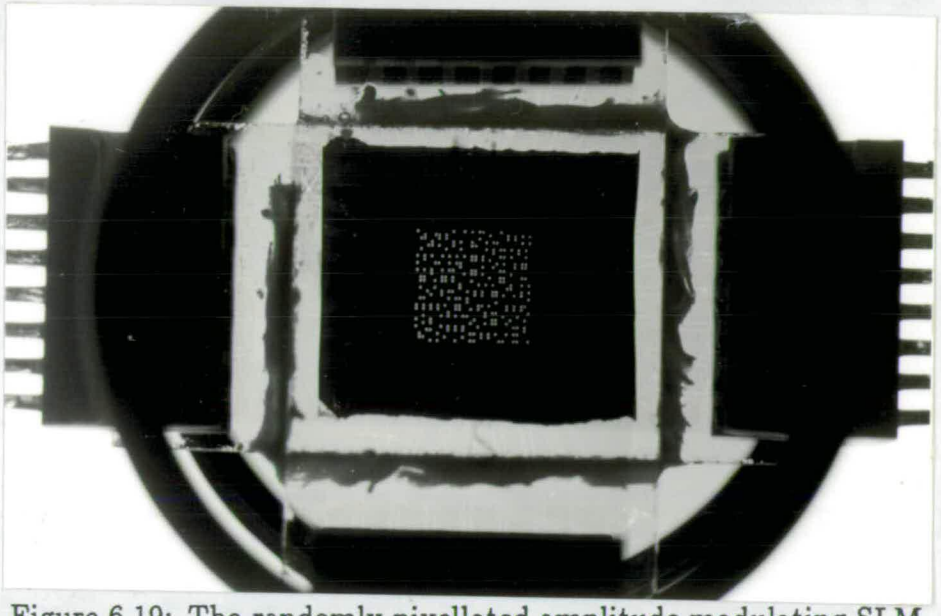

Figure 6.19: The randomly pixellated amplitude modulating SLM

6.5.3 Physical dimensions of the SLMs

The SLMs were 16×16 pixel devices with the pixels on a $(502 \pm 3)\mu$m pitch, giving an array size of 8mm×8mm. The pixels were $(182 \pm 3)\mu$m square, less than the ideal of 250μm because of the relatively crude contact printing used to define the electrodes. The width of the lines separating the pixels was $\sim 70\mu$m, which should be compared with line widths of $\sim 1\mu$m readily achievable in the lithographic processes of the semiconductor industry. The measurements were carried out using a travelling microscope.

6.6 Summary

A laboratory process to fabricate transmissive matrix addressed SLMs has been described. The technology used may be adapted to make a wide range of devices, using either nematic or ferroelectric liquid crystal to make phase or amplitude modulators. Electrical interfacing of a nematic device to a BBC microcomputer has been described. Transmissive matrix addressed nematic liquid crystal SLMs have been constructed and characterised for use in the following chapter.

Chapter 7

Correlation experiments

Chapter 5 has shown that by randomising the pixel positions of an amplitude filter with all pixels 'on', it is possible to remove all the well defined spectral orders in the filter's power spectrum other than the zero order, and consequently the replicas that would appear in the output of an optical processor. The motivation for this has been to allow the use of an optical correlator with a filter matched in space-bandwidth product to the target to be detected in the input scene, as opposed to that of the much larger overall input scene. This becomes possible when the distinct replicas which lead to aliasing are removed.

This chapter generalises the concept of randomising pixel positions to an actual filter, in this case the 16×16 SLMs of Chapter 6, to demonstrate its validity and to show how it now is possible to use a Fourier plane filter with a lower space-bandwidth product than the input plane. The potential benefits to be accrued from this include lower cost, higher speed and reduced complexity of addressing. Amplitude- and phase-only modulators are used, and the results compared with computational simulations.

The analysis so far has concerned itself with the case where all pixels are 'on', and the discrete position randomisation scheme has been arrived at through an analytical treatment of general randomised pixel schemes. That it should be possible to use the array as a filter by switching pixels 'off' (e.g. changing the

transmission state from 1 to 0) while still preventing aliasing occurring becomes apparent when we consider the following more physical interpretation of why the discrete position randomisation removes the odd spectral orders. With reference to figure 5.8, four positions for the pixel within the $\alpha \times \alpha$ cell are possible, labelled position 0, 1, 2, or 3. If all the pixels in the array are in the same position in each cell we have a regular array with well-defined spectral orders in its Fourier transform. Naming the regular array of pixels in position 0 as $r_0(\xi,\eta)$, the regular array of pixels in position 1 as $r_1(\xi,\eta)$ etc. it is apparent that

$$\begin{aligned} r_1(\xi,\eta) &= r_0(\xi - \frac{\alpha}{2},\eta) \\ r_2(\xi,\eta) &= r_0(\xi - \frac{\alpha}{2},\eta + \frac{\alpha}{2}) \\ r_3(\xi,\eta) &= r_0(\xi,\eta + \frac{\alpha}{2}) \end{aligned} \tag{7.74}$$

Fourier transforming, and applying the shift theorem gives

$$\begin{aligned} \mathcal{F}[r_0(\xi,\eta)] &= R_0(x,y) \\ \mathcal{F}[r_1(\xi,\eta)] &= R_0(x,y)e^{-i2\pi\frac{\alpha}{2}x} \\ \mathcal{F}[r_2(\xi,\eta)] &= R_0(x,y)e^{-i2\pi(\frac{\alpha}{2}x - \frac{\alpha}{2}y)} \\ \mathcal{F}[r_3(\xi,\eta)] &= R_0(x,y)e^{+i2\pi\frac{\alpha}{2}y}. \end{aligned} \tag{7.75}$$

At the positions of the odd spectral orders $x = x' = (2n+1)/\alpha$, $y = y' = (2m+1)/\alpha$, with n, m integers, giving

$$\begin{aligned} R_1(x',y') &= -R_0(x',y') \\ R_2(x',y') &= +R_0(x',y') \\ R_3(x',y') &= -R_0(x',y') \end{aligned} \tag{7.76}$$

i.e. positions 1 and 3 give contributions which are in antiphase to those of positions 0 and 4. Thus, so long as there are equal numbers of 'on' pixels in each of the four positions, the net effect should remain as cancellation of the odd spectral orders. Whether or not there are equal numbers will depend on various factors – the actual pixel positions chosen in the array, the algorithm used to decide the binarisation, and the reference object from which the filter is calculated. Even for simple objects the phase distribution in the Fourier plane is fairly complicated so it is reasonable to expect that once a filter has been computed, the 'on' pixels will be evenly distributed among the four positions. Any slight mismatch in the numbers of 'on' pixels for each position will lead to the presence of odd spectral orders, albeit very faint– though this may be significant if the zero order itself is of the same magnitude, due to the total number of transmitting 'on' pixels being small. In practice, each of the filters computed had approximately half of the total number of pixels 'on', with an equal distribution between the four positions allowed in the case of the discretely randomised filters. As a consequence they performed well with regard to preventing aliasing and replication as desired, as will be seen.

7.1 Simulations

Before the optical experiments took place, computational simulations were carried out to demonstrate the viability of using the 16x16 filters, and to gain benchmarks for the idealised situation with which the optical results could be compared. The simulations again were based on 256×256 arrays of numbers representing the object, filter, and image planes with the Fourier transforms between these planes carried out by an FFT routine incorporated in the FORTRAN program **corcon**. This program was also used to generate the filters by applying a binarisation algorithm to the complex amplitude in the Fourier plane. Once the filters were generated, the pattern information was written to a file as a 16×16 array of ones and zeroes, representing the binary state of each pixel. This file could be read back into **corcon** at a later stage and used to generate a filter which is multiplied by the Fourier transform of the input scene, the filtering operation being achieved

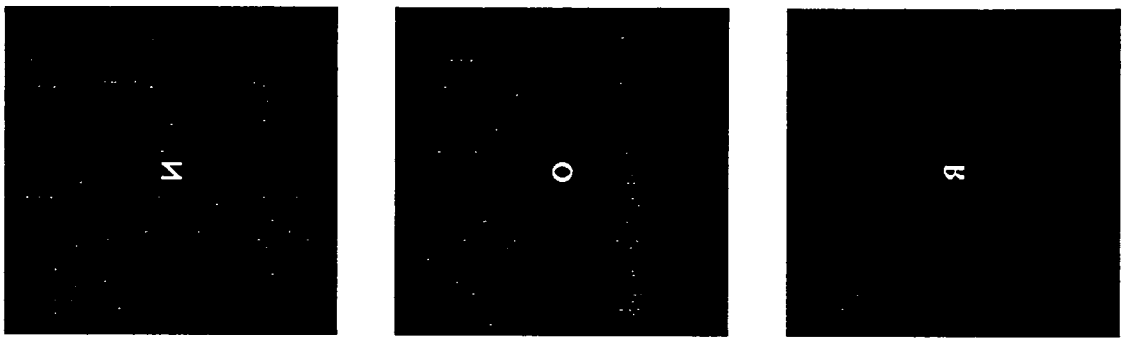

Figure 7.1: Letters used as inputs in the correlation experiments. Their orientation is not significant

by point to point multiplication of the two complex arrays[1]. The result is Fourier transformed again, to give the correlation output as a 256×256 array of numbers which is squared to give intensities before being written to a file.

The input

The targets to be detected by the correlator are necessarily simple, because of the low SBP of the filters. As a 15x15 pixel area of each SLM was used to perform the filtering (an outside row and column were left switched off so that the high frequency cutoffs to either side of the pixel centered on the origin were the same), the targets chosen were defined on a 15x15 array of points, corresponding to square pixels with a 100% fill factor on the actual input transparencies used in the optical experiments. Thus the individual targets are matched in SBP with the filter. Simple letters were chosen, as shown in figure 7.1, with the N and O exhibiting symmetry properties, and the R displaying no symmetry but sharing a resemblance to the O through its curved part.

The targets were chosen to be binary for ease of comparison with the optical experiments, where the input to the correlator was achieved by mounting a uniformly illuminated transparency in plane P1. Variations in a transparency's amplitude

[1]In practice the complex numbers necessary to represent the amplitude were coded as 2 arrays representing the real and imaginary parts respectively

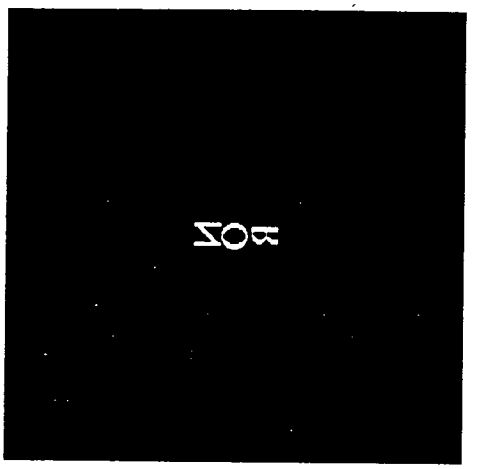

Figure 7.2: Multiple letter input used in the correlation experiments

transmission may lead to corresponding variations in phase due to the differing densities of the emulsion necessary to achieve shades of grey. These phase variations must be characterised unless the transparency has only two states, one opaque, the other uniformly transmitting, as is the case here. The letters all had the same transmissive area (88 pixels out of the 15x15) so that cross-correlations of any letter with another should be less than their autocorrelations.

A range of input scenes was created from various combinations of the target letters, with attention focussed on the particular combination shown in figure 7.2. Here the N, the O, and the R are placed next to each other at a pitch of 16 points, corresponding to the separation of the spectral orders in the transform of the appropriate filter based on a regular pixel array. Thus the orders present in the spectrum of a regular array should lead to the first order replicas of the correlation signals from one letter coinciding with the zero order correlation signals from its neighbours, and vice versa: the coherent sum of these contributions may give rise to misleading results, which we hope will <u>not</u> be the case when the corresponding filters based on the discretely randomised arrays are used.

Binarisation of the filters

As will be described later, simulations and optical results were obtained for amplitude encoded binary phase-only filters (AEBPOFs) and for binary phase-only filters (BPOFs). In both cases, binarisation of the phase to be ascribed to each pixel of the filter was necessary. This is obviously critically dependent on the target. The first stage of the filter binarisation is to place the uncluttered reference target at the input of the simulated optical processor, equivalent to placing the 15×15 array of 1's and 0's representing the target letter at the centre of a 256×256 array of zeroes representing corcon's input plane. Centering the target is important, especially with regard to the symmetrical letters N and O, given that the filter generated will also detect for a 'false image' as described in Chapter 2. Later results will demonstrate this for the letter O as a reference object, centered correctly, and with its centre displaced one point away from the centre of the input plane.

With the reference placed at the centre of the input plane, its Fourier transform is calculated, and from this a complex amplitude is ascribed to each point in the 256×256 point Fourier plane. Each pixel in the Fourier plane filter covers a number of points so 'averaging' of the phase must take place over the pixel to obtain a phase-only filter. This was achieved here by a coherent sum of the complex amplitudes over each pixel, giving a single complex number $z_{n,m}$ for the $(n,m)^{th}$ pixel. It should be noted that an 'average' phase per pixel can be calculated by other methods which may give different results. Various methods have been tried [106], including calculation of the phase at each point on the pixel and averaging; calculation of the phase at each point, thresholding then averaging; averaging of the phases at each point, including a weighting factor at each point related to its intensity– but the results are not conclusive as to whether any of them are better than the coherent sum of complex amplitudes.

The quantisation of the 'average' phase for each pixel to one of two values was carried out by application of a threshold condition, which was the same for both AEBPOF and BPOF. This is represented in figure 7.3, which shows the complex plane on which the coherent sum for each pixel, $z_{n,m}$, could be plotted: the two

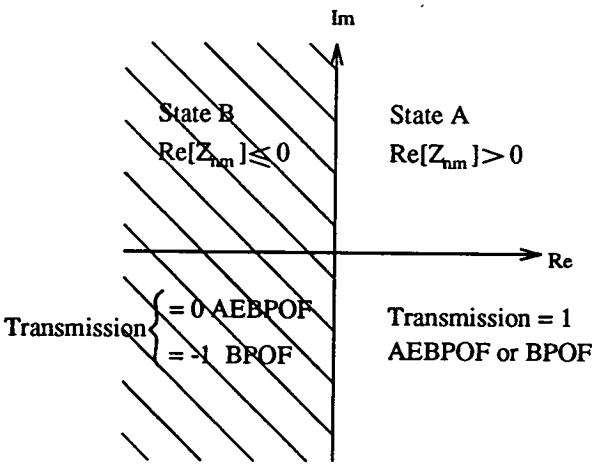

Figure 7.3: Complex plane representation ofthe filter binarisation.

states are chosen by dividing the complex plane along the imaginary axis.

This is the same binarisation algorithm as that described by Psaltis et al. [26] for the BPOF, and by Flavin and Horner [29] for the AEBPOF.

7.1.1 Simulation of AEBPOFs

Filters were generated to detect the N, the O, and the R, using a filter based on the special case of the discretely randomised array shown in figure 5.15 and the corresponding regular array with the same fill factor (25%). To critically sample the transform of each input letter the filters must have their pixels on an underlying pitch of 16 points (relative to the 256×256 arrays used in the computational simulations), with the pixels occupying 8×8 points. This gives a large number of points per pixel over which the coherent sum can be computed, but as the pixel occupies an even number of points it is not possible to centre the middle pixel of the array exactly symmetrically at the centre of the Fourier plane as desired. In practice the slight mismatch seemed to have little effect: ideally the the regular filters generated should have twofold rotation symmetry due to the binarisation algorithm's relation to the cosine transform[2]– as can be seen from figure 7.4 the

[2]For a function $g(x, y)$, if its cosine transform is $C(\xi, \eta)$ then $C(\xi, \eta) = C(-\xi, -\eta)$.

regular filters are very close to this symmetry. The filters generated based on the 'special case' array are not obviously symmetrical due to the randomisation of the pixel positions. These filters are shown in figure 7.5. All the filters were calculated from the Fourier transforms of the individual letters correctly centered at the input, with the white areas in the figures representing 'on' pixels with uniform transmission, and the black areas representing 'off' pixels or the interpixel areas with zero transmission.

7.1.2 Results of AEBPOF simulations

The results of the simulations are presented below, for single and multiple letter inputs to the correlator. As the computational simulations use an idealised model no estimates of uncertainty in the calculated values have been included. Also note that as the detection of correlation peaks is generally effected by thresholding at some intensity value, it is the maximum peak height that is important rather than the area under the peak.

Single letter as input

With a single letter as input the SBPs of the input and filter are matched, and we can note if the randomisation degrades the fundamental performance of the correlator. With the letter centered in the input plane this corresponds to a case where the region of interest in the output plane can be restricted to a 15×15 point region which excludes the replicas.

Figure 7.4 shows the correlation results for the regular case in columns headed by the filter used, with the second row the results for N as input, the third for O as input and the fourth for R as the input letter.

Figure 7.5 is the equivalent figure for the randomised array, with the maximum correlation value for each simulation in the two figures tabulated in Table 7.1. The

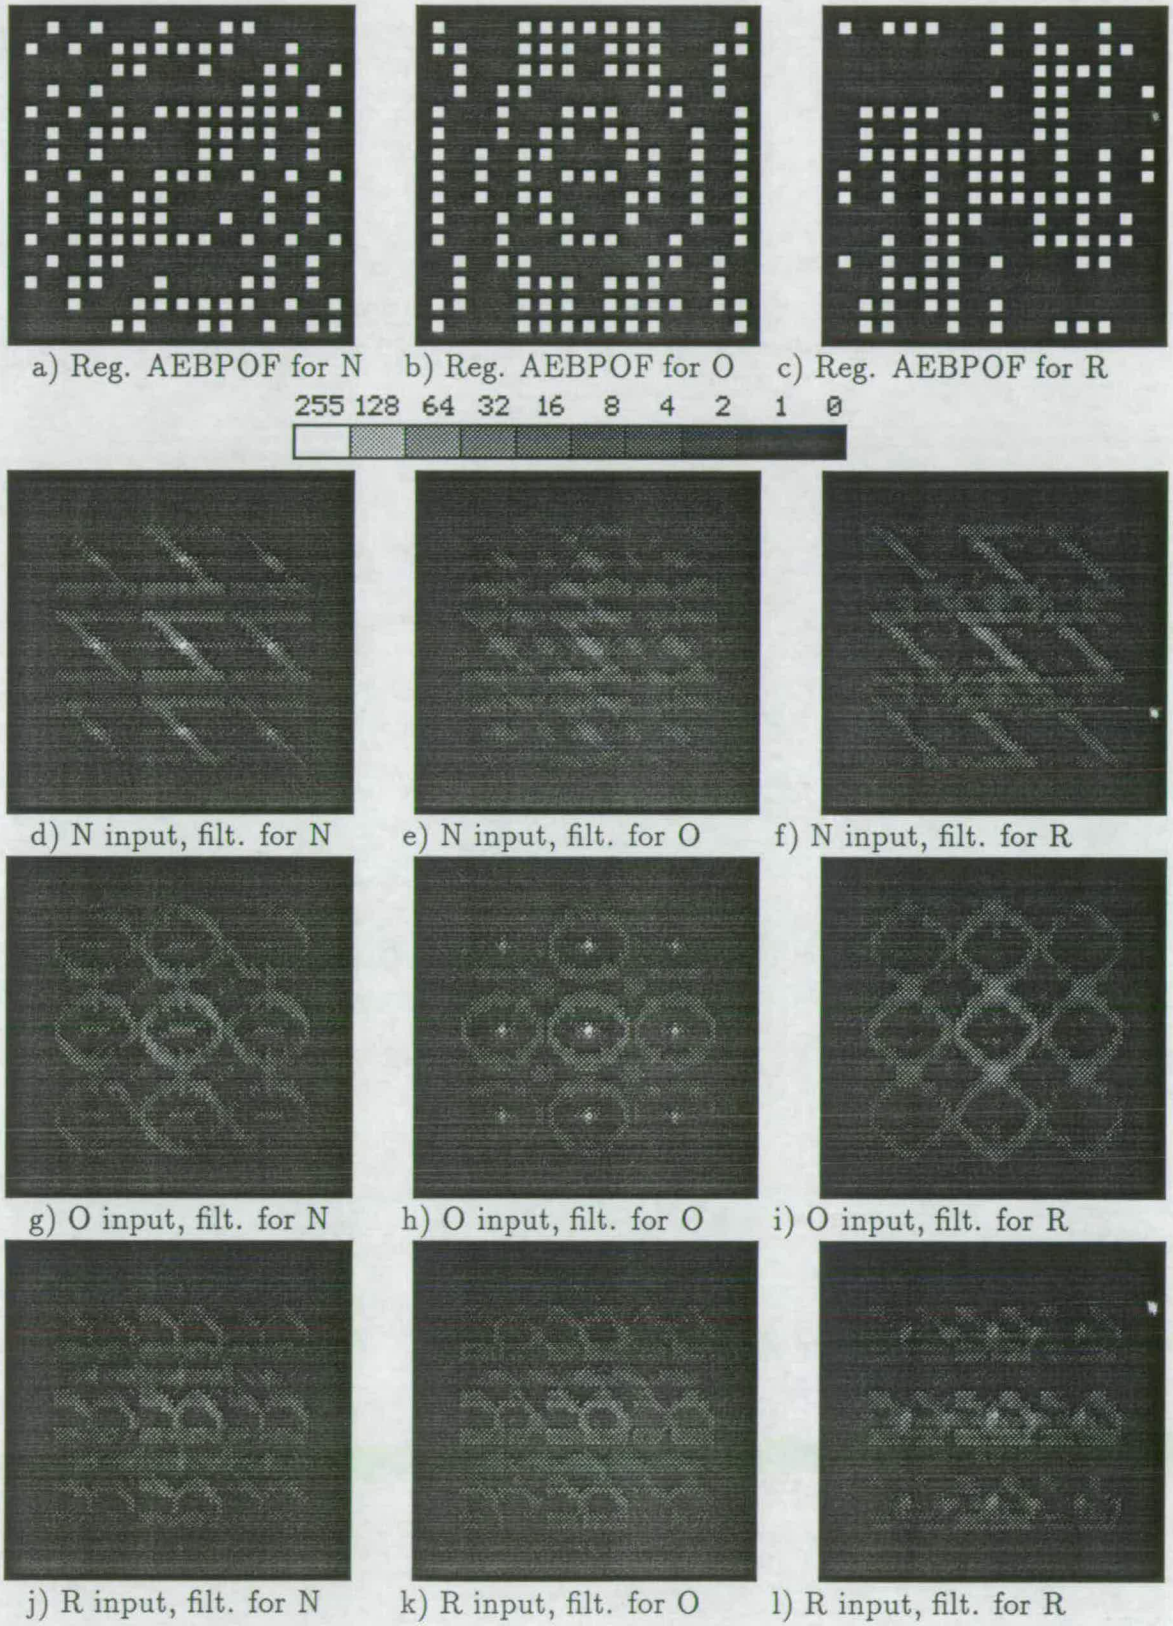

Figure 7.4: Filters for regular case AEBPOF, and correlation outputs using single letters N, O and R as input.

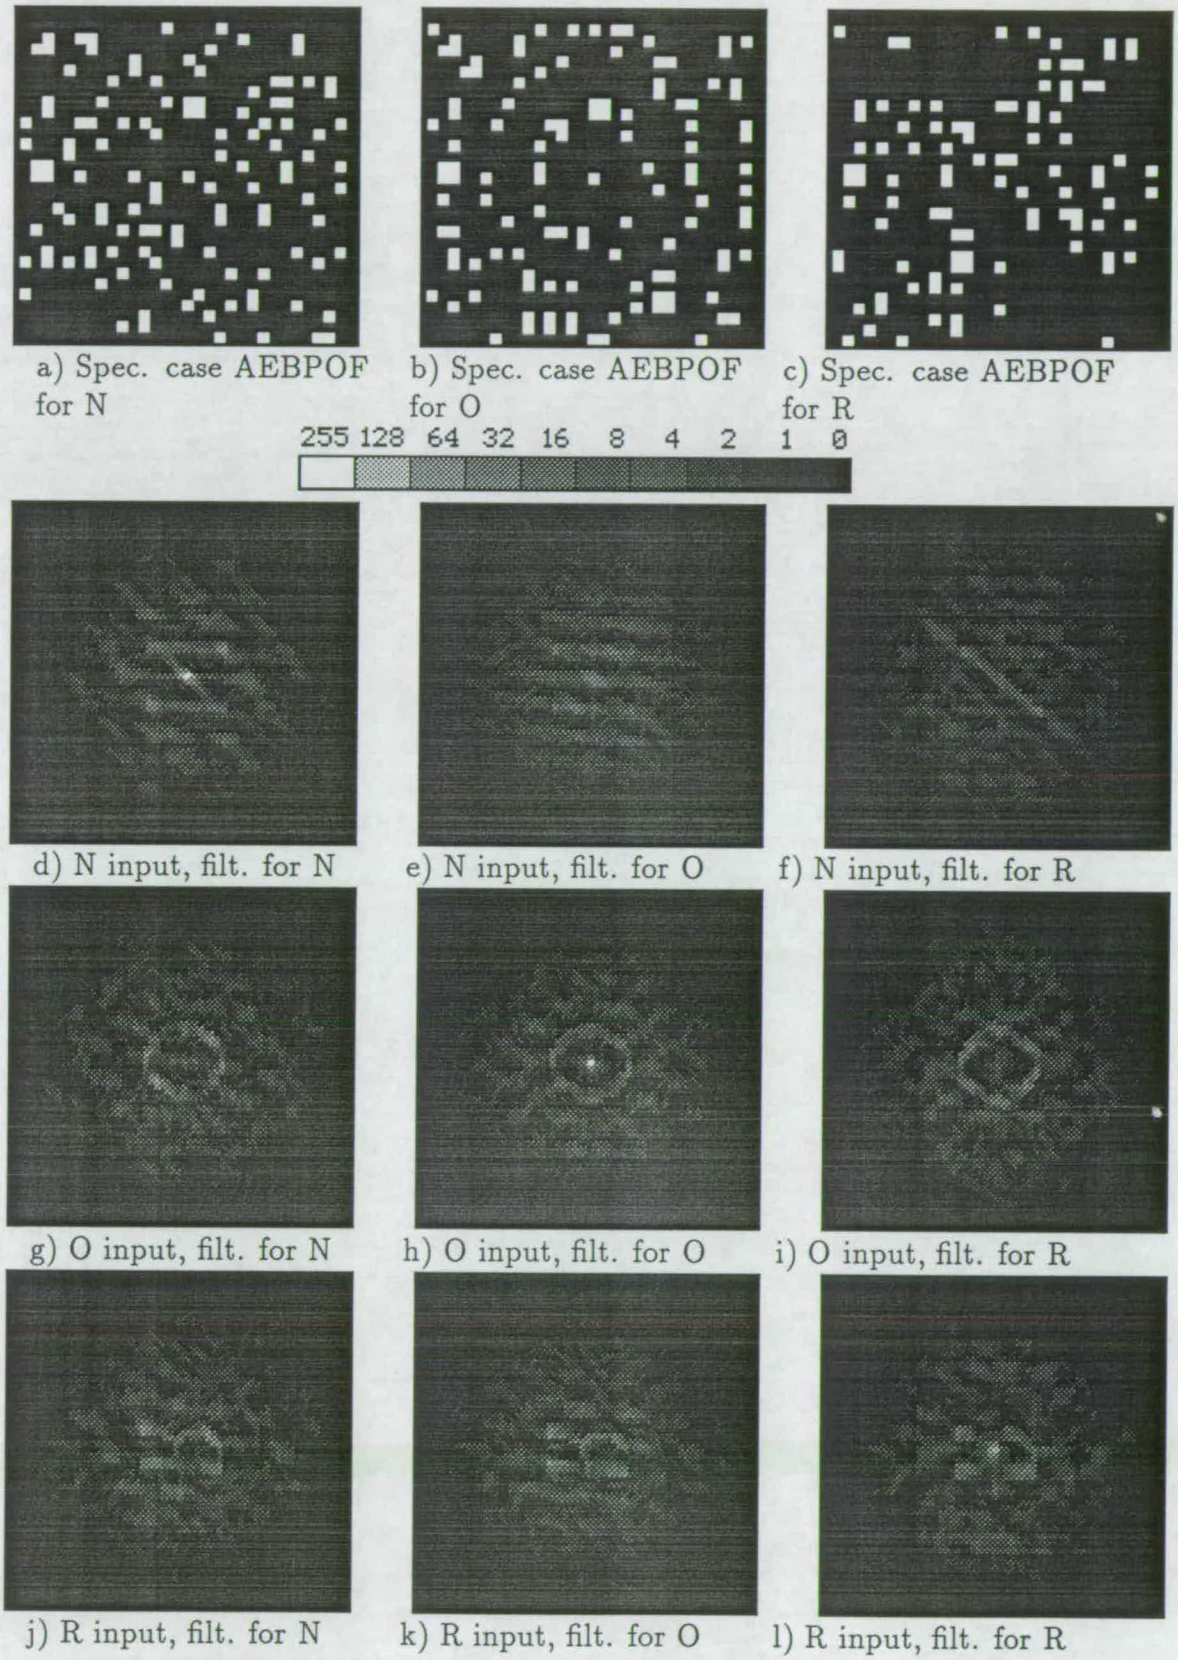

a) Spec. case AEBPOF for N
b) Spec. case AEBPOF for O
c) Spec. case AEBPOF for R
d) N input, filt. for N
e) N input, filt. for O
f) N input, filt. for R
g) O input, filt. for N
h) O input, filt. for O
i) O input, filt. for R
j) R input, filt. for N
k) R input, filt. for O
l) R input, filt. for R

Figure 7.5: Filters for 'special case' AEBPOF, and correlation outputs using single letters N, O and R as input.

Input	Regular filter for...			Randomised filter for...		
	N	O	R	N	O	R
N	0.402	0.088	0.088	0.379	0.068	0.045
O	0.073	0.360	0.083	0.107	0.407	0.070
R	0.071	0.058	0.160	0.101	0.125	0.183

Table 7.1: Results of correlation simulations for single letter input, using regular and 'special case' AEBPOFs

correlation outputs in the figures show intensities, and have been displayed using the image display package xv with the gamma value of the each image set at 1.5. The effect of this is to emphasise the lower intensities and prevent loss of detail when the figures are reproduced. The figures do allow comparisons to be made however, as the mapping of intensity to grey scale is the same for both figures, with white corresponding to the maximum value in table 7.1 of 0.407. The figures show the central 64×64 points of the correlation plane.

The table of results helps show that the correlation performance of the regular and randomised pixel filters is not significantly different in the case of a single centered target constituting the entire input scene– for each letter the maximum correlation value came from its corresponding filter, and for each filter the maximum correlation came from the corresponding letter, for both the regular and randomised cases. In addition, the maximum intensities for the non-target cross correlations are also similar for both regular and randomised cases.

The replicated outputs are clearly visible in the regular case, and show additional correlation peaks when the filter is matched to the target. These additional peaks are largest for the for the first order replicas and are 42% of the zero order for the N, 42% for the O and 45% for the R. The coordinate reversed input scene is also evident in the output, as expected with an AEBPOF [29], and this too is replicated when the filter pattern is on a regular array of pixels.

In the random case replication is not noticeable, with only one sharp, unambiguous, correlation peak for each letter when it is matched to its filter. The clearly defined replicas are replaced by a diffuse background which does not hinder the recognition task.

Note that the letters with twofold rotational symmetry, the O and N, have the largest correlation signals. This is as expected for filters generated by this binarisation algorithm, as the correlation signals resulting from the primary and false images in the filter's spectrum add when the filter is generated from a centered object with twofold rotation symmetry, as is the case (see chapter 2).

Multiple letters as input

When multiple letters as shown in figure 7.2 are used as the input, we have the case where the SBP of the input is greater than that of the filter, which is matched to each letter but not to the total input scene containing the three letters. Thus replication will occur within the region of the correlation plane where we wish to detect the correlation signals, when a regular array is used. Further, as the letters are adjacent, their first order replicas will coincide, potentially leading to large false correlation signals. Using the filters with the randomised pixel positions we would expect that the absence of distinct replicas would allow the position of each letter to be still be distinguishable.

The correlation results for the regular and randomised cases are shown in figures 7.6 and 7.7 respectively. Again the two figures are comparable, with the maximimum grey scale value of white corresponding to the highest intensity present. Ideally the correlation peaks for each letter should be obvious, with the correlation peak for the N to the left, for the O in the centre, and for the R to the right. The figures show the central 64x64 point region of the correlation plane.

The correlation outputs of figure 7.6 illustrate the replication problem well, with particularly large, sharp, false peaks occuring using the filters for N and O, where the filters are as in figure 7.4. In the case of the output when the filter for the R is used, the correlation peak, though sharp, is lower in intensity than the noise in the vicinity of the O.

The results of the simulations using the randomised filter show a distinct improvement, with unambiguous sharp correlation peaks occuring for the letters that each

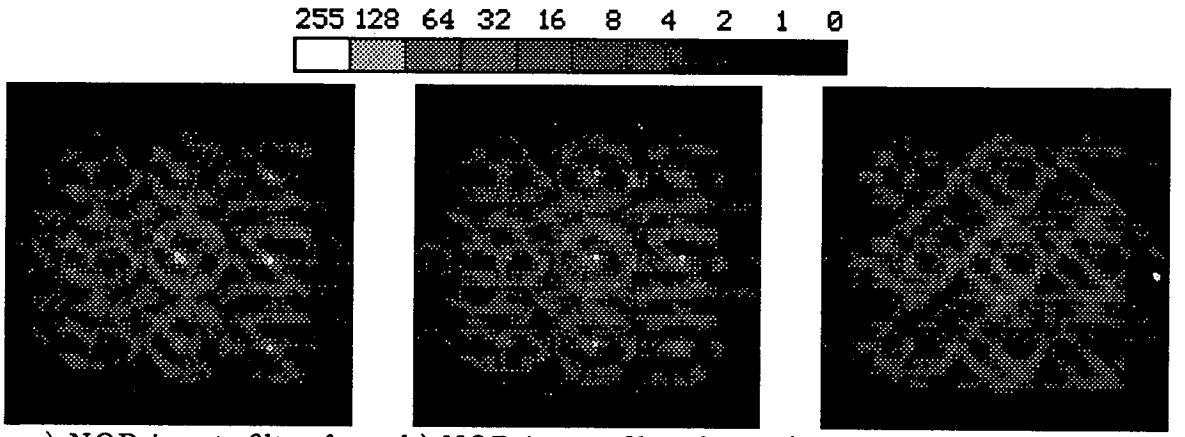

a) NOR input, filter for N

b) NOR input, filter for O

c) NOR input, filter for R

Figure 7.6: Output for correlation simulations using regular AEBPOF and multiple letter input. The correlation peak for the N is expected at the the right, for the O, in the centre, and for the R, at the left.

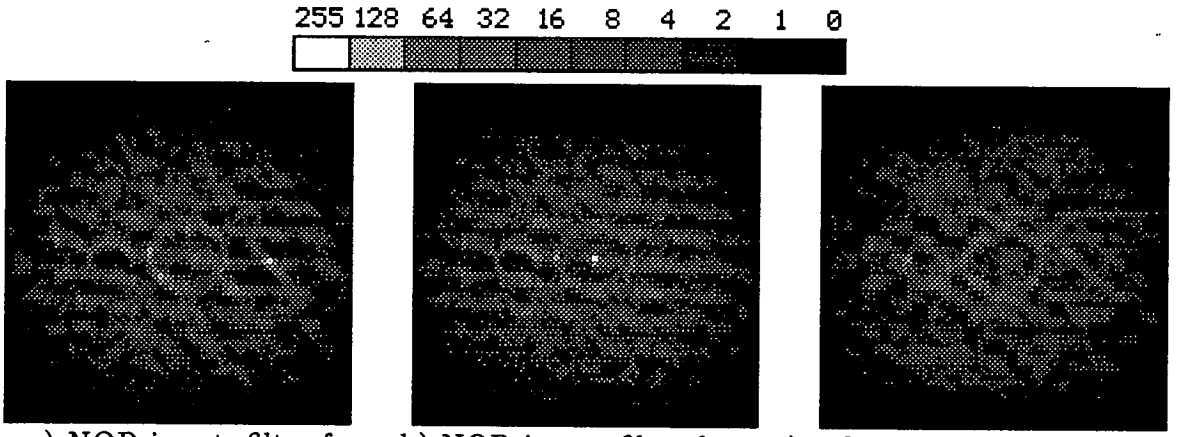

a) NOR input, filter for N

b) NOR input, filter for O

c) NOR input, filter for R

Figure 7.7: Output for correlation simulations using 'special case' AEBPOF and multiple letter input

Input	Regular filter for...			Randomised filter for...		
	N	O	R	N	O	R
N	0.502	0.278	0.159	0.396	0.138	0.101
O	0.294	0.499	0.201	0.228	0.586	0.147
R	0.124	0.179	0.167	0.140	0.141	0.174

Table 7.2: Results of correlation simulations for multiple letter input, using regular and 'special case' AEBPOFs

filter sets out to detect, with the correlation peak being the highest intensity point in the region of interest. Table7.2 shows the maximum correlation values in the vicinity of each of the three letters, for each filter used. The correlation values in the table are simple measures of maximum intensity in the region of each letter in the output and need not be related to a sharp peak, though for the true correlations they are. This should be noted when comparing the results for the filters to detect the N– the figures in the table suggest that the regular filter outperforms the random filter, while taking no account of the fact that the maximum false correlation value is much more sharply peaked for the regular filter and more likely to be mistaken for a true correlation. For the filters to detect the O and the R, the filters with the randomly positioned pixels clearly perform best, with higher actual correlation values and lower, broader, 'false' correlation values, which are easier to discriminate against by thresholding.

7.1.3 Effect of incorrectly centering the reference target

The simulations above show a distinct sharp peak for the O, which has been attributed to the reference letter being centered so that it exhibits twofold rotation symmetry around the optic axis, in order that the primary and false images in the filter's spectrum will be coincident. This can be demonstrated to be plausible by correlating with an input letter O when the filter has been generated from an O displaced from its symmetrical position by 1 point. The filter and the correlation output are shown in figure 7.8. The randomised filter is used so that we can verify these results by an optical experiment later.

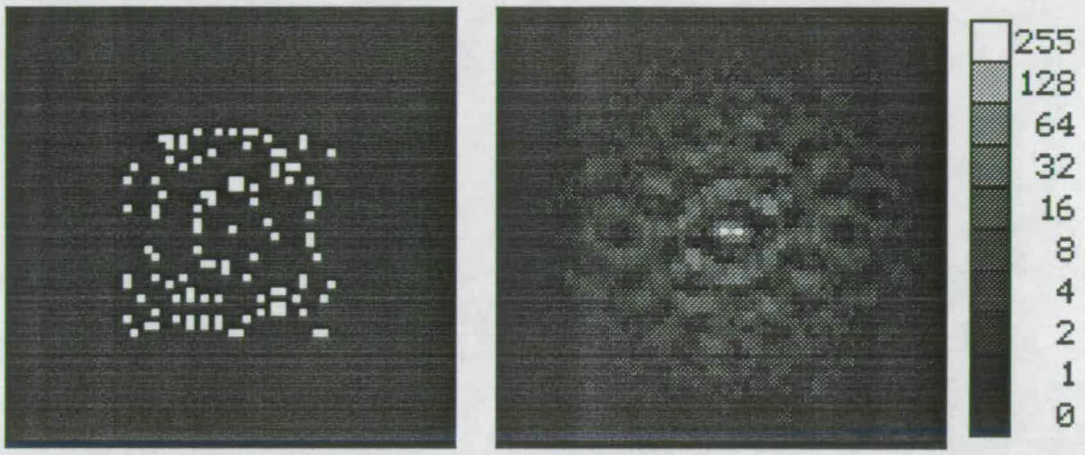

Figure 7.8: Filter used in, and output of, simulation of correlation with misaligned reference

Two correlation peaks are visible, each with an intensity of 0.172 units. This is still higher than the cross-correlations of the same filter with the N (0.075) or the R (0.070), but substantially lower than the value of 0.407 obtained earlier. Further, the correlation peaks are separated by 2 points, corresponding to the separation of the centres of the primary and false images.

7.1.4 Simulation of BPOFs

Simulations were carried out in a similar manner to those described above, except this time to demonstrate binary phase-only filtering. The filters simulated had a 25% fill factor as before, with the filters based on a regular array, and the discrete position randomisation described in Chapter 5. Note that the special case of the discrete position randomisation scheme was not used here because it was not necessary in order to achieve the optical implementation.

The binarisation of the filters was carried out by the same process and algorithm as described in the previous section, except that the binary states were a transmission of +1 if the real part of the coherent sum of the Fourier transform over the pixel was >0, and a transmission of -1 otherwise. The filters corresponding to the letters N, O and R for the regular case are shown in figure 7.9 and for the

Input	Regular filter for...			Randomised filter for...		
	N	O	R	N	O	R
N	0.989	0.116	0.103	0.947	0.116	0.114
O	0.168	1.296	0.210	0.123	1.304	0.169
R	0.202	0.134	0.587	0.090	0.118	0.694

Table 7.3: Results of correlation simulations for single letter input using regular and discretely randomised BPOFs

discrete randomisation case in figure 7.10. The transmission states of the pixels are indicated by a white pixel if the transmission is +1, and a grey pixel if the transmission is -1.

7.1.5 Results of BPOF simulations

The results for the simulations of the BPOFs are presented below, and compared with the results of the simulations for the AEBPOFs.

Single letter as input

The results of using a single letter as the input are shown in figures 7.9 and 7.10, corresponding to the regular and discretely randomised cases respectively,and in Table 7.3. As before the figures are displayed using xv with its gamma value set to 1.5, and showing the central 64 × 64 point region of the correlation plane. The figures are comparable with respect to intensities, with white corresponding to the maximum intensity obtained of 1.304 units.

With reference to the figures, the first thing that is apparent is that a version of the the input does not appear in the ouput as before. However the replicated correlation peaks and cross correlations are still clearly visible for the regular case filters as before, with the replicated first order correlation peaks 41.7% of the zero order for input of the N matched with its filter, 42.1% for the O and 44.9% for the R matched with its filter. For the discretely randomised filters this is not so

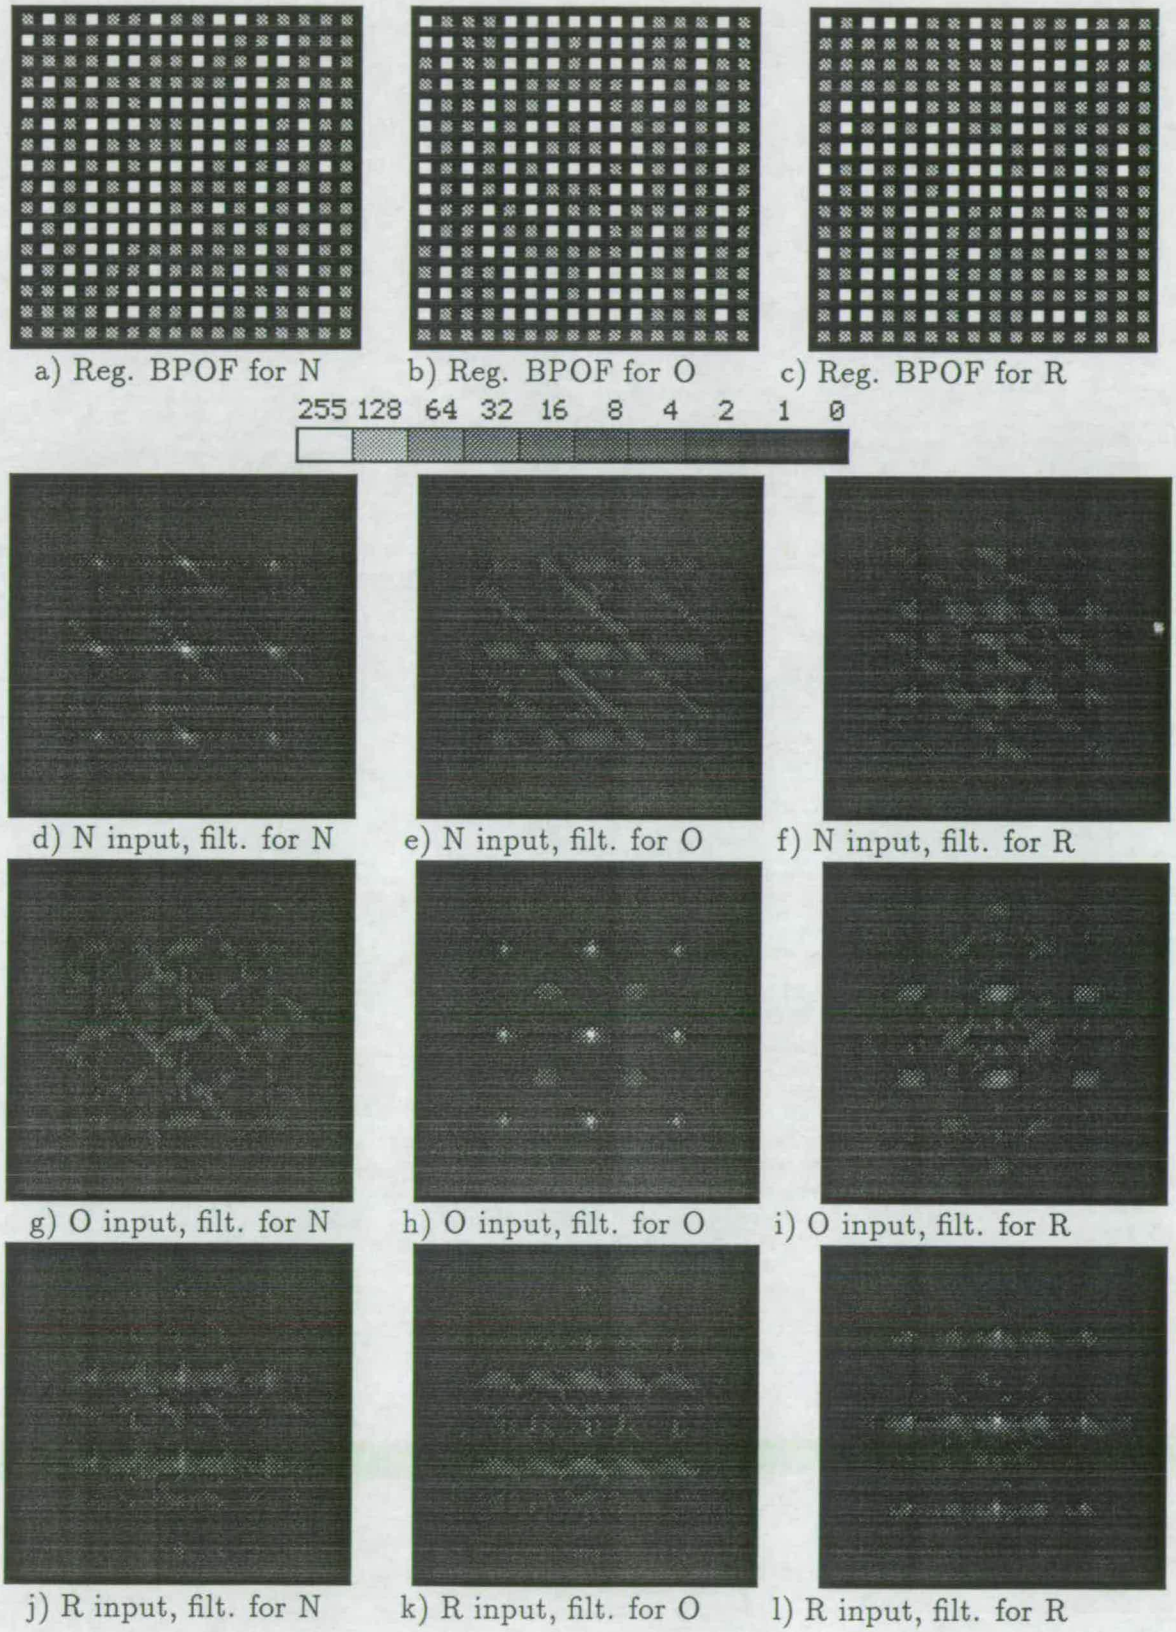

a) Reg. BPOF for N b) Reg. BPOF for O c) Reg. BPOF for R

d) N input, filt. for N e) N input, filt. for O f) N input, filt. for R

g) O input, filt. for N h) O input, filt. for O i) O input, filt. for R

j) R input, filt. for N k) R input, filt. for O l) R input, filt. for R

Figure 7.9: Correlation outputs for regular case BPOF with single letter input.

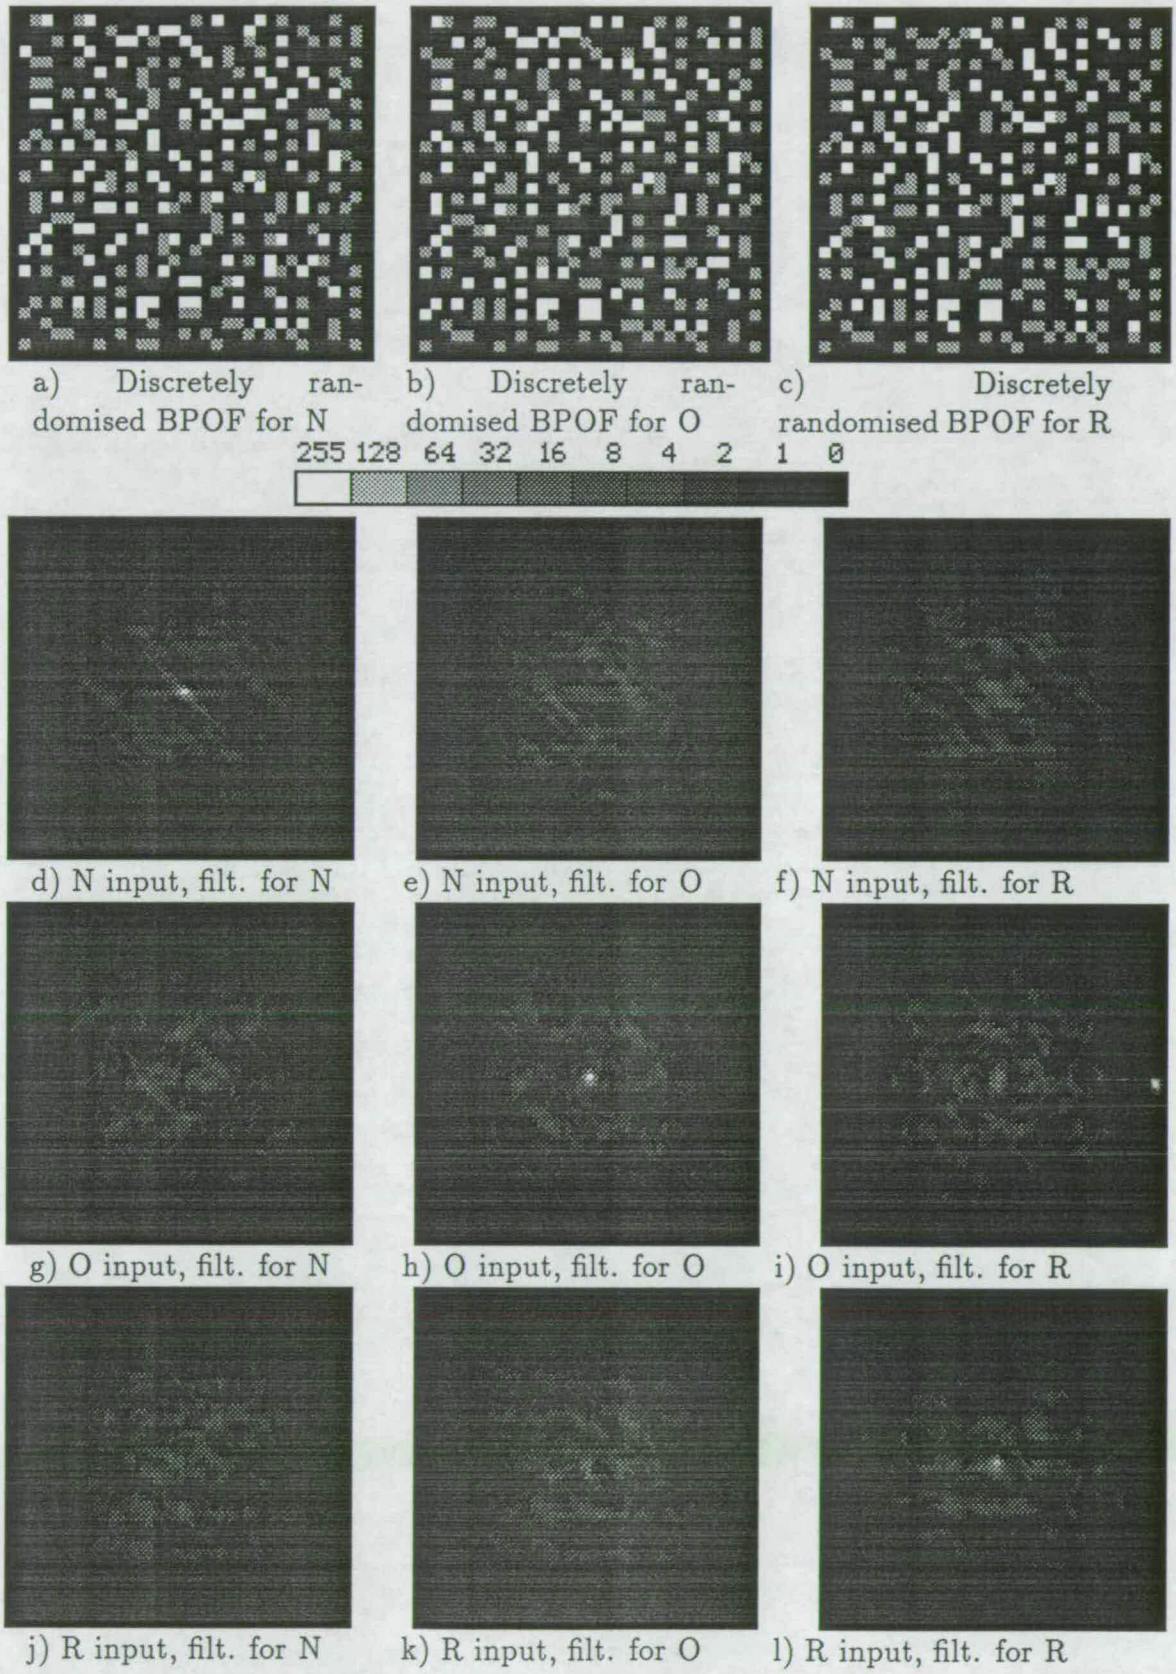

a) Discretely randomised BPOF for N
b) Discretely randomised BPOF for O
c) Discretely randomised BPOF for R
d) N input, filt. for N
e) N input, filt. for O
f) N input, filt. for R
g) O input, filt. for N
h) O input, filt. for O
i) O input, filt. for R
j) R input, filt. for N
k) R input, filt. for O
l) R input, filt. for R

Figure 7.10: Correlation outputs for discretely randomised BPOF with single letter input. Contrast with figure 7.9.

Input	Regular filter for...			Randomised filter for...		
	N	O	R	N	O	R
N	1.385	0.983	0.322	1.219	0.193	0.209
O	0.912	1.719	0.446	0.201	1.587	0.319
R	0.210	0.653	0.642	0.158	0.174	0.666

Table 7.4: Results of correlation simulations for multiple letter input using regular and discretely randomised BPOFs

with the desired correlation peaks showing up clearly and unambiguously. For each letter the largest correlation peak is obtained when its corresponding filter is used, and for each filter the largest correlation peak is obtained from its corresponding letter, for both regular and randomised cases. The magnitudes of the correlation peaks are very similar for both the regular and randomised cases as are the cross correlations with non-target inputs, indicating as before that randomisation of the pixel positions does not degrade the fundamental performance of the correlator.

Comparing the results of the simulations for the AEBPOFs and the BPOFs, it is apparent that the correlation peaks for the BPOFs are larger– this is as expected given that all pixels in the BPOF are transmissive. Again, the O and the N have significantly larger correlation peaks than the R, due to their rotational symmetry.

Multiple letters as input

The experiment with the multiple letter of the N, O, and R next to each other was repeated for the BPOFs, and the results shown in figures 7.11 and 7.12 corresponding to the regular and discretely randomised cases of the filters respectively, and in Table 7.4.

Again the mapping of grey scale to intensity is the same for each figure so that they are comparable. The central 96×64 point region of the correlation plane is shown. The entries in the table are simple measures of maximum intensity in the region of each letter in the output and need not be related to a sharp peak, though for the true correlations and results relating to the regular BPOF, they are.

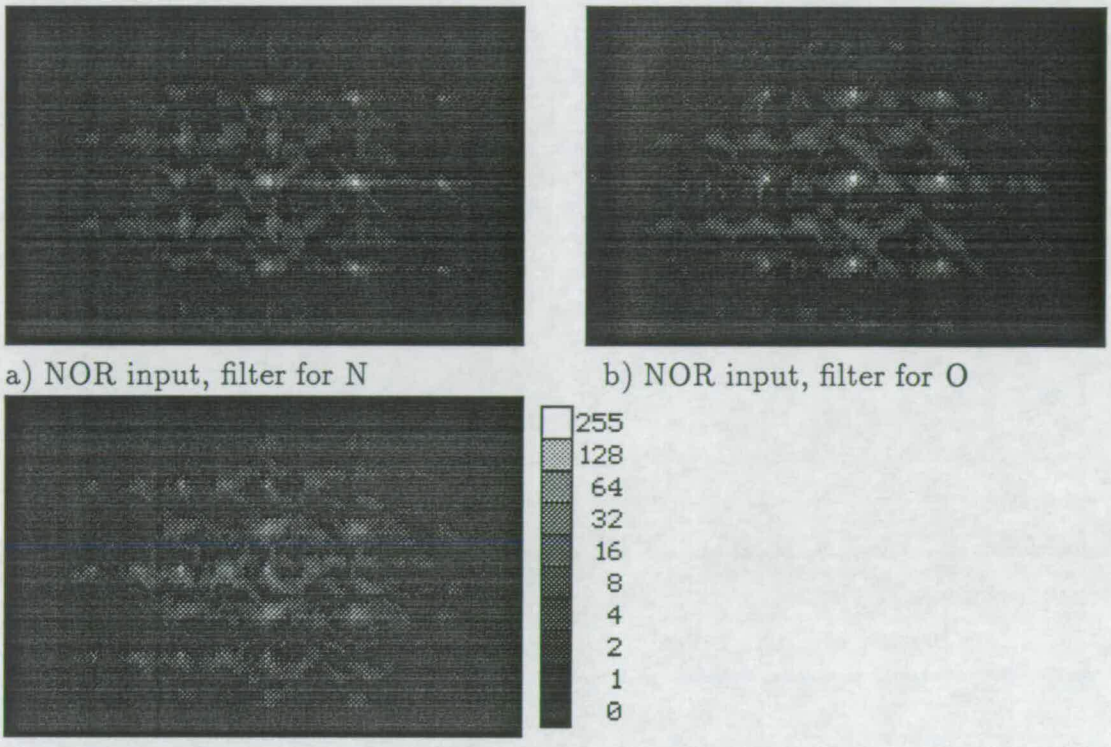

a) NOR input, filter for N

b) NOR input, filter for O

c) NOR input, filter for R

Figure 7.11: Correlation outputs for regular case BPOF with multiple letter input.

The figures show qualitatively the difficulty in isolating a correlation peak when aliasing occurs. For the regular case particularly large spurious correlation peaks occur at the position of the O when filtering for the N (shown in figure 7.11a)), and vice-versa (shown in figure 7.11b)), as a result of addition of the first order replica of the true correlation peak and the cross correlation between the letter and the transform of the filter. More quantitatively, the spurious peak at the O when filtering for the N has an intensity of 0.912 units, or 66% of that of the true peak, compared with the first order replicated correlation peak to the right of the N which has an intensity of 0.418 units, or 30% of that of the true peak. Clearly the interaction of the replicas and the cross correlations is significant and leads to problems in identification of multiple letters when using a regular BPOF matched to an individual letter's space-bandwidth product.

Identification of the letters is straightforward using the discretely randomised BPOF, as the correlation plane exhibits only one sharp peak at the position of the

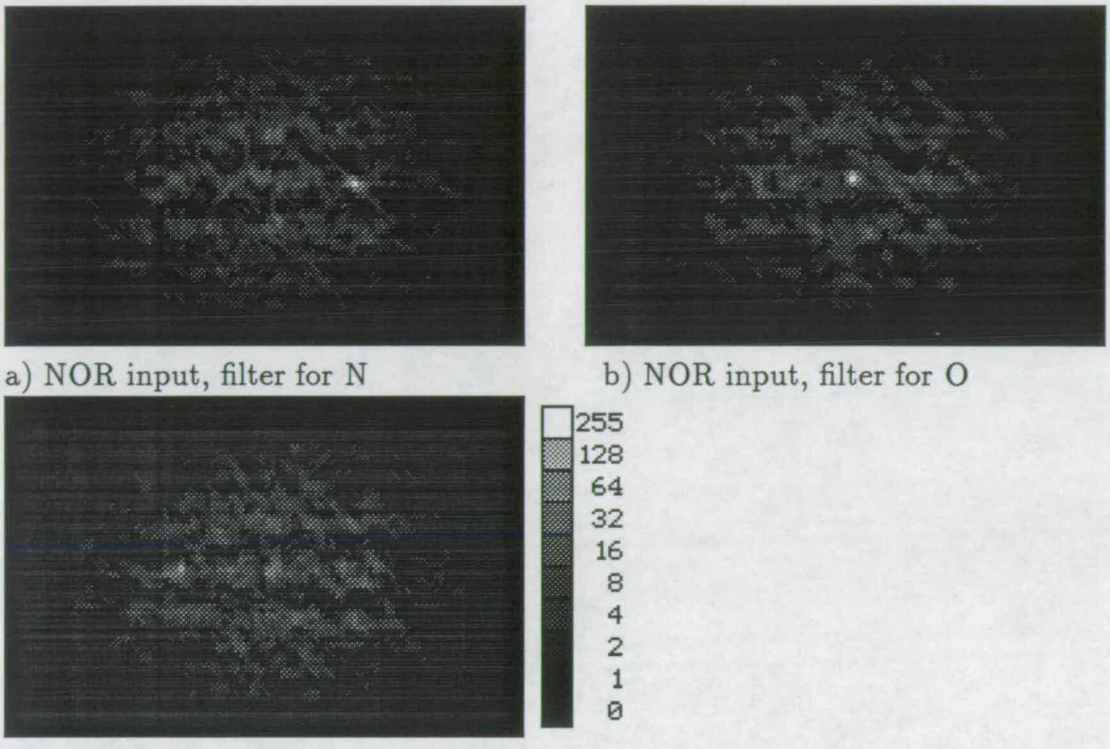

a) NOR input, filter for N

b) NOR input, filter for O

c) NOR input, filter for R

Figure 7.12: Correlation outputs for random case BPOF with multiple letter input. Contrast with figure 7.11.

letter to be detected, which is the highest intensity point in that plane for each simulation. In addition, the background noise in the simulations shows no distinct features, and is significantly lower in intensity than the correlation peaks, with a maximum background noise intensity of 0.234 units (19.2% of the correlation peak) when filtering for the N, and 0.193 units (12.2% of the correlation peak) when filtering for the O. The only exception to this is a blunt peak of 0.319 units (47.8% of the correlation peak) when filtering for the R– however this should be compared to when the regular BPOF for the R is used: 8 distinct peaks from spectral orders and 6 broader high intensity peaks from cross correlations with the O result, the highest 5 of which are 69.4%, 68.2%, 65.1%, 50.1%, and 47.1% of the correlation peak.

As expected, the correlation peaks for the BPOFs were higher than for the AEBPOFs, and the correlation outputs for the BPOFs did not contain a ver-

sion of the input scene.

7.2 Optical implementation

The correlation experiments simulated above were repeated optically where possible to verify the results. A 4-f correlator was used, with input through a uniformly illuminated binary transparency, and the Fourier plane filters implemented using the SLMs described in Chapter 6. As a 16×16 amplitude modulating regular pixel SLM of the appropriate dimensions was not available, the AEBPOF experiments which used the regular case filter were not repeated. The optical implementation used the SLMs decribed in Chapter 6: note that the special case of the discrete position randomisation scheme was not used for the randomly pixellated BPOF because it was not necessary in order to achieve the optical implementation– as an opaque blocking layer is necessary in the inter-pixel regions to prevent light passing through areas of the filter which cannot be modulated in a controlled manner, the pixels on the SLM were made as large as possible and the blocking layer used to define the transmissive pixels according to the randomisation scheme. For the amplitude modulating configuration the unswitched areas do not transmit (in the idealised case) and the opaque blocking layer is unnecessary if the crossing of electrodes defines the pixels, as is possible when the special case of the discrete randomisation scheme is used.

7.2.1 The optical correlator

The correlator used is shown schematically in figure 7.13, and is a folded version of the classical 4-f correlator architecture described in Chapter 2.

The lenses used were 10cm diameter achromatic doublets, with 1m focal lengths. The long focal lengths make folding of the optical path necessary if the correlator is to fit into a reasonable area. Illumination is provided by a 10mW He-Ne laser whose

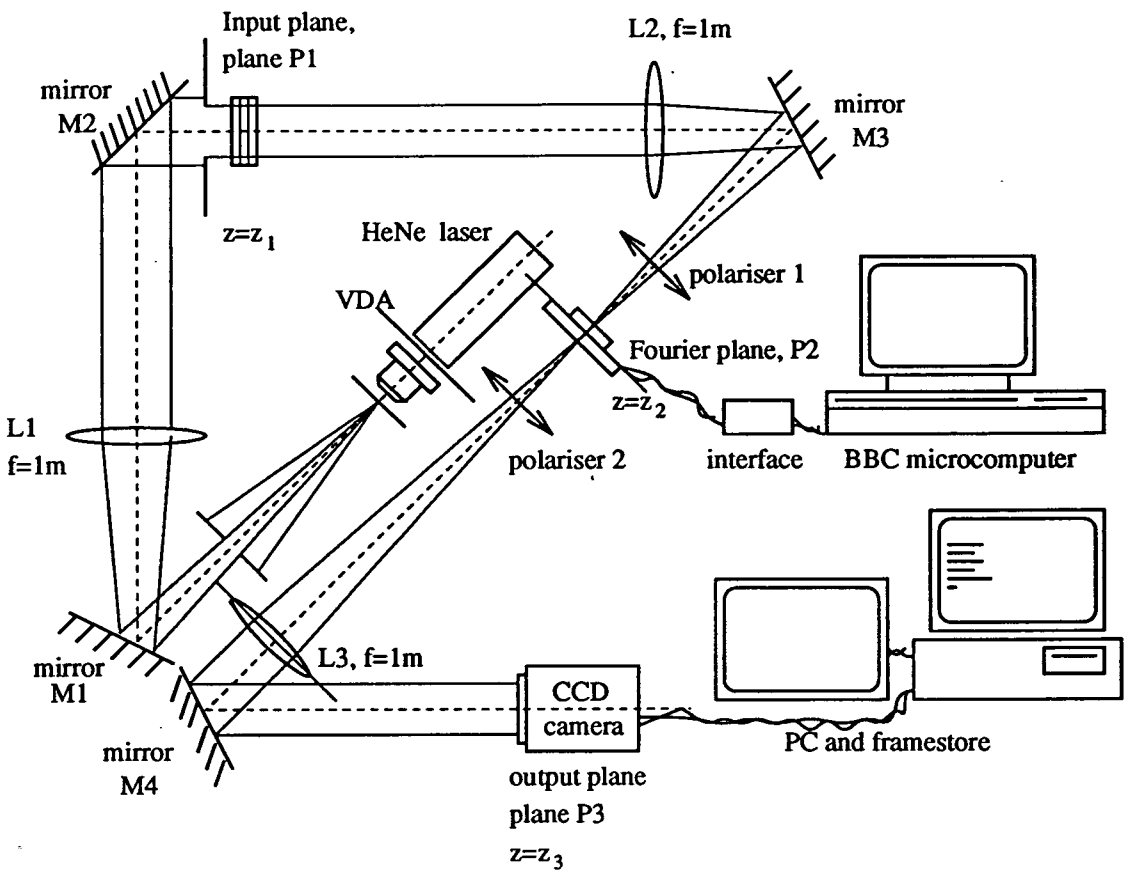

Figure 7.13: Optical correlator– folded 4-f configuration

unexpanded beam first passes through a variable density attenuator (VDA), included to provide convenient control of the illumination intensity. The beam is expanded by a spatial filter assembly, and passed through a circular aperture stop approximately 30cm in front of the spatial filter which passes only the central region (approximately 2cm in diameter) of the expanded beam. Following reflection from mirror M1 this illumination is collimated by lens L1 and reflected by mirror M2 towards the input plane. A further stop in front of the input plane restricts illumination to the area of the input transparency, which is mounted in the liquid gate described in Chapter 4. The Fourier transform of the input is formed at plane P2 by lens L2 and lies on the filter SLM, which has its central pixel (relative to the 15×15 region used for correlation purposes) centered on the optic axis of the correlator. Linear polarisers 1 and 2 respectively polarise the light incident on the filter SLM, and analyse the light leaving the SLM, with the orientation of

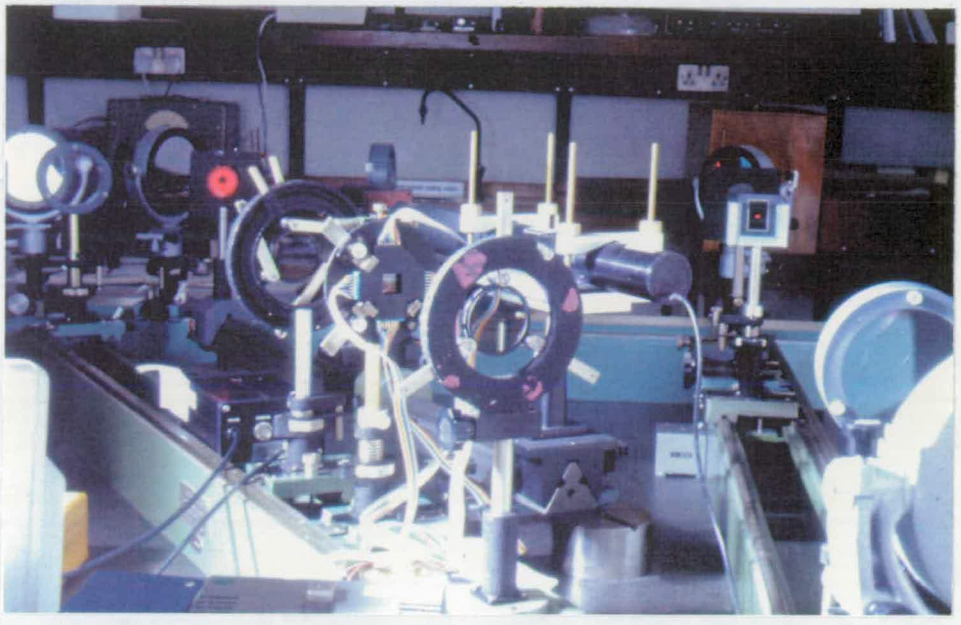

Figure 7.14: Detail of filter plane

their axes of polarisation dependent on whether the SLM is a phase or amplitude modulator. Figure 7.14 is a photograph showing this part of the processor.

Filter patterns are written to the SLM from a BBC Master microcomputer, controlling the SLM interface. A further Fourier transform is performed by lens L3, and the resulting output image constituting the output of the correlator is formed on the focal plane of the calibrated 512×512 pixel AVCAM 405 CCD array camera. The intensity distribution measured by the camera is stored using a PIP512 framestore mounted in an IBM PS/2 personal computer, and displayed on an additional monitor.

The intensities quoted for the optical results are in arbitrary units, being the 8-bit greyscale values read directly from the framestore. Therefore 0 is equivalent to zero intensity, and 255 is equivalent to the maximum intensity that the camera can detect without saturating, with a linear mapping of intensity to greyscale in between. An uncertainty of ± 4 units is estimated for the intensity measurements.

7.2.2 Alignment of components

Careful alignment of the components is necessary, e.g. to prevent aberrations introduced by lenses not centered on the optic axis, and scale mismatches if the optical path between object and transforming lens is not f. It is also crucial to align the input with the SLM so that the (x, y) and (ξ, η) axes that they define line up– there is a rotational degree of freedom which must be accounted for in addition to centering the SLM on the optic axis. This can be done by placing either the liquid gate at the input or the SLM on a rotatable mount. If the SLM is placed on a rotatable mount, any rotational adjustments necessitate adjustment of polarisers 1 and 2, and adjustment of the SLM in the ξ and η directions if it is not perfectly centered on the mount. Instead the rotational adjustment was placed at the input, which removed the need to adjust the polarisers, or make (ξ, η) positional adjustments, by exploiting the shift invariance of the Fourier transform. If the liquid gate is wedged, however, leading to a linear variation in phase with position over the input, then the shift theorem shows that the Fourier transform of the input will be offset from the optic axis, resulting in the need to make (ξ, η) adjustments of the SLMs position when the liquid gate is rotated. In practice these adjustments proved to be very small.

The input objects

Input transparencies showing the letters N, O, and R individually, or side by side as in figure 7.2 were fabricated by photoreduction of a large negative master onto high resolution, high contrast lithographic film. The masters were prepared using a laser printer, and from the same files used to represent the inputs in the simulations.

The photoreduction scale is chosen so that each letter has its Fourier transform critically sampled by the SLM in the Fourier plane. For an SLM with pixel pitch $\alpha = 500\mu$m, the separation of the spectral orders in the output will be $\lambda f/\alpha$ (see Chapter 4), where f is the focal length of the transforming lens and λ is

the wavelength of the coherent illumination. In the correlator used, $f = 1$m and $\lambda = 632.8$nm, giving the separation of the spectral orders as 1.266mm. As this correlator is a unit magnification optical system the required height of the photoreduced letters is also 1.266mm. In practice the letters on the transparencies were measured to be (1.23±0.02)mm high, leading to a slight mismatch of the Fourier transform with the SLM.

7.2.3 The filter patterns

The filter patterns written to the SLMs were the same as used in the simulations. These patterns are appropriate to an idealised optical system– in practice the system deviates from the ideal and we can expect a degradation in the correlation performance, such as a reduction in the correlation peak height relative to the background noise. These deviations will be considered in relation to the optical results.

7.2.4 Optical results for AEBPOFs

With the components of the bench carefully aligned (to within small fractions of a mm when aligning the SLM relative to the optic axis) a series of experiments were carried out where the filter patterns shown in figure 7.5 were written in turn to the amplitude modulating 'special case' discretely randomised SLM for inputs of each individual letter and the three letters side by side.

The intensity distributions in the output plane were recorded using the calibrated array camera and framestore, with the camera controller set with gamma=1 and on manual gain to prevent nonlinear and unquantified scaling of intensities. The outputs are displayed using xv with gamma=1.5, for comparibility with the simulations.

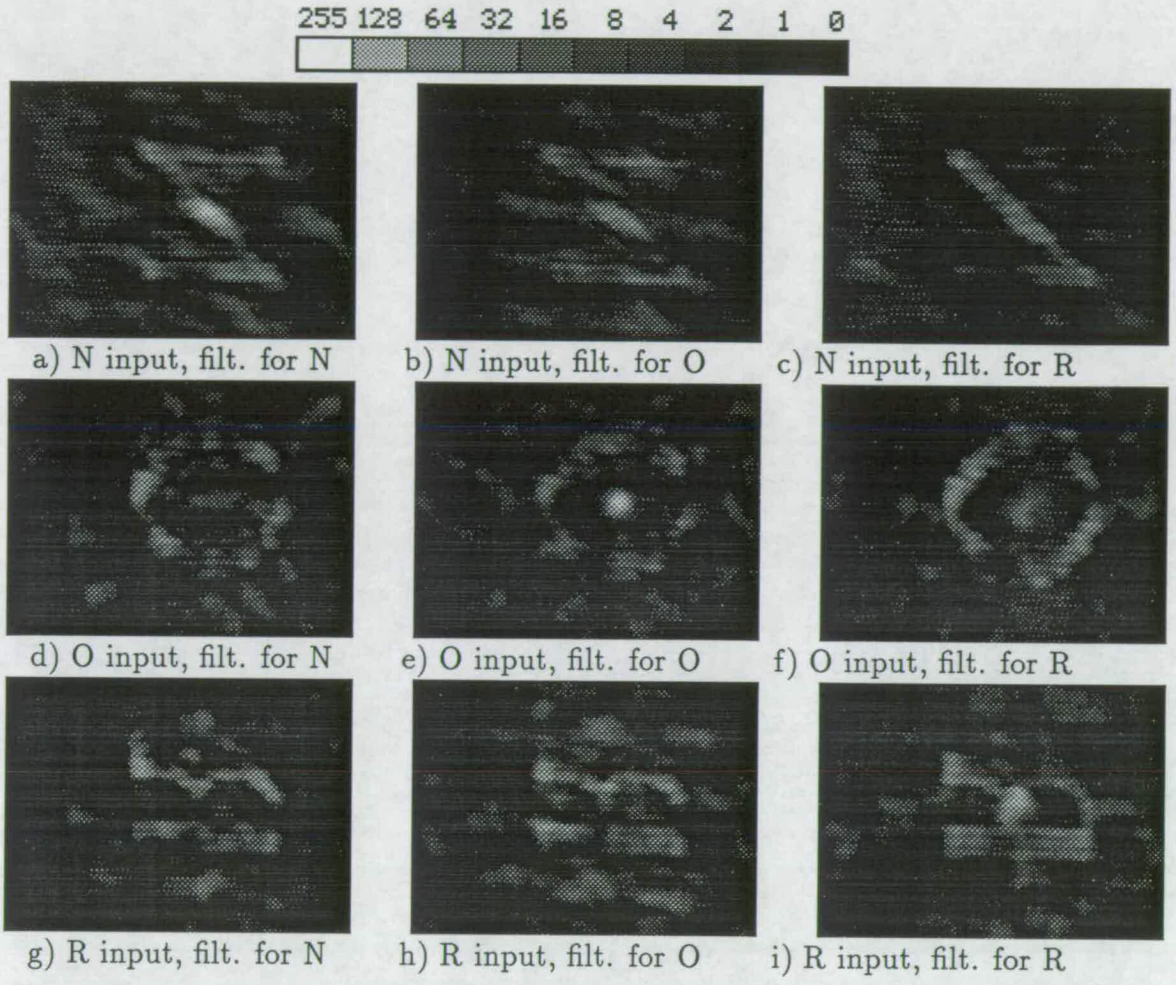

Figure 7.15: Optical correlation outputs for special case AEBPOF with single letter input. Compare with figure 7.5.

Single letter as input

The correlation outputs with the single letter inputs are shown in figure 7.15 and should be compared with figure 7.5. The gain on the array camera was held constant throughout. The left column shows the correlation outputs for the various letters with the filter for N, the middle column with the filter for O, and the right column with the filter for R. The figures show the central 256×256 pixels of the camera's output.

The optical correlation results are in very good visual agreement with the simulations–

Input	Randomised filter for...		
	N	O	R
N	188±4	59±4	56±4
O	69±4	151±4	74±4
R	93±4	85±4	118±4

Table 7.5: Table of optical correlation results for single letter input using AEBPOF

the rotated version of the input appears in the output as before, correlation peaks are apparent for the letters matched with their corresponding filters, and the background detail is very similar for the simulated and optical results.

Table 7.5 shows the maximum intensities in each of the correlations in figure 7.15. As expected the maximum intensity for each letter occurs at the correlation peak when its corresponding filter is used, and the maximum intensity for each filter occurs at the correlation peak for its corresponding letter.

The true correlation peaks are not as high relative to the cross correlations with non-target letters as they are in the simulations. This is reasonable given imperfections in the correlator, some difficult to quantify such as scattering from dust or scratches on the optical components, misalignment of components, slight non-uniformity of illumination of the target, or others which we can measure such as thickness variation of the SLM (measured in Chapter 6), or the mismatching of the input's Fourier transform to the SLM due to the input being slightly undersized.

The mismatch in scale alone is enough to account for the relative reduction in size of the correlation peaks, recalling that Casasent and Psaltis [31] have noted a reduction in signal to noise ratio (SNR) of as much as 4 times for a 2% scale mismatch. Here the mismatch is (3.0±1.6)%, and regarding the mean of the non target cross correlation for each filter as a crude measure of noise (plausible with reference to the figures) the SNR when filtering for N is $\approx$ 2.5 (3.5 in the simulation), for the O $\approx$ 2 (4 in the simulation) and for the R $\approx$ 2 (3 in the simulation), rounding to the values to the nearest half to reflect their qualitative nature.

The largest correlation peak for the optical results is for the N, as opposed to for

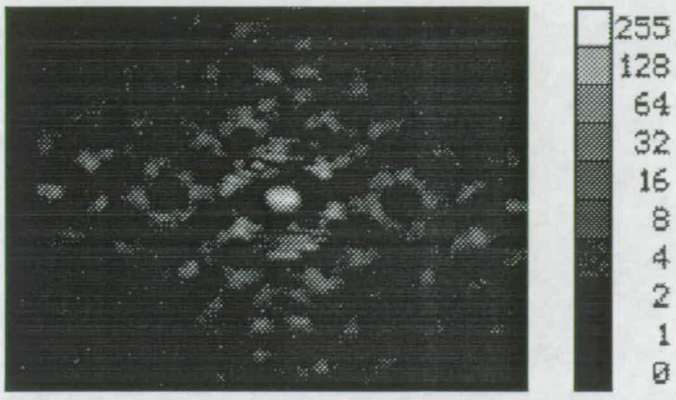

Figure 7.16: Optical correlation output for input of O with filter generated from incorrectly centered reference O. Compare with figure 7.8.

the O in the simulations. Qualitatively this can also be explained by the mismatch in scale: the main feature of the Fourier transform of the N is a diagonal line which lies along a diagonal line of predominantly 'on' pixels on the filter– here the correlation is more sensitive to rotational rather than scale mismatch; the main features of the transform of the O is a series of approximately annular rings which correspond to a similar pattern of 'on' pixels on the filter– here the transform is more sensitive to radial (and thus scale) mismatch than rotational mismatch.

The slight linear wedging of the SLM noted in Chapter 6 should not affect the heights of the correlations, only introducing a shift of the position of the output.

Incorrectly centered reference target

The filter generated with the reference letter O displaced by one point from its symmetrically centered position was used to detect the input letter O. The output is displayed in figure 7.16. and shows the two correlation peaks arising from the correlations with the 'real' and 'false' images in the filter's spectrum. This experiment was carried out with a different gain setting on the camera controller, but for comparison a correlation using the filter from the correctly centered reference was carried out under the same conditions. The correlation peak for the incorrectly centered case was 107±4 compared with 226±4 for the centered case, a ratio of

Input	Randomised filter for...		
	N	O	R
N	176±4	101±4	142±4
O	124±4	249±4	135±4
R	115±4	100±4	165±4

Table 7.6: Table of optical correlation results for multiple letter input using AEBPOF

0.47±0.03, which is in reasonable agreement with the simulation where the ratio was 0.42. Note that the uncertainties in the intensity measurements attempt to account for measurement errors, but take no account of variations due to the slight scale mismatch of the input letters with the Fourier plane filter.

Multiple letter input

The experiments for the randomly pixellated AEBPOF with multiple letter input were implemented optically. The correlation outputs are shown in figure 7.17 and should be compared with those in figure 7.7. Similarly Table 7.6 should be compared with Table 7.2. The figures show the entire output region detected by the array camera, with the gain of the camera controller held constant throughout.

These optical results are in good visual agreement with the computational simulations, and as required show only one distinct correlation peak correctly identifying the appropriate letter in each case.

The comments for the single letter input regarding the relative values of the correlation peaks and the non-target cross correlations are equally valid here. Using the non-target cross correlations as a measure of noise, we find that the resulting SNR using the filter for the N is ≈ 1.5 (2 in the simulations), for the O is ≈ 2.5 (4 in the simulations) and for the R is ≈ 1 (1.5 in the simulations).

A feature to note in the optical simulations is that the indistinct background noise extends a greater distance away from the positions of the three letters. This is as a consequence of the pixels in the SLM being slightly less than half the size of the

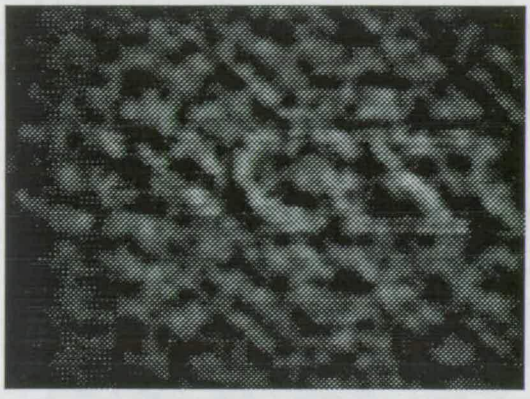

a) NOR input, filter for N

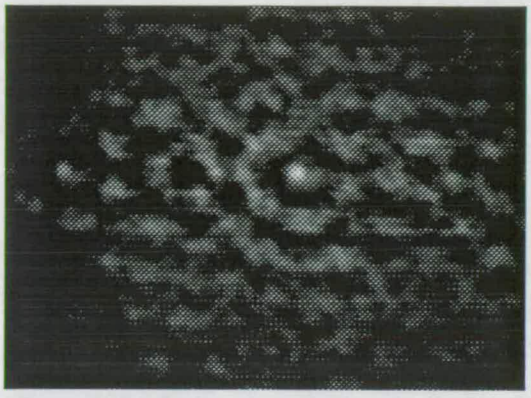

b) NOR input, filter for O

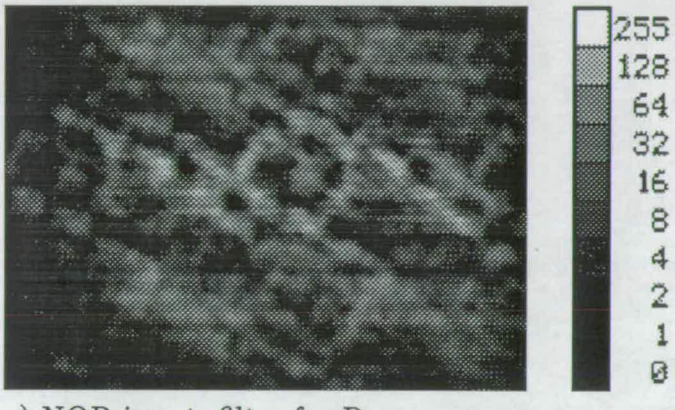

c) NOR input, filter for R

Figure 7.17: Optical correlation outputs for special case AEBPOF with multiple letter input. Compare with figure 7.7.

underlying pixel pitch, leading to intensity contributions from the resulting faint second order replicas. Note that no first order replicas are apparent, as expected.

7.2.5 Optical results for BPOFs

The experiments for the BPOFs were implemented optically using the phase modulating SLMs with regular and discretely randomised pixel arrays described in Chapter 6. These SLMs were used to display the filter patterns shown in figures 7.9 and 7.10. The correlator architecture is as for the AEBPOFs, other than the SLM and the orientation of polarisers 1 and 2, which are aligned as appropriate

to the phase modulators (see chapter 3, section 3.3.2).

The results are not as impressive as for the AEBPOFs, due to the limitations of the SLMs used. The regular pixel SLM suffered from having 2 adjacent columns (the 3rd and 4th from the right, with respect to the filter pattern figures) which did not switch due to breaks in the ITO electrodes, and from significant thickness variation and curvature across its width which gave the SLM undesired lens-like properties. The AEBPOF also has the advantage that the two binary states can be achieved with little difficulty due to the sharp threshold characteristic of the transmitted intensity with respect to applied voltage, described in Chapter 3, which gives well defined 'off' and 'on' states over a range of several volts either side of the threshold voltage. However, the BPOFs which use the field induced birefringence effect are highly sensitive to variations in thickness and applied voltage, as the difference between states is proportional to $\Delta n.d$, with d thickness, and Δn (the difference between the two states in refractive index encountered by the incident light) a smoothly varying function of applied voltage and thus sensitive to non uniform resistance of the electrodes or poor electrical contacts.

The discretely randomised SLM also had a column which did not switch, but was comparitively flat. In addition it exhibited a slight misalignment of the switchable pixel areas with the aluminised blocking layer which was meant to mask off the interpixel regions. This is shown in figure 7.18. The non-modulating transmissive areas such as the faulty columns, spaces between electrodes visible for adjacent pixels, and regions left visible by misalignment of the blocking layer will all lead to features of the input being transmitted unmodulated, and give rise to the rotated versions of the inputs faintly visible in the outputs below.

Single letter as input

The correlation outputs for the single letters as input are shown in figures 7.19 and 7.20 for the regular and discretely randomised cases respectively. The outputs show the entire output region detected by the array camera, to show the extent of the aliasing for the regular case.

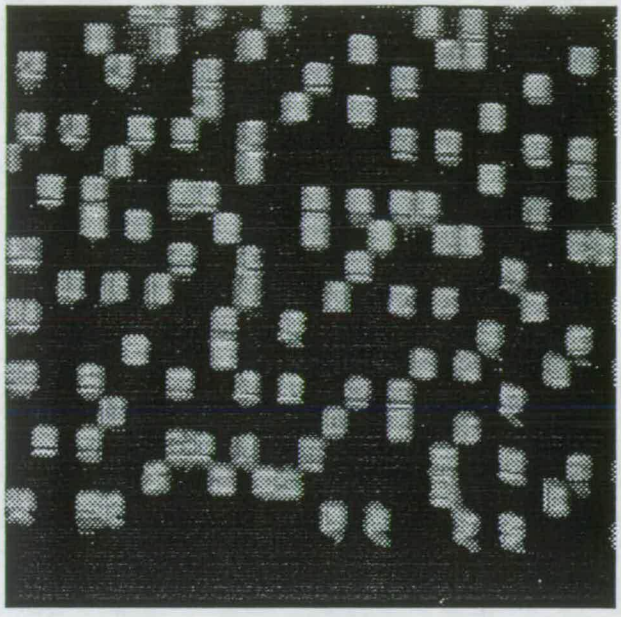

Figure 7.18: Misalignment of pixels and aluminised mask

Considering the outputs using the regular BPOF first, it is possible to discern peaks for the N and the O using their appropriate filters. No output images for the R were framegrabbed as no discernible peak or features were evident. This is consistent with earlier results where the correlations were stronger for the symmetrical letters O and N. The gain of the camera was constant for the results with the N as input, and constant at a different value with O as the input. Comparing the optical results with the simulations in figure 7.9 we can note that replication of the correlation output does take place in the optical results, but no comparison of the low level background detail is possible due to the presence of versions of the input in the output, as a result of the non-modulating regions of the SLM. When filtering for the O even the correlation peak was swamped by the spurious versions of the input, as is apparent from Table 7.7.

The results for the discretely randomised BPOFs are much more consistent with the idealised simulations, with reference to both figure 7.20 and Table 7.7. These results were all obtained with the gain of the array camera held constant at a single value, and the figure shows the the detail of the central 256×256 pixel region of the camera's output. Visually the results are in broad agreement with the simulations, with correlation peaks evident when the inputs are matched with the

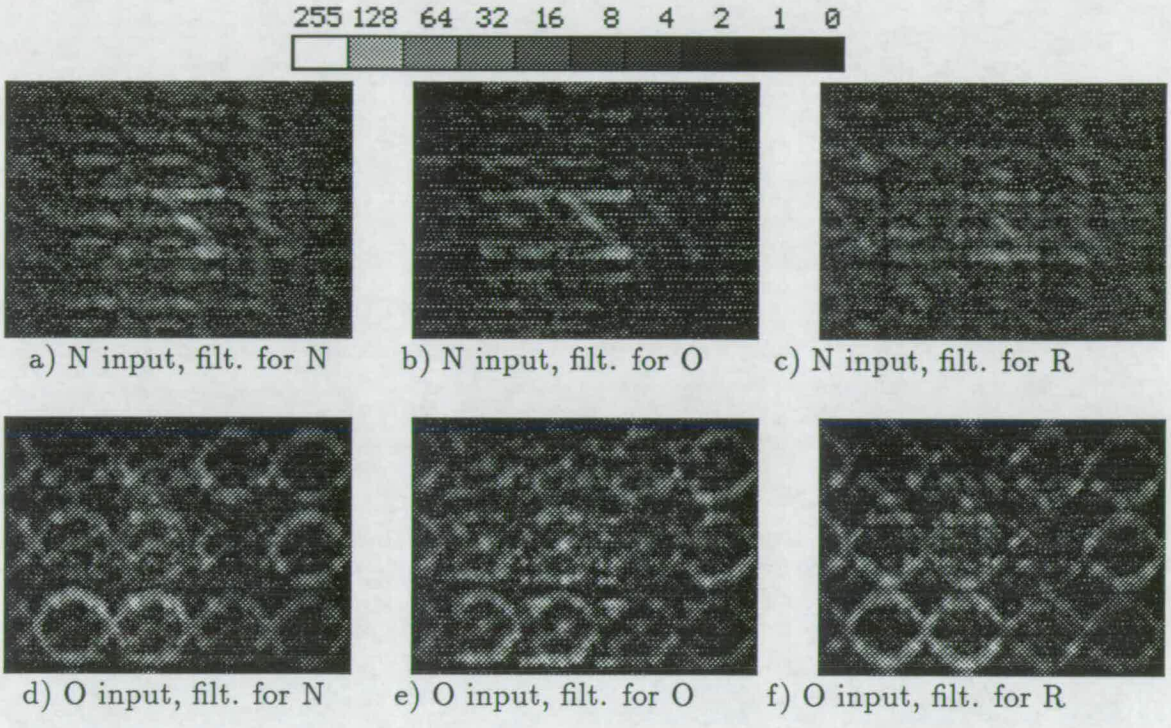

Figure 7.19: Optical correlation outputs for regular case BPOF with single letter input. No recognisible features were detected using the filter for R. Compare with figure 7.9.

Input	Regular filter for...			Randomised filter for...		
	N	O	R	N	O	R
N	188±4	129±4	123±4	162±4	101±4	111±4
O	213±4	174±4†	200±4	94±4	187±4	88±4
R	No discernible features			99±4	118±4	152±4

†Noise maximum nearby of 181.

Table 7.7: Table of optical correlation results for single letter input using BPOFs

appropriate filters, but with faint versions of the input apparent. The correlation peak intensities also show the same trend as shown by the simulations in Table 7.3, with the correlation for the O using its filter greater than the correlation for the N using its filter, which in turn is greater than the correlation for the R using its filter. The correlation for each letter is also greatest using its corresponding filter, and for each filter is greatest with its corresponding letter. Using the non-target cross correlations as a measure of noise we note that the SNR is severely

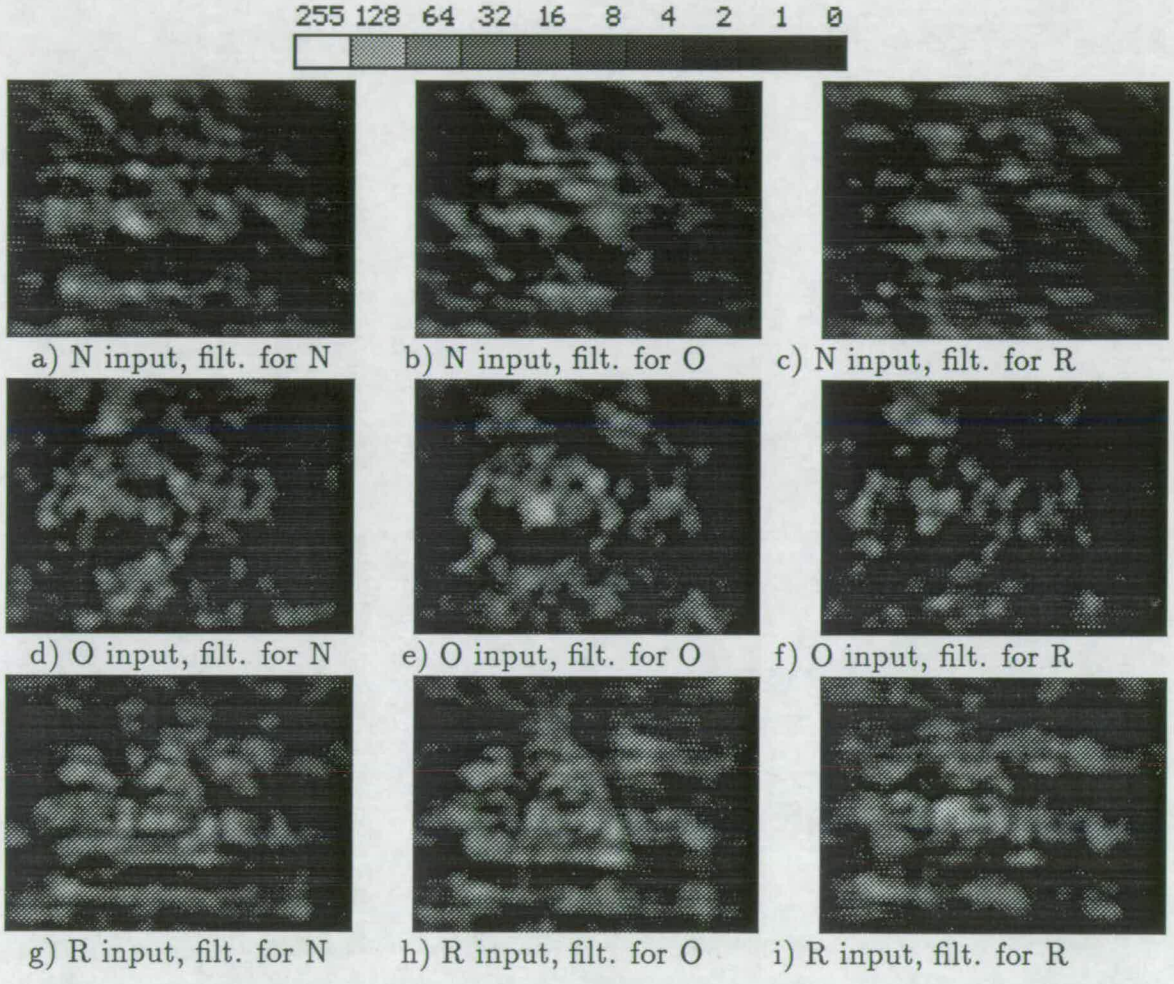

a) N input, filt. for N b) N input, filt. for O c) N input, filt. for R

d) O input, filt. for N e) O input, filt. for O f) O input, filt. for R

g) R input, filt. for N h) R input, filt. for O i) R input, filt. for R

Figure 7.20: Optical correlation outputs for discretely randomised case BPOF with single letter input. Compare with figure 7.10.

degraded for the optical results– SNR when filtering for the N is ≈ 1.5 (9 in the simulations), for the O ≈ 1.5 (11 in the simulations), and for the R ≈ 1.5 (4.5 in the simulations)– but this is as expected given the scale mismatch and the non-modulating transmissive regions of the SLM which serve to increase the noise and reduce the correlation efficiency.

Input	Randomised filter for...		
	N	O	R
N	145±4	74±4	102±4
O	118±4	154±4	103±4
R	106±4	86±4	163±4

Table 7.8: Table of optical correlation results for multiple letter input using discretely randomised BPOF

Multiple letter input

The correlation outputs using the input of the three letters N, O and R side by side are shown in figures 7.21 and 7.22 for the regular and discretely randomised BPOFs respectively.

Again the optical results for the regular BPOF are inconclusive, due to the failings of the SLM. The replication is readily apparent, but the replicated versions of the input swamp the correlation peaks that should appear in the output. Nevertheless, when using the filter for the N a correlation peak of 113 units is present at the position of the N, which is larger than the peak occuring at the O of 103 units. However both these peaks are *less* than surrounding regions of noise which have intensities of 144 and 119 units. No distinguishable features were apparent when the filter for the R was used.

A complete set of results was obtained for the discretely randomised BPOFs which were in reasonably good agreement with the simulation. The gain of the array camera was constant throughout these results. Comparing figure 7.22 and the simulations in figure 7.12 we can see that the sets of results are in very good visual agreement. Distinct correlation peaks identify the positions of the letters without any perceptible aliasing, and the background detail is also very similar, though a faint letter O is apparent.

The numerical results in Table 7.8 are also broadly as expected, though it is unclear why the R should have the largest correlation value. The table shows that the highest intensity recorded for each filter is from the correlation peak at

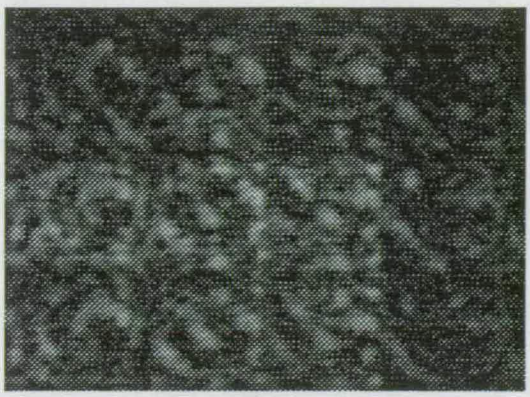

a) NOR input, filter for N

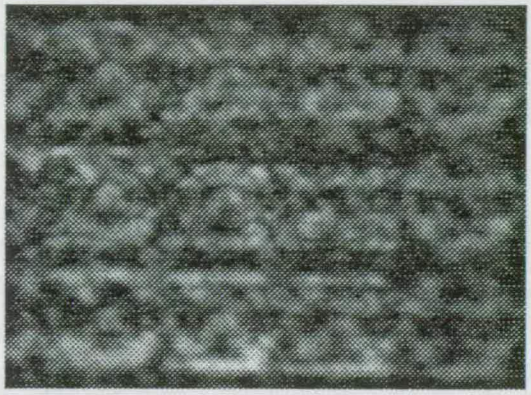

b) NOR input, filter for O

Figure 7.21: Optical correlation outputs for regular case BPOF with multiple letter input. Again, filtering for the R was unsuccessful. Compare with figure 7.11.

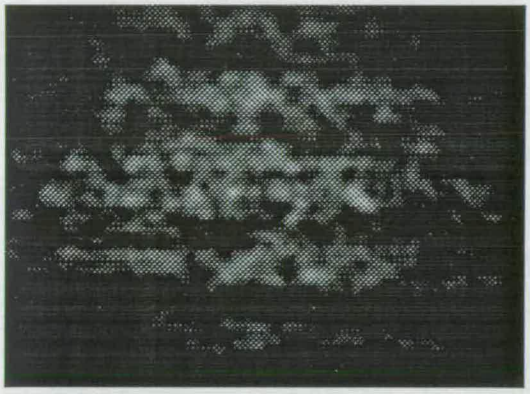

a) NOR input, filter for N

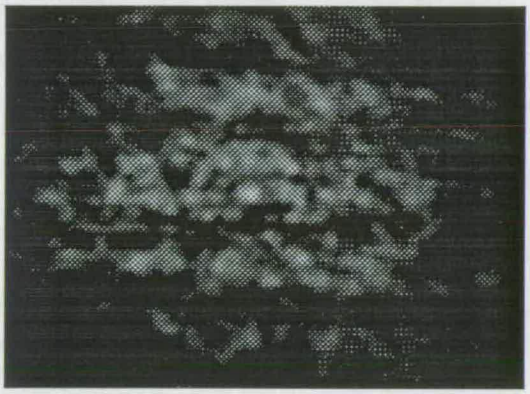

b) NOR input, filter for O

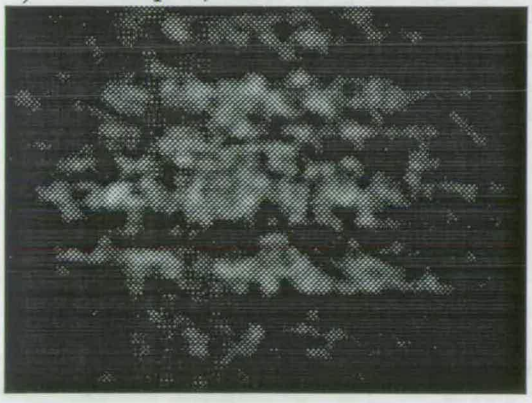

c) NOR input, filter for R

Figure 7.22: Optical correlation outputs for discretely randomised case BPOF with multiple letter input. Compare with figure 7.12.

its corresponding letter, and for each letter from the correlation peak obtained using its corresponding filter. The intensities listed in the table for the non-target cross correlations are the maximum values of the generalised background noise in the region of that letter. Averaging them for each filter and dividing the true correlation intensity by the average gives us a measure of the SNRs– ≈ 1.5 using the filter for N (7 in the simulations), ≈ 2 using the filter for O (8.5 in the simulations) and ≈ 1.5 using the filter for R (2.5 in the simulations)– which are reasonable figures given the scale mismatch of the input and the failings of the SLM.

7.3 Conclusions and summary

This chapter has generalised the concept of randomising pixel positions to generate filters which can successfully detect targets matched in SBP to the filter, in large input scenes with a greater SBP than that of the filter, without aliasing or replication of the correlation output as is conventionally found when the filter's pixels are on a regular grid. This has been demonstrated by simulations and optically for both amplitude-only and phase-only filters.

The computational simulations with the single letters matched in SBP to the filter have shown that the randomisation of the pixel positions does not degrade the performance of the correlator with respect to correlation peak height and width. Also, the redistribution of the energy in the spectral orders of the filter's spectrum into a diffuse low-level background does not lead to distinct replicas being formed, aiding identification of the correlation peak. Conversely, filtering with the regularly pixellated filters shows replicated outputs, including spurious correlation peaks which are more than 40% of the height of the true peak.

The simulations with the multiple letter input with a SBP well in excess of that of the filter have demonstrated the problems arising from aliasing. Use of the regular filters gave rise to replicated correlation peaks which interacted with coincident cross correlations, resulting in spurious correlation peaks which lay within the

region of interest in the correlation plane and were sometimes even in excess of the true correlation peak. The randomised filters performed much better, with no visible replication of the correlation peaks which always correctly identified the appropriate target. The computational simulations also showed, as expected, the superior performance of BPOFs compared with AEBPOFs, *viz.* higher correlation peaks, with no version of the input appearing in the output plane. The importance of correctly centering the reference target has also been demonstrated, both in simulation and optically.

The optical results were in good agreement with the computational simulations, with discrepancies explicable by imperfections in the experimental apparatus. The optical results using the AEBPOFs matched the simulations particularly closely, reflecting the fact that uniform binary states are easier to achieve on the amplitude modulating device than the phase modulating nematic LC SLMs, which are highly sensitive to variations in thickness and applied voltage. The optical results demonstrated that a low resolution SLM may be used successfully in a correlator, and more importantly, helped verify that by use of an SLM with pixels of size and positioning specified by a suitable randomisation scheme, the SBP of the filter need be matched only to that of the target to be detected rather than to that of the entire input scene.

Chapter 8

Discussion

In this chapter the work of the thesis is summarised, and the possible directions of future research considered.

8.1 Summary of the work

8.1.1 Consequences of pixellation

Currently available spatial light modulators suitable for use in optical correlators are almost always pixellated, with identical pixels regularly spaced on a square grid, and binary in operation. The analysis of the effects of this pixellation has shown that it gives rise to spectral orders in the Fourier spectrum of the filter which leads to replication in the output of the optical correlator. This in itself need not be problematic if the replicas do not overlap, as interest can be restricted to the central replica. This corresponds to the filter sampling the Fourier transform of the input at the Nyquist frequency or greater, or equivalently if the input is also pixellated, to the filter having a space bandwidth product greater than or equal to that of the input. If this is not the case the replicas will overlap and aliasing will occur, distorting the information in the output.

The analysis of the Fourier spectrum of an array of regularly spaced pixels has also shown that the width of the spectral orders is dependent on the overall size of the array, and that the size of the higher spectral orders relative to the zero order is dependent on the spacing of the pixel centres and the pixel size. The latter point led to consideration of randomising the pixel positions as a further means of reducing or removing the higher spectral orders, thus reducing or removing the effects of aliasing if the filter has a lower space bandwidth product than the input. This is desirable if the target to be detected does not occupy the entire input, as a filter with the lower space bandwidth product of the target would be adequate to identify it, were it not for aliasing of the correlation peak.

8.1.2 Analysis of random pixel arrays

By incorporating a random displacement of the pixel position into the transmission function of an array of transmissive pixels, all 'on', an analytical expression has been derived for the power spectrum of the average of the ensemble of possible arrays. This expression contains factors modulating the higher spectral orders, which are determined by the probability distribution function govering the random displacement of the pixel. These factors are ≤ 1, showing that further attenuation of the higher spectral orders is possible, by a suitable choice of randomisation scheme. The energy removed from the higher spectral orders is redistributed in a diffuse halo, with a *light depletion zone*, surrounding the zero order. This should not cause problems with aliasing for any filter of a reasonable size as it is at least Q^2 lower in magnitude than the zero order, where Q^2 is the number of pixels in the filter. Two probability distribution functions were investigated. The first had the pixels positioned with uniform probability at any point within a square cell of side equal to the pixel pitch, subject to the constraint that all of the pixel should lie within the cell– this helped attenuate the higher spectral orders, but still left significant first orders. The second (the discrete case) had the pixels restricted to a choice of one of four positions in the cell, and allowed all the distinct higher spectral orders to be removed when used with square pixels of side half the underlying pixel pitch.

These results were derived for the ensemble averaged power spectra of the pixel arrays and verified by simulation using ensembles of 50 arrays. Individual members of the ensembles were also investigated, including a special case of the discretely randomised case chosen to simplify fabrication of an amplitude modulating SLM. Their power spectra were found to be in close agreement with the ensemble averages, as desired. Further simulations of letters of various sizes being imaged through the arrays showed the effectiveness of the discrete case randomisation at removing replicas and preventing aliasing. Transparencies were generated from the individual members of the ensemble and successfully used to verify the results optically.

8.1.3 Fabrication of transmissive matrix addressed nematic liquid crystal SLMs

Custom SLMs were fabricated to use as Fourier plane filters in optical correlation experiments. The SLMs were 16×16 arrays of pixels, formed by the overlap of sets of transmissive conducting electrodes sandwiching a layer of nematic liquid crystal, and electrically addressed using a multiplexed drive scheme. Amplitude-only and phase-only SLMs were fabricated which implemented regular and randomised arrays.

Fabrication of the SLMs became reliable and relatively straightforward. The technology required was quite simple, with the sputtering of the ITO requiring the most sophisticated equipment: however, even this could be avoided by etching glass obtained already coated with ITO, instead of the lift off procedure used here. Assembly of the SLM demanded great care, both in registering the pixels accurately, and ensuring uniformity of thickness.

The performance of the amplitude modulator fabricated was extremely good, with a contrast ratio of (115 ± 12):1, all pixels working, and only approximately one wavelength of curvature and wedging visible over the array of pixels. The performance of the phase modulators was less satisfactory, with one and two columns

not switching for the regular and randomised SLMs respectively. A modulation depth of π was demonstrated, but the sensitivity of the modulating effect (field induced birefringence) to variations in thickness and voltage was thought to lead to lack of uniformity over the arrays. The interferometric measurements showed particularly large amounts of curvature present over the regular SLM, with as much as 3 fringes of difference ($\sim 2\mu$m) from side to side of the array.

8.1.4 Correlation experiments

Removal of aliasing and replication in the output of an optical correlator when the SLM is acting as a matched filter with pixels 'off' as well as 'on' was demonstrated to be plausible, by means of a more physical interpretation of the underlying process. This was subsequently supported both by computational simulation and optical implementation using the SLMs in a series of correlation experiments. The interpretation used notes that the contributions to the spectral orders from pixels placed at two of the four sites allowed in the discrete positions randomisation are in antiphase to the contributions from the other two allowed sites: thus a necessary condition for removal of the spectral orders when using this randomisation scheme is that the 'on' pixels should be evenly distributed over the 4 allowed positions.

The experiments were carried out for binary inputs of single letters matched in space-bandwidth product to the filters, and for three letters side by side, with filtering carried out by amplitude modulating AEBPOFs or phase modulating BPOFs. The simulations for the single letter input show that the fundamental performance of the correlator is not degraded by pixel position randomisation. The simulations using the multiple letter input illustrate the problems caused by aliasing when a regular filter of too low a space-bandwidth product is used, and show the success of the corresponding discretely randomised filter in avoiding these problems, in agreement with the earlier work of the thesis. Further conclusions are in agreement with the work of other authors, such as the superior performance of the BPOFs [29], and higher correlation peaks for the symmetrical objects [26]. The optical results were in good agreement with the simulations, with any discrepancies explicable in terms of the experimental apparatus.

8.2 Future work

8.2.1 Improvements to the SLMs

The construction

Among possible improvements to the SLMs would be the use of a higher resolution lithography process when defining electrodes, to reduce further the dead space between adjacent pixels; the use of optically flat glass to make the liquid crystal cells; tighter control of cell thickness by use of 12μm spacer balls mixed in the liquid crystal (or glue holding the halves of the cell together) rather than the polyester sheet spacers; and use of more advanced liquid crystals.

The liquid crystal layer

Ferroelectric liquid crystal exhibits a much faster electro-optical effect and would reduce switching times compared with the nematic. Also, with regard to phase modulation, a method exists where binary phases of 0 and precisely π may be obtained which is tolerant of thickness variation: however, thickness variation will now lead to undesired amplitude modulation. Nematic liquid crystals still show promise with regard to 'grey scale' operation, as their electro-optical response is approximately proportional to the rms voltage applied– thus by frame addressing at such a rate that the response of the liquid crystal is integrated over N frames, N 'grey levels' may be possible, either of phase or amplitude. Grey scale operation has been shown to improve correlation performance [25], and will reduce the sensitivity of the filter to rotated versions of the input.

The device technology

SLMs with a larger space bandwidth product than used here will be better implemented as the next generation of reflective liquid-crystal-over-silicon devices, as the number of external connections can be reduced by addressing circuitry being incorporated on-chip. At present such devices suffer from poor reflectivity and flatness, but recent work on device planarisation[107] offers a solution to this problem. Planarisation also offers the promise of improved fill factors by patterning mirrors on the planarised surface, which are limited in size only by the design rules of the silicon process governing metal-metal proximity. VLSI technology also offers fast frame rates [108], necessary for rapid identification of targets, particularly if there is a large bank of filters with which to compare the input. The possibility exists to implement randomised pixel arrays on such devices. Such work is in progress, and forms part of the paper 'Specification of an optical correlator for road sign recognition' included in Appendix C. Here constraints on the positioning of pixels arise from the VLSI design rules, as opposed to ensuring continuity of electrodes in the matrix addressed case.

8.2.2 Further work with random pixel SLMs

If more time had been available it would have been interesting to evaluate the performance of the randomised pixel filter SLMs with noisy or grey scale inputs,and to attempt to optimise their discrimination performance by some optimisation procedure such as simulated annealing. Some more speculative work is suggested below.

Use of alternative randomisation schemes to provide amplitude weighting

Having demonstrated that aliasing can be removed without each pixel having a 100% fill factor, the use of variants of the discrete position randomisation already

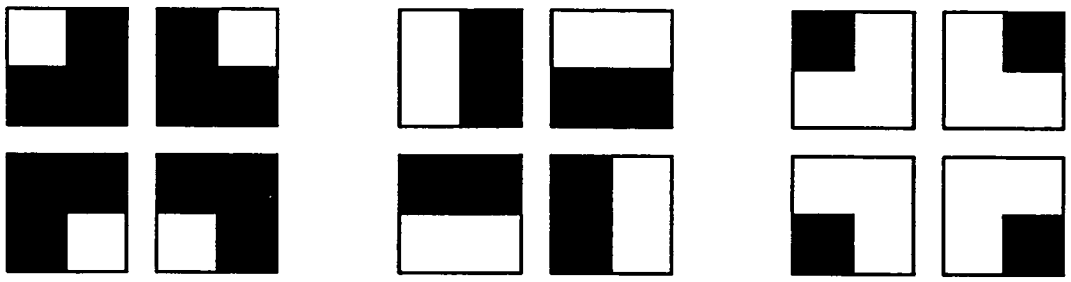

Figure 8.1: Possible 2 sub-pixel from 4, and 3 sub-pixel from 4, configurations. The 1 from 4 case described in the thesis is also shown.

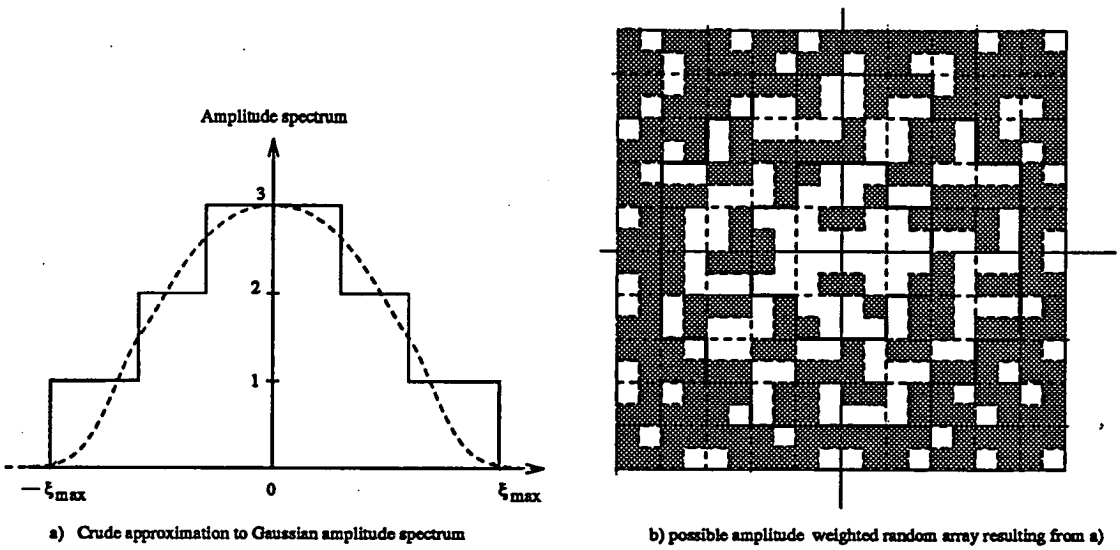

Figure 8.2: Possible amplitude weighted random array, achieved by approximating to Gaussian amplitude spectrum.

described may be considered. In particular, preliminary simulations suggest that randomly choosing blocks of two or three sub-pixels per $\alpha \times \alpha$ cell as shown in figure 8.1 should be equally effective. This in turn leads to the possibility of combining the randomisation schemes on an SLM as shown in figure 8.2, to achieve a crude weighting of the amplitudes of the transmitted spatial frequencies. That this is desirable is suggested by Gheen [109], who states that the performance of a phase only filter can be drastically improved by adding a few amplitude levels; that a single crude approximation should be equally appropriate for all target objects is suggested by Oppenheim and Lin [17], when they note that the image

reconstructed from an object's phase spectrum weighted by a simple Gaussian closely resembles the original object.

Application to binary logic/switching

An earlier attempt at generating a discretely randomised array resulted in an array where the pixels showed obvious periodicity when the array was viewed along one of the axes, and which had a power spectrum with distinct spectral orders extending along one of its axes. The suggestion here is that a means exists whereby information can be routed via certain of the first order replicas, dependent on the distribution of the transmitting subpixel in each cell of the array- with the discretely randomised array described all first orders would be 'off'; by using arrays with periodicity with respect to the axes the first orders on the x-axis, or y-axis (or both, if the array is regular) would be 'on'. Figure 8.3 helps illustrate this.

There are several points to note. First, is that the idea lends itself to existing SLMs: using the 16x16 silicon backplane SLM designed by Underwood [67], a 2×2 block of pixels could be used to code each of the α by α cells, giving an 8x8 array of cells with one pixel 'on' in each cell. This would then give a background illumination due to the 'off' state of the first order $\sim$ 64 times less than the 'on' state, a difference in states which should be simple to threshold. Further, the SLM would be capable of routing 8x8 pixel arrays when used in the Fourier plane of a 4-f optical processor, offering the benefits of parallelism. Lastly, all the 'on' first order replicas have the same intensity regardless of whether there are zero, two or four of them (to a good approximation): this would not be the case if the routing were achieved simply by writing stripe patterns on the SLM to approximate diffraction gratings, as the total number of transmitting pixels would vary depending on the desired number of replicas.

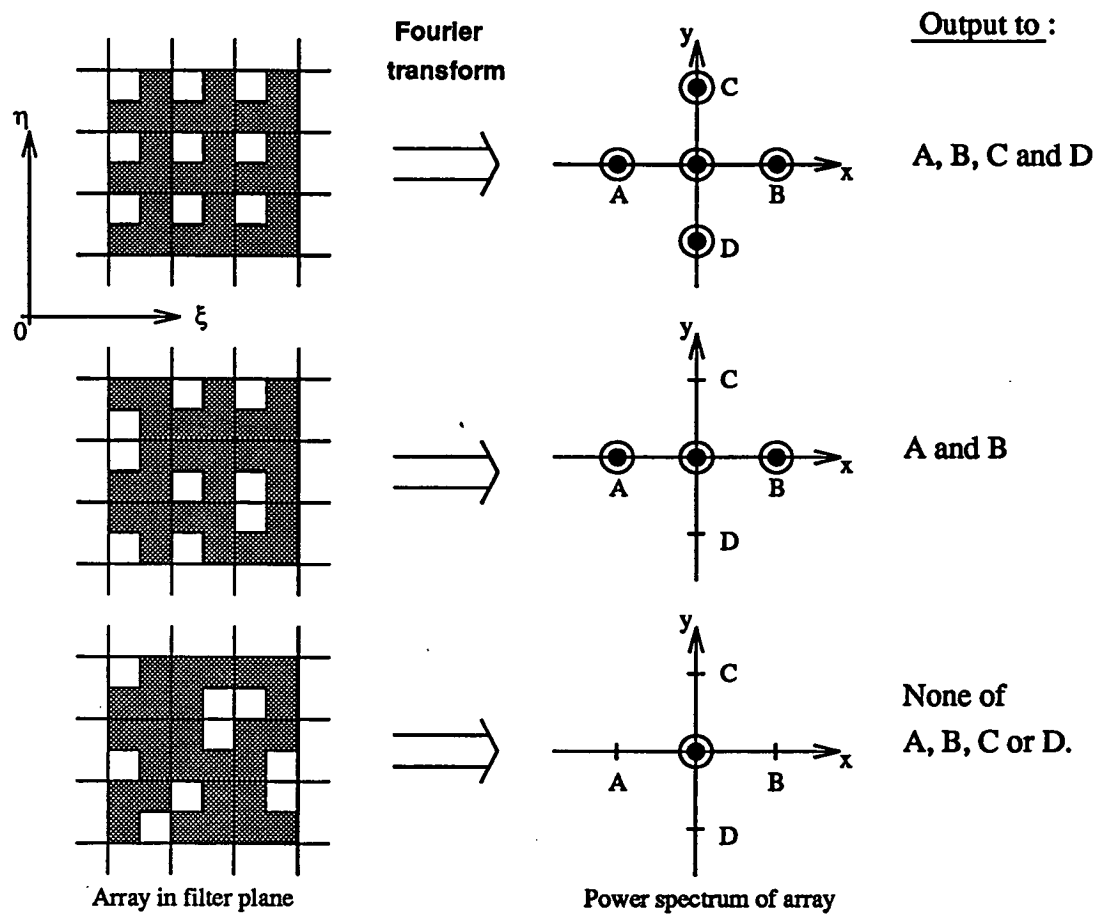

Figure 8.3: Selective removal of aliasing to provide selective fanout/routing

8.3 Summary

A study of optical correlation using pixellated spatial light modulators has been carried out. A particular problem– the requirement to match the space-bandwidth product of the filter SLM to that of the entire input rather than that of the target to be detected, to prevent aliasing in the output– has been identified, and attributed to spectral orders in the Fourier spectrum of the SLM. Randomisation of the pixel positions has been proposed as a means of reducing the spectral orders, and an analytical expression derived to show its effect. This has shown that by careful choice of the randomisation scheme, elimination of the distinct spectral orders is possible, which has been verified optically and in simulations.

Optical correlation experiments have been successfully carried out using custom amplitude-only and phase-only SLMs. The results of these experiments are in good agreement with correlation simulations which illustrate the aliasing problem when using regularly pixellated SLMs, and demonstrate the effectiveness of the chosen discretely randomised pixel array in removing the problem. Some areas of future research in related areas have been identified.

Bibliography

[1] J. Weigelt. Binary logic by spatial filtering. *Optical Engineering*, 26:28–33, 1987.

[2] B.S. Wherrett and F.A.P. Tooley, editors. *Optical Computing, Proceedings of the thirty-fourth Scottish Universities Summer School in Physics.* SUSSP Publications, Edinburgh, 1989.

[3] H.J. Caulfield and G. Gheen, editors. *Selected papers on Optical Computing.* SPIE Optical Engineering Press, Bellingham, 1989.

[4] J.W. Hardy. Active optics: a new technology for the control of light. *Proceedings of the IEEE*, 66:651–697, 1978.

[5] G.O. Reynolds, J.B. DeVelis, G.B. Parrent, Jr., and B.J. Thompson. *The New Physical Optics Notebook: Tutorials in Fourier Optics.* SPIE Optical Engineering Press, Bellingham, 1989.

[6] E.G. Steward. *Fourier Optics an introduction.* Ellis Horwood, Chichester, second edition, 1987.

[7] J. Gaskill. *Linear Systems, Fourier Transforms and Optics.* Wiley, New York, 1978.

[8] R.N. Bracewell. *The Fourier Transform and its Applications.* McGraw-Hill, New York, 1986. second edition, revised.

[9] J.W. Goodman. *Introduction to Fourier Optics.* McGraw-Hill, New York, 1968.

[10] K.G. Birch. Spatial filtering in optical data-processing. *Rep. Prog. Phys.*, 35:1265–1314, 1972.

[11] H. Stark. *Applications of Optical Fourier transforms.* Academic Press, New York, 1982.

[12] L.J. Cutrona, E.N. Leith, C.J. Palermo, and L.J. Porcello. Optical data processing and filtering systems. *IRE Transactions on Information Theory*, pages 386–400, 1960.

[13] A. Vander Lugt. Coherent optical processing. *Proceedings of the IEEE*, 62:1300–1319, 1974.

[14] A.W. Lohmann and D.P. Paris. Computer generated spatial filters for coherent optical data processing. *Applied Optics*, 7:651–655, 1968.

[15] J.L. Horner. Light utilization in optical correlators. *Applied Optics*, 21:4511–4514, 1982.

[16] A.A.S. Awwal, M.A. Karim, and S.R. Jahan. Improved correlation discrimination using an amplitude-modulated phase-only filter. *Applied Optics*, 29:233–236, 1990.

[17] A.V. Oppenheim and J.S. Lim. The importance of phase in signals. *Proceedings of the IEEE*, 69:529–541, 1981.

[18] J.L. Horner and P.D. Gianino. Phase-only matched filtering. *Applied Optics*, 23:812–816, 1984.

[19] B.V.K.V. Kumar and Z. Bahri. Phase-only filters with improved signal-to-noise ratio. *Applied Optics*, 28:250–257, 1989.

[20] M.W. Farn and J.W. Goodman. Optimal binary phase-only matched filters. *Applied Optics*, 27:4431–4437, 1988.

[21] J.L. Horner and H.O. Bartelt. Two-bit correlation. *Applied Optics*, 24:2889–2893, 1985.

[22] R.D. Juday. Optimal realizable filters. *Journal of the Optical Society of America– A*, 1993. preprint.

[23] D.L. Flannery, J.S. Loomis, and M.E. Milkovich. Design elements of binary phase-only correlation filters. *Applied Optics*, 27:4231–4235, 1988.

[24] J.W. Goodman and A.M. Silvestri. Some effects of Fourier-domain phase quantization. *IBM Journal of Research and Development*, pages 478–484, 1970.

[25] J.L. Horner and J.R. Leger. Pattern recognition with binary phase-only filters. *Applied Optics*, 24:609–610, 1985.

[26] D. Psaltis, E.G. Paek, and S.S. Venkatesh. Optical image correlation with a binary spatial light modulator. *Optical Engineering*, 23:698–704, 1984.

[27] D.M. Cottrell, R.A. Lilly, J.A. Davis, and T. Day. Optical correlator performance of binary phase-only filters using Fourier and Hartley transforms. *Applied Optics*, 26:3755–3761, 1987.

[28] F.M. Dickey, J.J. Mason, and K.T. Stalker. Analysis of binarized Hartley phase-only-filter performance with respect to stochastic noise. *Optical Engineering*, 28:8–12, 1989.

[29] M.A. Flavin and J.L. Horner. Amplitude encoded phase-only filters. *Applied Optics*, 28:1692–1696, 1989.

[30] R. Cormack, K.M. Johnson, L. Zhang, and W.T. Cathey. Optical inspection of manufactured glass using adaptive Fourier filtering. *Optical Engineering*, 27:358–364, 1988.

[31] D.P. Casasent and D. Psaltis. Position, rotation, and scale invariant optical correlation. *Applied Optics*, 15:1795–1799, 1976.

[32] D. Casasent. Optical pattern recognition and artificial intelligence. *ICA-LEO Proceedings*, November 1986.

[33] W.J. Hossack, A.M. Darling, and A. Dadhouh. Coordinate transformations with multiple computer-generated optical elements. *Journal of Modern Optics*, 34:1235–1250, 1987.

[34] D. Casasent and D. Psaltis. Deformation invariant, space-variant optical pattern recognition. In E. Wolf, editor, *Progress in Optics XVI.* North Holland, 1978.

[35] C.F. Hester and D. Casasent. Multivariant technique of multiclass pattern recognition. *Applied Optics*, 19:1758–1761, 1980.

[36] R.R. Kallman. Construction of low noise optical correlation filters. *Applied Optics*, 25:1032–1033, 1986.

[37] A. Mahalanobis, B.V.K.V. Kumar, and D. Casasent. Minimum average correlation energy filters. *Applied Optics*, 26:3633–3640, 1987.

[38] G. Ravichandran and D.P. Casasent. Generalized in-plane rotation-invariant minimum average correlation energy filter. *Optical Engineering*, 30:1601–1607, 1991.

[39] J.L. Horner and P.D. Gianino. Applying the phase-only filter concept to the synthetic discriminant function correlation filter. *Applied Optics*, 24:851–854, 1985.

[40] D.A. Jared and D.J. Ennis. Inclusion of filter modulation in synthetic-discriminant-function construction. *Applied Optics*, 28:232–239, 1989.

[41] M.S. Kim and C.C. Guest. Simulated annealing algorithm for binary phase only filters in pattern classification. *Applied Optics*, 29:1203–1208, 1990.

[42] M.S. Kim and C.C. Guest. Experiments on annealed binary phase only filters fabricated with electron beam lithography. *Applied Optics*, 29:3380–3386, 1990.

[43] T.D. Wilkinson, D.C. O'Brien, and R.J. Mears. Scale invariant binary phase-only matched filter using a ferroelectric liquid crystal spatial light modulator. *Optics Letters*, 1993. Preprint, submitted February 1993.

[44] J.E. Rau. Detection of differences in real distributions. *Journal of the Optical Society of America*, 56:1490–1494, 1966.

[45] B. Javidi and S.F. Odeh. Multiple object identification by bipolar joint transform correlation. *Optical Engineering*, 27:295–300, 1988.

[46] D.A. Gregory and J.A. Loudin. Joint transform correlator limitations. *Proc. SPIE*, 1053:198–207, 1989.

[47] F.T.S. Yu, Q.W. Song, Y.S. Cheng, and D.A. Gregory. Comparison of detection efficiencies for VanderLugt and joint transform correlators. *Applied Optics*, 29:225–236, 1990.

[48] J.A. Davis, E.A. Merrill, and D.M. Cottrell. Effects of sampling and binarization in the output of the joint Fourier transform correlator. *Optical Engineering*, 29:1094–1100, 1990.

[49] D Casasent. Spatial light modulators. *Proc. IEEE*, 65:143–157, 1977.

[50] A.D. Fisher and J.N. Lee. The current state of two-dimensional spatial light modulator technology. *Proc. SPIE*, 634:352–371, 1986.

[51] K.M. Johnson, D.J. McKnight, and I Underwood. Smart spatial light modulators using liquid crystals on silicon. *IEEE Journal of Quantum Electronics*, 29:699–714, 1993.

[52] *OSA 1993 Technical Digest Series, Volume 6*, 1993.

[53] J. Grinberg, A. Jacobsen, W. Bleha, L. Miller, L. Fraas, D. Boswell, and G. Myer. A new real-time non-coherent to coherent light image converter–the hybrid field effect liquid crystal light valve. *Optical Engineering*, 14:217–225, 1975.

[54] U. Efron, J. Grinberg, P.O. Braatz, M.J. Little, and R.N. Schwartz. The silicon liquid crystal light valve. *Journal of Applied Physics*, 57:1356–1368, 1985.

[55] S.G. Latham and M.P. Owen. A silicon liquid-crystal spatial light modulator. *GEC Journal of Research*, 4:219–222, 1986.

[56] R.L. Hillman, Jr., G.A. Melnik, T.N. Tsakiris, F. Leard, R. Jurgilewicz, and C. Warde. Electron-beam-addressed lithium niobate spatial light modulators. *Proc.SPIE*, 1562, 1991.

[57] W.E. Ross, D. Psaltis, and R.H. Anderson. Two-dimensional magneto-optic spatial light modulator for signal processing. *Optical Engineering*, 22:485–490, 1985.

[58] N.H. Harhat and Z.Y. Shae. Scheme for enhancing the frame rate of magnetooptic spatial light modulators. *Applied Optics*, 28:4792–4800, 1989.

[59] W.E. Ross and D.N. Lambeth. Fabrication and switching sensitivity of the new magneto-optic spatial light modulator. *OSA 1993 Technical Digest Series*, 6:32–35, 1993.

[60] L.J. Hornbeck. Deformable mirror spatial light modulators. *SPIE Critical Reviews Series*, 1150:86–102.

[61] M.A.Z. Rejman-Greene, E.G. Scott, P. Skevington, and R.P. Webb. Two dimensional arrays of InGaAs/InP MQW modulators and InGaAs photodiodes. In *Colloquium on Two-dimensional optoelectronics device arrays, Digest no. 1191/158*, October 1991.

[62] J.H. Marsh and K. McIlvaney. Issues in the design of MQW arrays. In *Colloquium on Two-dimensional optoelectronics device arrays, Digest no. 1191/158*, October 1991.

[63] M.E. Prise, N.C. Craft, M.M. Downs, R.E. LaMarche, L.A. D'Asaro, L.M.F. Chirovsky, and M.J. Murdocca. Optical digital processor using arrays of symmetric self-electrooptic devices. *Applied Optics*, 30:2287–2296, 1991.

[64] J.A. McEwan, A.D. Fisher, P.B. Rolsma, and J.N. Lee. Optical processing characteristics of a low-cost liquid crystal display device. *Journal of the Optical Society of America, A*, 2:P8, 1985.

[65] M. Young. Low-cost LCD video display for optical processing. *Applied Optics*, 25:1024–1026, 1986.

[66] E.C. Tam, S. Wu, A. Tanone, F.T. Yu, and D.A. Gregory. Closed loop binary phase correction of an LCTV using a point diffraction interferometer. *IEEE Photonics Technology Letters*, 2:143–146, 1990.

[67] I Underwood. *An nMOS addressed liquid crystal spatial light modulator.* PhD thesis, University of Edinburgh, 1987.

[68] M.J. Ranshaw. *Phase-modulating spatial light modulators.* PhD thesis, University of Edinburgh, 1988.

[69] D.J. McKnight, D.G. Vass, and R.M. Sillitto. Development of a spatial light modulator: a randomly addressed liquid crystal over nMOS array. *Applied Optics*, 28:4757–4762, 1989.

[70] I. Underwood, D.G. Vass, R.M. Sillitto, G. Bradford, N.E. Fancey, A.O. Al Chalabi, M.J. Birch, W.A. Crossland, A.P. Sparks, and S.G. Latham. A high performance spatial light modulator. *Proc. SPIE*, 1562:107–115, 1991.

[71] D. Vass. Electrically programmed optical logic arrays. UKDTI Patent 8910168.7.

[72] P.W. McOwan, J. Gourlay, M.W.G. Snook, D.G. Vass, and W.J. Hossack. SASLM smart optical array processing. Technical report, University of Edinburgh, 1992.

[73] T. Slagel and K. Wagner. Competitive optical learning with winner-take-all optical modulators. *OSA 1991 Technical Digest Series*, 6:280–283, 1991.

[74] A.R. MacGregor. *Modelling the electro-optic properties of liquid crystals.* PhD thesis, University of Edinburgh, 1989.

[75] L.M. Blinov. *Electro-optical and magneto-optical properties of liquid crystals.* Wiley-Interscience, 1983.

[76] S. Chandrasekhar. *Liquid crystals.* Cambridge University Press, Cambridge, 1992.

[77] M.G. Clark, K.J. Harrison, and E.P. Raynes. Liquid crystal materials and devices. *Physics Technology*, 11:232–240, 1980.

[78] K. Skarp and M.A. Handschy. Ferroelectric liquid crystals. material properties and applications. *Mol. Cryst. Liq. Cryst.*, 165:439–509, 1988.

[79] J. Cognard. Alignment of nematic liquid crystals and their mixtures. *Mol. Cryst. Liq. Cryst. Supplement Series*, 1, 1982.

[80] BDH Limited, Advanced Materials Division, Broom Road, Poole, BH12 4NN. *Electro-optic Liquid Crystals.* datasheets.

[81] C.H. Gooch and H.A. Tarry. The optical properties of twisted nematic liquid crystal structures with twist angles $\leq 90°$. *Journal of Physics, D*, 8:1575–1584, 1975.

[82] J. Grinberg and A.D. Jacobsen. Transmission characteristics of a twisted nematic liquid-crystal layer. *Journal of the Optical Society of America*, 66:1003–1009, 1976.

[83] P.M. Alt and P. Pleshko. Scanning limitations of liquid crystal displays. *IEEE Transactions on Electron Devices*, ED-21:146–155, 1974.

[84] B. Needham. Practical limits on addressing twisted nematic displays. *Phil. Trans. R. Soc. Lond. A*, 309:179–188, 1983.

[85] E. Hecht and A. Zajac. *Optics.* Addison Wesley, Reading, Mass., 1974.

[86] L. Brillouin. *Science and Information Theory.* Academic Press, New York, 1956.

[87] D.C. Champeney. *Fourier Transforms and their Physical Applications.* Academic Press, London, 1973.

[88] John Bradley. *xv- an interactive image display for the X Window System.* bradley@cis.upenn.edu. Public domain software, version 2.21.

[89] H.F. Yau, Y.P. Tong, C.M. Lee, C.C. Chang, and M.W. Chang. Pattern recognition with circular harmonic filtering and phase-only circular harmonic filtering using a vander lugt filter. *Japanese Journal of Applied Physics*, 31:49–53, 1992.

[90] G. Gheen, E. Washwell, and D. Armitage. The effect of filter pixelation on optical correlation. *OSA 1990 Technical Digest Series*, 14:161–165, 1990.

[91] H. Stark. Diffraction patterns of nonoverlapping circular grains. *Journal of the Optical Society of America*, 67:700–703, 1977.

[92] H. Stark. Some spectral properties of random arrays of grains: Comment. *Journal of the Optical Society of America*, 68:1785–1786, 1978.

[93] C. Aime. A morphological interpretation of the spatial power spectrum of the solar granulation. *Astronomy and Astrophysics*, 67:1–6, 1978.

[94] F. Martin and C. Aime. Some spectral properties of random arrays of grains. *Journal of the Optical Society of America*, 68:1782–1785, 1978.

[95] *Uniras+ 2000, scientific data visualisation software.* 376 Gladsaxevej, DK2860, Søborg, Denmark, 1990.

[96] J.T.M. Stevenson and A.M. Gundlach. The application of photolithography to the fabrication of microcircuits. *J. Phys. E: Instrum*, 19:654–667, 1986.

[97] Shipley Europe Limited, Herald Way, Coventry. *Microposit 1800 series photo-resist.* datasheets, 17/3/86.

[98] Shipley Europe Limited, Herald Way, Coventry. *Microposit 351 developer.* datasheets, 2/12/86.

[99] E. Davidson. A laboratory method of producing indium tin oxide layers on a glass substrate. Technical report, University of Edinburgh, 1991.

[100] S.M. Sze. *VLSI technology.* McGraw Hill International Student Editions, Singapore, 1983.

[101] BDH Limited, Advanced Materials Division, Broome Road, Poole, BH12 4NN. *Twisted nematic display.* Construction details.

[102] Hitachi Europe Limited, 21 Upton Road, Watford, Herts. *Hitachi LCD controller/driver LSI data book.*

[103] I.R. Redmond and D.J. McKnight. Spider LCD driver software for BBC microcomputer. Private communication, June 1991.

[104] W.H. Steel. *Interferometry.* Cambridge University Press, Cambridge, second edition, 1986.

[105] R. Guenther. *Modern Optics.* John Wiley & Sons, New York, 1990.

[106] D.J. Potter. *Phase-only optical information processing.* PhD thesis, University of Edinburgh, 1992.

[107] A. O'Hara, J.R. Hannah, I. Underwood, D.G. Vass, and R.J. Holwill. Mirror quality and efficiency improvements of reflective spatial light modulators using dielectric coatings and chemical-mechanical polishing. *Applied Optics.* Accepted June 1993.

[108] N. Collings, W.A. Crossland, P.J. Ayliffe, D.G. Vass, and I. Underwood. Evolutionary development of advanced liquid crystal spatial light modulators. *Applied Optics*, 28:4740–4747, 1989.

[109] G. Gheen, E. Washwell, and C. Huang. Problems facing optical correlators. *Proceedings of the SPIE*, 1772:96–103, 1992.

Appendix A

Linearity of array camera

A brief experiment was carried out to determine the linearity of the array camera used throughout the thesis. The camera under test was an Avcam 405, number 185, a 512×512 pixel CCD array camera. The experimental apparatus is shown in figure A.1. The illumination was collimated, and attenuated by means of the rotatable variable density attenuator to a level which almost saturated the camera when the polarisers were aligned. The maximum intensity in this state was recorded. The second polariser was then rotated so that its axis was at an angle θ relative to the fixed polariser, and the angle and corresponding maximum intensity recorded. This was repeated at 5^0 intervals until $\theta = 90^0$. By Malus' Law the intensity transmitted by the polarisers is proportional to $(\cos\theta)^2$, so a graph of $(\cos\theta)^2$ versus measured intensity should be linear and pass through the origin

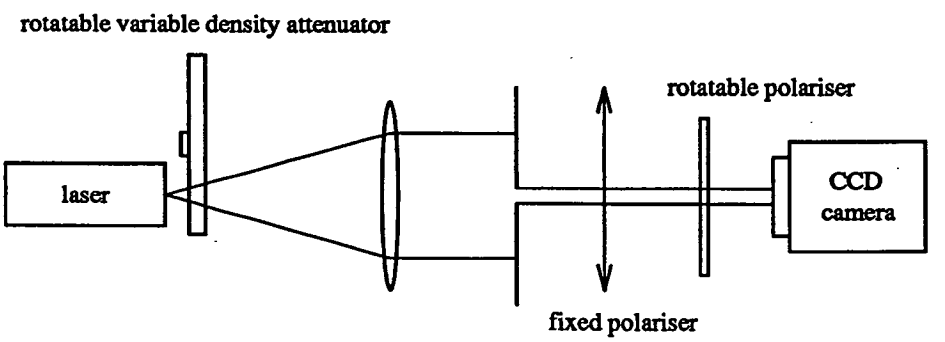

Figure A.1: Experimental apparatus to determine linearity of array camera

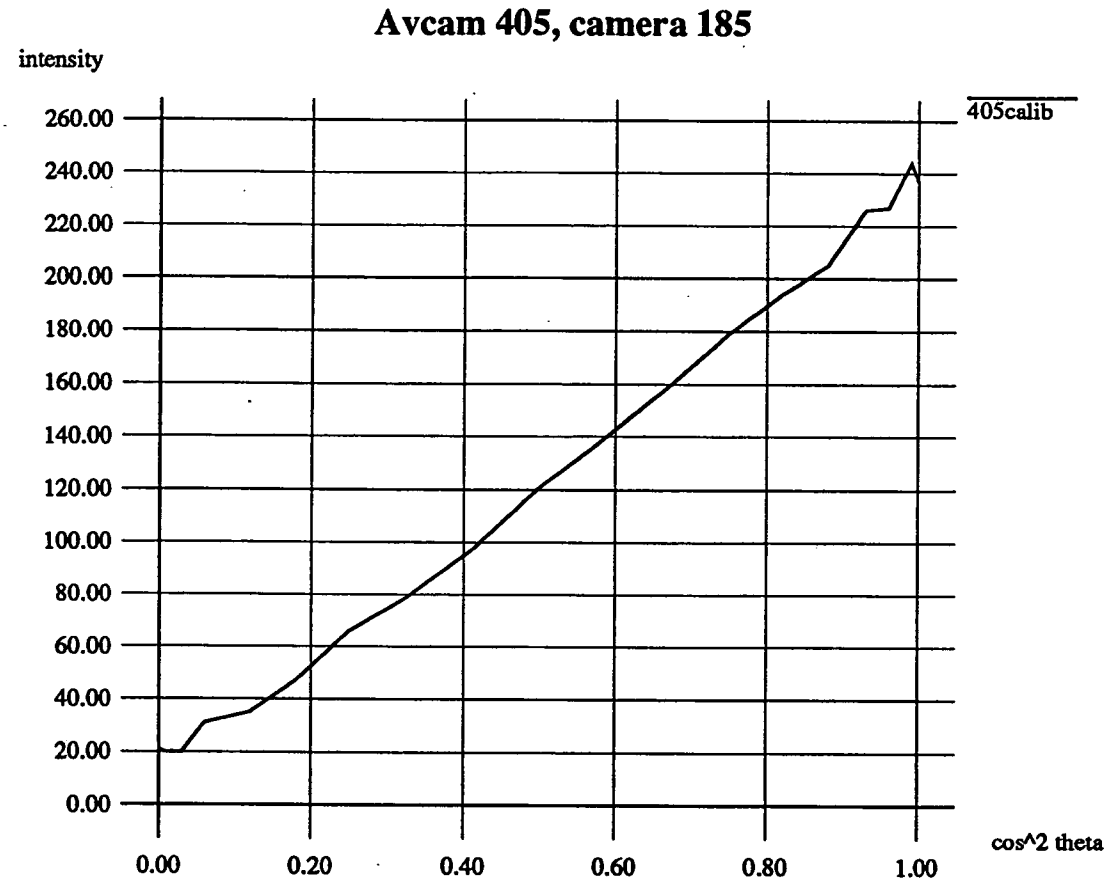

Figure A.2: Calibration curve for Avcam 405 CCD camera, number 185

if the camera is perfectly linear.

The results are shown in figure A.2, and show that the camera is indeed extremely linear over most of its range (from intensity values of 30 to 210). An offset of 3 intensity units results when a straight line is fitted to the data— this is accounted for in the results quoted throughout the thesis.

Appendix B

Listing of pixel.f

```
c23456
      program pixel
c
c     Program to generate 16x16 pixel arrays, at various scales
c     with various randomisation schemes.
c     Steven Heddle.
c
      character*9 outfile
      real diff,scale
      integer pflag,offflag,arrayt
      integer spxlx(0:16,0:16),spxly(0:16,0:16),patt(0:16,0:16)
      real ampl(0:300,0:300)
      real ampl2(256,256)
      print*,'Enter name of output file'
      read(*,300)outfile
300   format(a9)

      print*,'Input pixel size,sub pixel size in mm.'
      read*,psize,spsize

      print*,'Is array regular (1) or randomised (2) '
      print*,'                 or the SPECIAL CASE!!! (3)'
```

```
      read*,arrayt
      print*,'Is offset of array required (1) or not (2)'
      read*,offflag
      print*,'Input background transmission,pixel transmission (0-16)'
      read*,bg,on

 10   print*,'Array scaled smaller (max 16x)? Enter scale factor'
      read*,scale
      if((scale.lt.1).or.(scale.gt.16))goto 10
      spsize=(16.0*spsize/psize)/scale
      psize=16.0/scale
      ipsize=anint(psize)
      ispsize=anint(spsize)
      diff=psize-spsize

      print*,'Is pattern superimposed on array (1) or not (2)'
      read*,pflag
      call pattern(pflag,patt,diff)
c
c     Superimposes pattern on the 16x16 output if that is necessary
c
      if(arrayt.eq.1)call regular(spxlx,spxly,psize)
      if(arrayt.eq.2)call random(diff,spxlx,spxly,psize)
      if(arrayt.eq.3)call special(spxlx,spxly,psize)
c
c     Sets top corner position of the pixels
c
      if(offflag.eq.1)then
      call offset(spsize,spxlx,spxly,xoff,yoff)
      else
      continue
      endif
c
c     prompts for offset, then establishes what it should be
c     and then implements it s.t. the zero order in the fourier plane
c     gets through.
```

```
c
      call writeout(outfile,patt,spxlx,spxly,ispsize,bg,on,xoff
     *,yoff,scale)
c
c     ...and thats it.
c
      stop
      end

c********SUBROUTINES START HERE**********************

      subroutine pattern(pflag,patt,diff)
      integer pflag
      integer patt(0:16,0:16)
      character*32 patfile
      if(pflag.eq.1)then
        print*,'Enter name of 16x16 pattern array file'
        read*,patfile
        open(unit=1,file=patfile,status='old')
        do 23,j=0,15
        read(1,13)(patt(i,j),i=0,15)
 23     continue
        close(1)
 13   format(16i1)
      else
      do 33,j=0,15
       do 34,i=0,15
        patt(i,j)=1
 34    continue
 33   continue
      endif
      return
      end

      subroutine regular(spxlx,spxly,psize)
      integer spxlx(0:16,0:16),spxly(0:16,0:16)
```

```
      open(unit=1,file='extreg',status='new')
        do 10,j=0,15
          do 20,i=0,15
          spxlx(i,j)=anint(psize*i)
          spxly(i,j)=anint(psize*j)
          write(1,*)spxlx(i,j)
          write(1,*)spxly(i,j)
 20     continue
 10   continue
      close(1)
      return
      end

      subroutine random(diff,spxlx,spxly,psize)
      integer spxlx(0:16,0:16),spxly(0:16,0:16)
      integer plag,dpos
      double precision x,y
      real diff,offx,offy
      external G05CBF
      double precision G05CAF
      external G05CAF
      print*,'Uniform distbn. of pixels (1) or discrete positions (2)?'
      read*,pflag
      if(pflag.ne.1)then
       print*,'How many discrete positions then?'
       read*,dpos
       if(diff.gt.(psize/2).and.dpos.eq.2)diff=psize/2
      else
       continue
      endif
      open(unit=1,file='extrads',status='new')
        call G05CBF(0)
        do 10,j=0,15
          do 20,i=0,15
          x=G05CAF(x)
          y=G05CAF(y)
```

```
        if(pflag.ne.1)then
         x=int(x*dpos)
         y=int(y*dpos)
         offx=x*diff/(dpos-1)
         offy=y*diff/(dpos-1)
        else
         offx=x*diff
         offy=y*diff
        endif
        spxlx(i,j)=anint(psize*i+offx)
        spxly(i,j)=anint(psize*j+offy)
        write(1,*)spxlx(i,j)
        write(1,*)spxly(i,j)
20     continue
10   continue
     close(1)
     return
     end

     subroutine writeout(outfile,patt,spxlx,spxly,ispsize,bg,on,xoff
    *,yoff,scale)
     real ampl(0:300,0:300)
     real ampl2(256,256)
     real scale
     integer xoff,yoff
     integer spxlx(0:16,0:16),spxly(0:16,0:16),patt(0:16,0:16)
     character*9 outfile,pattern
       do 30,k=0,255
         do 30,l=0,255
         ampl(k,l)=0.0
30   continue
     islmx=xoff
     lslmx=xoff+anint(256.0/scale)-1
     islmy=yoff
     lslmy=yoff+anint(256.0/scale)-1
       do 60, i=islmx,lslmx
```

```
          do 60, j=islmy,lslmy
          ampl(i,j)=bg
 60     continue
      open(unit=8,file=outfile,status='new',form='unformatted')
        do 40,j=0,15
          do 40,i=0,15
            if((patt(i,j)).eq.1)then
            do 41,m=0,ispsize-1
              do 41,n=0,ispsize-1
              k=spxlx(i,j)+m
              l=spxly(i,j)+n
              ampl(k,l)=on
 41         continue
          else
          continue
          endif
 40     continue
         do 50,j=1,256
          do 50,i=1,256
          ampl2(i,j)=ampl(i-1,j-1)
 50   continue
      write(8)ampl2
      close(8)
      return
      end

      subroutine offset(spsize,spxlx,spxly,xoff,yoff)
      integer spxlx(0:16,0:16), spxly(0:16,0:16)
      integer xoff,yoff
      xoff=int(129-spxlx(7,7)-spsize/2)
      yoff=int(129-spxly(7,7)-spsize/2)
      print 100,xoff,yoff
100   format('Horizontal offset is ',i4,' Vertical offset is ',i4)
      do 20,j=0,15
        do 20,i=0,15
        spxlx(i,j)=spxlx(i,j)+xoff
```

```
        spxly(i,j)=spxly(i,j)+yoff
 20   continue
      return
      end

      subroutine special(spxlx,spxly,psize)
      integer parray(-1:15,-1:15),spxlx(0:16,0:16),spxly(0:16,0:16)
      integer pflag,p0,p1,p2,p3,pout
      double precision x
      double precision G05CAF
      external G05CAF
      external G05CBF
      call G05CBF(0)
      do 20,j=-1,15
       do 20,i=-1,15
       x=G05CAF(x)
       parray(i,j)=int(3.99*x)
 20   continue
      open(unit=1,file='extspec',status='new')
      do 10,j=0,15
       do 10,i=0,15
       a=parray(i,j-1)
       b=parray(i-1,j)
 30    x=G05CAF(x)
        if(((a.eq.0).or.(a.eq.1)).and.((b.eq.0).or.(b.eq.3)))then
          call choose(1,1,1,1,pout,x,pflag)
        elseif((a.eq.2).and.((b.eq.0).or.(b.eq.3)))then
          call choose(0,1,1,1,pout,x,pflag)
        elseif((b.eq.2).and.(a.ne.3))then
          call choose(0,1,1,1,pout,x,pflag)
        elseif((a.eq.3).and.((b.eq.0).or.(b.eq.3)))then
          call choose(1,0,1,1,pout,x,pflag)
        elseif(((a.eq.0).or.(a.eq.1)).and.(b.eq.1))then
          call choose(1,1,1,0,pout,x,pflag)
        elseif((a.eq.2).and.(b.eq.1))then
          call choose(0,1,1,0,pout,x,pflag)
```

```
        elseif((a.eq.3).and.(b.eq.1))then
          call choose(1,0,1,0,pout,x,pflag)
        elseif((a.eq.3).and.(b.eq.2))then
          call choose(0,0,1,1,pout,x,pflag)
        else
        continue
       endif
       if(pflag.ne.1)goto 30
       parray(i,j)=pout
        if((pout.eq.0).or.(pout.eq.1))offy=0
        if((pout.eq.2).or.(pout.eq.3))offy=psize/2
        if((pout.eq.0).or.(pout.eq.3))offx=0
        if((pout.eq.1).or.(pout.eq.2))offx=psize/2
       spxlx(i,j)=anint(psize*i+offx)
       spxly(i,j)=anint(psize*j+offy)
       write(1,*)spxlx(i,j)
       write(1,*)spxly(i,j)
  10   continue
       close(1)
       return
       end

       subroutine choose(p0,p1,p2,p3,pout,x,pflag)
c
c      This routine sticks a subpixel into one of four boxes,
c      such that the pixels can be joined up in an actual
c      matrix addressed device. Some weighting of the selections
c      is performed due to the nature of the ensemble.
c
       integer p0,p1,p2,p3,pout,pflag
       real xx
       double precision x
       xx=x*4
        if(xx.le.1.33)then
        pout=0
        pflag=p0
```

```
      elseif(xx.le.2.33)then
      pout=1
      pflag=p1
      elseif(xx.le.3)then
      pout=2
      pflag=p2
      elseif(xx.le.4)then
      pout=3
      pflag=p3
      else
      continue
     endif
     return
     end
```

Appendix C

Publications

Some of the work described in this thesis has been submitted to the Journal of Modern Optics and published as

S.Heddle and R.M.Sillitto, **Reduction and removal of replication in an optical processor by randomisation of pixel positions in the Fourier plane filter,** *Journal of Modern Optics,* **40***(2), 299–313, 1993.*

A paper was presented to the SPIE's 1992 International Symposium on Optical Applied Science and Engineering: Optical Information Systems and Architectures IV, July 20-21 1992, San Diego. It is published as

S.Heddle, D.G.Vass and R.M.Sillitto, **Reduction of aliasing in correlation using a pixelated Spatial Light Modulator,** *Proceedings of the SPIE, volume 1772, 116–127, 1992.*

Both these papers are included here.

The specific task of correlating for road signs from a moving vehicle is addressed in a paper scheduled for presentation at the OSA Annual Meeting in Toronto, 3–8th October 1993, and includes some of the ideas of this thesis. A draft version of this paper is also included here. The abstract will be published in an issue of

the Journal of the Optical Society of America as

S.Heddle, W.J.Hossack and D.G.Vass, **Specification of an optical correlator for road sign recognition.**

JOURNAL OF MODERN OPTICS, 1993, VOL. 40, NO. 2, 299–313

Reduction and removal of replication in an optical processor by randomization of pixel positions in the Fourier plane filter

S. HEDDLE and R. M. SILLITTO

Department of Physics, University of Edinburgh, The Kings Buildings, Mayfield Road, Edinburgh EH9 3JZ, Scotland

(*Received 29 July 1992 and accepted 14 October 1992*)

Abstract. When a pixellated spatial light modulator (SLM) is used in an optical processor, replicated images occur in the output due to the underlying regular pixel structure of the SLM. These replicas arise through convolution with distinct spectral orders in the SLM's Fourier spectrum and in the case of a correlator can lead to false correlation signals. A means of suppressing or eliminating the spectral orders through randomization of the pixel positions is presented, with analysis of the Fourier spectrum of random arrays. Two specific randomization schemes are considered and the results are compared with that for a regular array.

1. Introduction

The use of spatial light modulators (SLMs) in optical processors has become widespread as a flexible means of general data input to the processor. Additionally, much interest has been shown in their use as Fourier plane filters in optical correlators [1–3].

When the pattern on the SLM is dependent on the output of some digitally computed algorithm [1–3] electrical addressing is necessary to enable the output to be conveniently written to the SLM. Suitable SLMs such as the MOSLM [4], active silicon backplane liquid crystal devices [5–7], or transmissive matrix addressed liquid crystal devices [8], in addition to being electrically addressable are pixellated. The regular pixel structure produces a series of replicas of the correlator output, whose intensities and spatial separation are dependent on the number of pixels and the pixel dimensions.

In this paper it is shown how these replicas arise and how they can lead to false correlation results. A means of greatly suppressing these replicas by randomizing the pixel positions is introduced. It is shown that the spectral orders in the diffraction pattern of the SLM, which lead to replication, can be eliminated by a suitable choice of pixel-position probability distribution function. Computational and optical simulations are presented, and followed by a discussion of the results.

2. Formation of replicas

In the standard optical processor configuration shown in figure 1 an input object of transmittance $t_1(x, y)$ is uniformly illuminated by collimated coherent light of wavelength λ, and the Fourier transform of the object is formed by lens L_2 in its back focal plane, plane P_2. Assuming illumination of unit amplitude the light field

0950–0340/93 $10·00

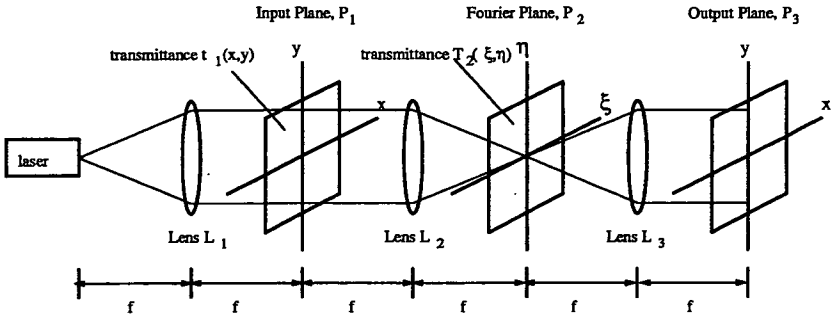

Figure 1. Optical processor, and coordinate system.

immediately to the left of P_2 can be written as $u_2^-(\xi,\eta)=T_1(\xi,\eta)$ where $T_1(\xi,\eta)$ is the Fourier transform of $t_1(x,y)$. The (ξ,η) are scaled spatial frequency coordinates with

$$\xi=u/\lambda f, \qquad \eta=v/\lambda f, \tag{1}$$

where u and v are spatial frequencies and f is the focal length of lens L_2. Denoting the transmittance of the SLM in plane 2 by $T_2(\xi,\eta)$, the light field immediately to the right of plane P_2 will be given by $u_2^+(\xi,\eta)=T_1(\xi,\eta)T_2(\xi,\eta)$. Lens L_3 then performs a second forward Fourier transform to give the light field in plane 3 as $u_3(x,y)=t_1'(x,y)*t_2(x,y)$, where $*$ denotes convolution, $t_1'(x,y)$ and $t_2(x,y)$ are the Fourier transforms of $T_1(\xi,\eta)$ and $T_2(\xi,\eta)$ respectively, and $t_1'=t_1(-x,-y)$, noting that two successive forward Fourier transforms produce a reversal of coordinates. Thus $u_3(x,y)=t_1(-x,-y)*t_2(x,y)$, demonstrating that the output is a coordinate reversed version of the input convolved with the Fourier transform of the filter function $T_2(\xi,\eta)$.

For a pixellated SLM with all pixels 'on' we can write the filter function as

$$T_2(\xi,\eta)=p(\xi,\eta)*\sum_{n=0}^{Q-1}\sum_{m=0}^{Q-1}\delta(\xi-n\alpha,\eta-m\alpha), \tag{2}$$

which denotes a $Q\times Q$ array of pixels each described by the pixel function $p(\xi,\eta)$ and with their centres placed on a regular square grid of points with pitch α. The filter function has been defined asymmetrically with respect to the origin purely for convenience—otherwise two definitions would be necessary to ensure that a pixel was centred on the origin: one symmetric, corresponding to Q odd; the other slightly displaced from a symmetric position, corresponding to Q even.

Fourier transforming gives

$$t_2(x,y)=\mathrm{P}(x,y)\sum_{n=0}^{Q-1}\exp(-2\pi in\alpha x)\sum_{m=0}^{Q-1}\exp(-2\pi im\alpha y) \tag{3}$$

$$=P(x,y)\frac{\sin(\pi Q\alpha x)}{\sin(\pi\alpha x)}\frac{\sin(\pi Q\alpha y)}{\sin(\pi\alpha y)}, \tag{4}$$

ignoring a phase factor resulting from the asymmetric definition of the filter with respect to the origin.

The diffraction pattern of the SLM can be shown more explicitly to be a 2-dimensional (2D) array of δ functions, broadened by convolution with narrow sinc(..) functions, by noting that [9]

$$\sum_{n=0}^{Q-1} \delta(\xi - n\alpha) = \sum_{n=-\infty}^{\infty} \delta(\xi - n\alpha) \operatorname{rect}\{[\xi - (Q-1)\alpha/2]/Q\alpha\}, \tag{5}$$

whence

$$t_2(x,y) = P(x,y)\left\{Q^2 \operatorname{sinc}(Q\alpha x, Q\alpha y) * \sum_{n=-\infty}^{\infty}\sum_{m=-\infty}^{\infty} \delta\left(x - \frac{n}{\alpha}, y - \frac{m}{\gamma}\right) \times \exp[-i\pi(Q-1)\alpha(x+y)]\right\}, \tag{6}$$

where $\exp[-i\pi(Q-1)\alpha(x+y)]$ is the phase factor ignored earlier.

Equation (6) shows that the replicas will be formed at the positions of the spectral orders defined by the 'broadened δ functions', separated by $\lambda f/\alpha$. These broadened δ functions define spectral orders of spatial frequency whose magnitudes are governed by the form of $P(x,y)$. In particular, if the pixels are square and of side β with amplitude transmittance 0 outside the pixel and 1 inside the pixel,

$$P(x,y) = \beta^2 \operatorname{sinc}(\beta x) \operatorname{sinc}(\beta y). \tag{7}$$

Figure 2 (*a*) shows a regular 16 x 16 array of square pixels of side $\beta = \alpha/2$. Figure 2 (*b*) shows $|t_2(x,y)|^2$, the squared modulus of its Fourier transform. This quantity, the power spectrum, is of course what is detected when the optical Fourier transform is observed. Figure 3 is a cross-section of the power spectrum along the x axis, and shows how the relative intensities of the spectral orders depend on $P(x,y)$—in particular note the absence of the second spectral orders.

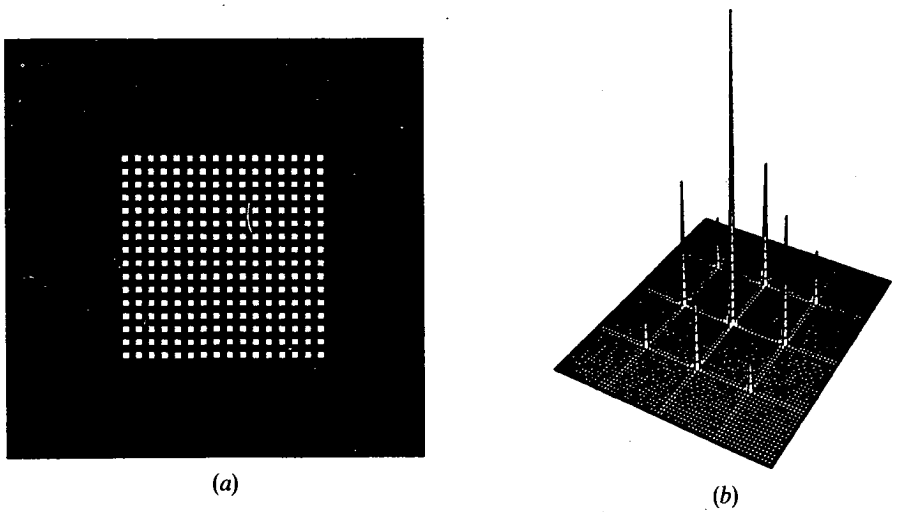

(*a*) (*b*)

Figure 2. (*a*) Regularly spaced 16 x 16 array, with $\beta = \alpha/2$. (*b*) Power spectrum, $|t_2(x,y)|^2$, of (*a*).

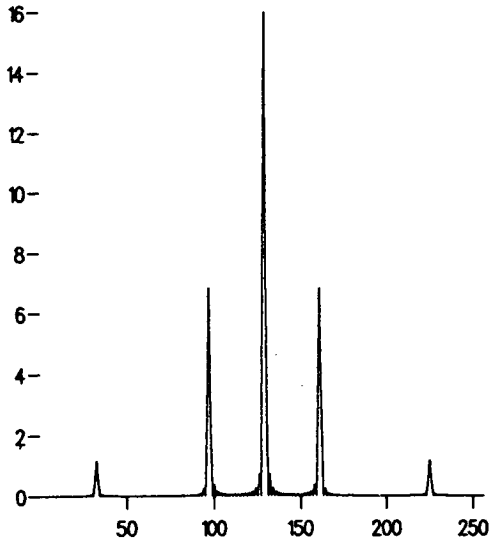

Figure 3. Power spectrum along the x axis, $|t_2(x, 0)|^2$, for regularly spaced array, with $\beta = \alpha/2$.

3. Consequences of replication

We now consider the consequencies of the replication on optical correlation experiments, where the objective is to detect the presence of a target object in an input scene. Whether or not the replicas constitute a problem will depend on the spectral content of the input scene, the nature of the target contained within, and the resolution of the filter SLM. In many applications the input to the processor will itself be implemented using an SLM. This is assumed to be the case here, with the input SLM labelled as SLM1 and the filter SLM as SLM2. It is convenient to characterize each SLM by its space-bandwidth product (SBP), equal to the number of pixels across the SLM, Q.

To prevent the replicas overlapping (i.e. aliasing), the Fourier transform must be sampled at frequency greater than or equal to the Nyquist frequency. This condition imposes a maximum size on the object—if the pixel pitch of SLM2 is α, the maximum size of the object will be $\lambda f/\alpha$. If this object then occupies $Q_1 \times Q_1$ pixels on SLM1 then the lowest spatial period that can be represented is $(2/Q_1)(\lambda f/\alpha)$. This in turn will transform to a point $Q_1\alpha/2$ distant from the position of the zero spatial frequency component of the transform, which passes through the centre of the central pixel of SLM2, requiring that the SBP of the SLM2 be at least as large as that of the target in order to represent accurately its transform or complex conjugate. In the case where the SBPs of SLM1 and SLM2 are equal (necessary if the target occupies all of SLM1), then the replication of the output will not be a problem as the replicas will fall outwith the region of interest in plane 3, which will be restricted to the area defined by the geometric image of the perimeter of SLM1. If however, the target or targets each occupy a 32×32 pixel region in a 256×256 input SLM for example, then a 32×32 pixel SLM in plane 2, though adequate to represent the target's Fourier transform, will produce replicas lying within the region of interest in the output plane [10, 11] which could give rise to peaks which may be misinterpreted as signifying correlation. Increasing the SBP of SLM2 so that it was a 256×256 pixel

device would solve the problem, but the advantages to be gained from using a low-resolution SLM in the Fourier plane (fast frame rates, compact size, modest memory size for the filter bank [11]) suggest that a worthwhile alternative approach to the problem is to try to eliminate or at least reduce the magnitudes of the spectral orders in $t_2(x, y)$ which lead to the replicas.

It is apparent that $P(x, y) = \beta^2 \operatorname{sinc}(\beta x) \operatorname{sinc}(\beta y)$ has zeros at

$$x = \frac{N}{\beta} \quad \text{or} \quad y = \frac{M}{\beta}, \quad N, M = \pm 1, \pm 2, \pm 3 \ldots . \tag{8}$$

So we see that if the size of the SLM pixels, β, was equal to the pixel pitch, α, all the spectral orders except the zero order could be eliminated, so removing all the replicas in plane 3—but only if all pixels were 'on'. This implies a fill factor of 100% [12], which is not realizable for electrically addressed devices as short-circuits would occur between adjacent pixels. Fill factors of available SLMs are quite likely to be as little as 54% [4] or 25% [5]. Nevertheless, the larger the fill factor the smaller the magnitude of the first spectral orders will be, and eliminating any given spectral order higher than the first is practicable by choice of some $\beta/\alpha < 1$.

As it is not possible to eliminate the first orders by the choice of β/α, some other means must be found: randomizing the pixel positions is one interesting possibility, which we shall now discuss.

4. The diffraction pattern of a random array

In this section β/α is chosen to be 0·5, thus placing a zero of the multiplying transformed pixel function at the centre of all the even orders, excepting of course the zero order. This is illustrated for the regular array in figure 3 which shows the cross-section across figure 2 (*b*).

We shall show that it is possible to eliminate the distinct odd spectral orders by displacing the pixels from their regularly spaced positions, by some displacement chosen randomly from a specified set of displacements.

The diffraction pattern of an array of square pixels displaced from positions on a regular grid must first be analysed. This problem has a long history of rediscovery and rederivation. It was considered by Debye who reported on it in a series of papers in 1913 [13] which considered the diffraction of X-rays by a simple cubic crystal whose atoms were displaced from their equilibrium positions by thermal agitation. More recent work on this general problem set it in the contexts of the diffraction of light by patterns of non-overlapping grains [14, 15] and the power spectra of solar granulation [16, 17], from which the following analysis is partially derived.

To simplify the analysis, we proceed first by postulating a filter function in only one dimension (1D):

$$T_2(\xi) = p(\xi) * \sum_{n=0}^{Q-1} \delta(\xi - n\alpha - \phi_n), \tag{9}$$

with $T_2(\xi)$ the transmission function, $p(\xi)$ the pixel function, and ϕ_n a random displacement of the pixel centre from the regular grid position.

Fourier transforming and using the convolution and shift theorems gives

$$t_2(x) = P(x) \sum_{n=0}^{Q-1} \exp(-i2\pi n\alpha x) \exp(-i2\pi \phi_n x). \tag{10}$$

To handle this equation analytically requires knowledge of all the ϕ_n values, and clearly will not result in a useful general expression. However, noting that the equation represents just one of an ensemble of possible pixel arrays, dependent on the choice of the ϕ_n, we can use an ensemble average to replace the part of the equation involving the random values, and this will yield an expression showing the form of the diffraction pattern. Since we detect the intensity the ensemble averaging is performed over $|t_2(x)|^2$.

Thus we wish to evaluate

$$\langle |t_2(x)|^2 \rangle = \left\langle |P(x)|^2 \sum_{n=0}^{Q-1} \sum_{n'=0}^{Q-1} \exp(-i2\pi n\alpha x) \exp(i2\pi n'\alpha x) \exp(-i2\pi\phi_n x) \exp(i2\pi\phi_{n'} x) \right\rangle, \tag{11}$$

where the angle brackets denote the ensemble averaging operation.

By removing the ϕ_n, $\phi_{n'}$ terms which are dependent on each other (i.e. when $n=n'$) from the summation,

$$\langle |t_2|^2 \rangle = |P(x)|^2 \times \left[Q + \sum_{\substack{n=0 \\ n\neq n'}}^{Q-1} \sum_{n'=0}^{Q-1} \times \exp(-i2\pi n\alpha x) \exp(i2\pi n'\alpha x) \langle \exp(-i2\pi\phi_n x) \rangle \langle \exp(i2\pi\phi_{n'} x) \rangle \right]. \tag{12}$$

This can be readily evaluated in general, as

$$\langle \exp(-i2\pi\phi_n x) \rangle = \int_{-\infty}^{\infty} P_{\phi_n}(\phi) \exp(-i2\pi\phi x)\, d\phi, \tag{13}$$

where $P_{\phi_n}(\phi)$ is the real, positive probability of finding ϕ_n with value ϕ, and is identical to $P_{\phi_{n'}}(\phi)$. Using this, and taking $P_{\phi_n}(\phi)$ to be even, allows us to write

$$\langle \exp(i2\pi\phi_{n'} x) \rangle = \int_{-\infty}^{\infty} P_{\phi_n}(\phi) \exp(i2\pi\phi x)\, d\phi \tag{14}$$

$$= \left[\int_{-\infty}^{\infty} P_{\phi_n}(\phi) \exp(-i2\pi\phi x)\, d\phi \right]^*, \tag{15}$$

giving

$$\langle \exp(-i2\pi\phi_n x) \rangle \langle \exp(i2\pi\phi_{n'} x) \rangle = |\rho(x)|^2, \tag{16}$$

where $\rho(x)$ is the Fourier transform of the probability distribution function, $P_{\phi_n}(\phi)$.

Substituting using

$$\underbrace{\sum_{n=0}^{Q-1} \exp(-i2\pi n\alpha x) \sum_{n'=0}^{Q-1}}_{n\neq n'} \exp(i2\pi n'\alpha x) = \frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)} - Q \tag{17}$$

gives

$$\langle |t_2(x)|^2 \rangle = |P(x)|^2 \times Q \left[1 - |\rho(x)|^2 + \frac{|\rho(x)|^2}{Q} \frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)} \right]. \tag{18}$$

This expression is clearly similar to that derived for the regularly pixellated case, and indeed is the same if $\rho(x)$ is identically 1, corresponding to the ϕ_n being identically zero, as would be expected. Other specific choices of probability distribution function will be considered when the expression is generalized to two dimensions, but it is useful at this stage to make some general observations regarding $\rho(x)$.

(1) Equation (13) shows that if $P_{\phi_n}(\phi)$ is properly normalized i.e. $\int_{-\infty}^{\infty} P_{\phi_n}(d\phi)$ $(d\phi = 1)$, then $\rho(0)$ will be equal to 1, regardless of the form of $P_{\phi_n}(\phi)$.

(2) As $0 \leqslant P_{\phi_n}(\phi) \leqslant 1$, an upper bound of 1 can be established for $\rho(x)$ using the following general result for Fourier transform pairs [18]:

$$|\rho(x)| \leqslant \int_{-\infty}^{\infty} |P_{\phi_n}(\phi)| \, d\phi. \tag{19}$$

(3) Combining the results in (1) and (2) above shows that $|\rho(x)|^2$ is a maximum at $x = 0$ and must either decrease or remain equal to 1 as we move away from the origin.

These points help to show quantitatively how the randomization will affect the form of $\langle |t_2(x)|^2 \rangle$, the power spectrum of $T_2(\xi)$, and how randomization may be used to suppress the peaks of higher spectral order which give rise to well defined replicas.

It is apparent that, with the exception of the zero order, the spectral peaks resulting from the $[\sin^2(\pi Q \alpha x)]/[\sin^2(\pi \alpha x)]$ term are attenuated by the $|\rho(x)|^2$ multiplier. This suggests that with a suitable choice of probability distribution function, $|\rho(x)|^2$ may be tailored to have zeros at the centres of some or all of the spectral orders, thus to a good approximation eliminating them, with the remaining orders certainly no larger than for the corresponding regularly pixellated case.

The magnitude of the zero order is proportional to the total transmissive area of the array, and is thus independent of the chosen randomization scheme so long as the pixels do not overlap and reduce the total transmissive area; this follows from the general relation for Fourier transform pairs that $f(0) = \int_{-\infty}^{\infty} F(\xi) \, d\xi$, and from $\rho(0) = 1$. As overlapping pixels are neither practical nor particularly desirable in an SLM, all the cases considered will constrain the pixels from overlapping.

The $[1 - |\rho(x)|^2]$ part of the equation gives rise to a *light depletion region* [14]. This refers to a dark patch surrounding the zero order (where $1 - |\rho(0)|^2 = 0$) with increasing illumination arising from where $1 - |\rho(x)|^2$ has non-zero values. This background illumination is a redistribution of the energy which is removed from the higher spectral orders by $|\rho(x)|^2$ and must be conserved (expressed here by Rayleigh's theorem: $\int_{-\infty}^{\infty} |\rho(x)|^2 \, dx = \int_{-\infty}^{\infty} |P_{\phi_n}(\phi)|^2 \, d\phi$).

Little more can be said about the form or extent of the light depletion zone without choosing specific examples. However, if it is in the form of diffuse, low-level background illumination it will not contribute to the formation of distinct replicas when $t_2(x)$ is convolved with $t_1(-x)$ to give the output of the optical processor.

For the regular case the analysis in 2D can be easily generalized from the analysis in 1D because of the separability of the Fourier transform in the x and y directions. This is not the case with a randomly pixellated array, as can be seen when the transmission function of the 2D filter is written down:

$$T_2(\xi, \eta) = p(\xi, \eta) * \sum_{n=0}^{Q-1} \sum_{m=0}^{Q-1} \delta(\xi - n\alpha - \phi_{n,m}, \eta - m\gamma - \psi_{n,m}). \tag{20}$$

Here $\phi_{n,m}$ and $\psi_{n,m}$ are ξ and η components of the displacement of the centre of the pixel labelled by n and m from the position it would have occupied were it regularly positioned with pitch α in the ξ direction and pitch γ in the η direction. Q, n and m are integers.

The problem here is that the random displacements are specified for each pixel in the $Q \times Q$ array by both labels n and m, which are counters labelling the pixels in the ξ and η directions respectively. Thus neither the transmission function of the filter nor its Fourier transform is easily separable. Therefore $T_2(\xi, \eta)$ must be Fourier transformed directly to give

$$t_2(x,y) = P(x,y) \times \sum_{n=0}^{Q-1} \sum_{m=0}^{Q-1} \times \exp(-i2\pi n\alpha x) \exp(-i2\pi\phi_{n,m}x) \exp(-i2\pi m\gamma y) \exp(-i2\pi\psi_{n,m}y). \quad (21)$$

When intensities are considered and ensemble averaging takes place to allow an analytic solution to the equation we arrive at

$$\langle |t_2(x,y)|^2 \rangle = \langle |P(x,y)|^2 \rangle \times \left\langle \sum_{n=0}^{Q-1} \sum_{m=0}^{Q-1} \sum_{n'=0}^{Q-1} \sum_{m'=0}^{Q-1} \{\exp[-i2\pi(n-n')\alpha x] \exp[-i2\pi(m-m')\gamma y] \times \exp(-i2\pi\phi_{n,m}x) \exp(i2\pi\phi_{n',m'}x) \exp(-i2\pi\psi_{n,m}y) \exp(i2\pi\psi_{n',m'}y)\} \right\rangle \quad (22)$$

This can be simplified greatly if the values of the random variables are mutually independent. This is the case except when $n = n'$ and $m = m'$, fixing the values of $\phi_{n',m}$ and $\psi_{n',m}$ if $\phi_{n,m}$ and $\psi_{n,m}$ are known. If these cases are explicitly removed from the summation then the averaging can be performed over each function of the random variables, to give

$$\langle |t_2(x,y)|^2 \rangle = |P(x,y)|^2 \times \left[Q^2 + \underbrace{\sum_{n=0}^{Q-1} \sum_{n'=0}^{Q-1} \exp[-i2\pi(n-n')\alpha x] \sum_{m=0}^{Q-1} \sum_{m'=0}^{Q-1} \exp[-i2\pi(m-m')\gamma y]}_{n,m \neq n'm'} \right] \times \langle \exp(-i2\pi\phi_{n,m}x) \rangle \langle \exp(i2\pi\phi_{n',m'}x) \rangle \langle \exp(-i2\pi\psi_{n,m}y) \rangle \langle \exp(i2\pi\psi_{n',m'}y) \rangle. \quad (23)$$

Substituting for the $\langle \ldots \rangle$ terms, and rearranging using

$$\underbrace{\sum_{n=0}^{Q-1} \sum_{n'=0}^{Q-1} \exp[-i2\pi(n-n')\alpha x] \sum_{m=0}^{Q-1} \sum_{m'=0}^{Q-1} \exp[-i2\pi(m-m')\gamma y]}_{n,m \neq n'm'} = \frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)} \frac{\sin^2(\pi Q\gamma y)}{\sin^2(\pi\gamma y)} - Q^2 \quad (24)$$

gives the general 2D equation for a $Q \times Q$ randomized pixel array:

$$\langle |t_2(x,y)|^2 \rangle = |P(x,y)|^2 \times Q^2 \left[1 - |\rho(x)|^2 |\rho(y)|^2 + \frac{1}{Q^2} \frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)} \frac{\sin^2(\pi Q\gamma y)}{\sin^2(\pi\gamma y)} |\rho(x)|^2 |\rho(y)|^2 \right], \quad (25)$$

where $\rho(y)$ is the Fourier transform of the probability distribution function $P_{\psi_{n,m}}(\psi)$.

This equation is similar to that of the 1D case and indeed is exactly the same along the axes. More important, however, is that the general form is the same, with the $|\rho(\,.\,.)|^2$ terms attenuating the spectral orders, the redistributed energy giving rise to a background illumination with a light depletion zone surrounding the zero order, and the zero order being unchanged in magnitude. Further, the ratio of background and zero order intensities is inversely related to the number of pixels i.e.

$$\frac{1-|\rho(x)|^2|\rho(y)|^2}{\langle|t_2(0,0)|^2\rangle} \leqslant \frac{1}{Q^2}, \tag{26}$$

with the mean background contribution proportional to $[1-|\rho(x)|^2|\rho(y)|^2]$.

Thus in the real 2D case it is possible that some or all of the replicas may be suppressed, depending on the randomization scheme used.

5. Specific examples of randomized arrays

Two different randomization schemes have been tried, both subject to the constraint that pixels do not overlap, which is necessary for practical implementation on an SLM. Choosing the pixels to be square with side β means that

$$-\frac{\alpha-\beta}{2} \leqslant \phi_{n,m}, \psi_{n,m} \leqslant \frac{\alpha-\beta}{2}, \tag{27}$$

and makes $p(\xi,\eta)=\text{rect}\,(\xi/\beta,\eta/\beta)$, and $P(x,y)=\beta^2\,\text{sinc}\,(\beta x,\beta y)$ as before.

5.1. *Uniformly random distribution*

The first randomization scheme considered has $\phi_{n,m}$ uniformly distributed over the allowed range $-(\alpha-\beta)/2 \leqslant \phi_{n,m} \leqslant (\alpha-\beta)/2$, and might intuitively be thought a good candidate for removing the spectral orders. This random distribution of the offsets conforms to a normalized probability distribution function

$$P_{\phi_{n,m}}(\phi)=\frac{1}{\alpha-\beta}\,\text{rect}\left(\frac{\phi}{\alpha-\beta}\right). \tag{28}$$

Fourier transformation of $P_{\phi_{n,m}}(\phi)$ gives $\rho(x)=\text{sinc}\,[(\alpha-\beta)x]$, and similarly $\rho(y)=\text{sinc}\,[(\alpha-\beta)y]$. Setting $y=0$ establishes the form of the diffraction pattern along the x axis as

$$\langle|t_2(x,0)|^2\rangle = Q^2\beta^4\,\text{sinc}^2\,(\beta x)\times\left\{1-\text{sinc}^2\,[(\alpha-\beta)x] + \frac{1}{Q^2}\frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)}\,Q^2\,\text{sinc}^2\,[(\alpha-\beta)x]\right\}. \tag{29}$$

The $1-\text{sinc}^2\,[(\alpha-\beta)x]$ part represents the light depletion zone. The spectral orders resulting from the $[\sin^2(\pi Q\alpha x)]/[\sin^2(\pi\alpha x)]$ part are attenuated by the $\text{sinc}^2\,(\beta x)$ factor from the transform of the pixel function as before, and additionally by the $\text{sinc}^2\,(\alpha-\beta)$ part due to the randomized scheme. Figure 4 shows an array which is a member of the ensemble of possible arrays, and its power spectrum, $|t_2(x,y)|^2$, for $Q=16$ and $\beta=\alpha/2$. Figure 5 shows the cross-section along the x axis of the squared modulus of the computed Fourier transform of the random array shown in figure 4. The length of the pixel-side is chosen to be $\alpha/2$ to eliminate the higher

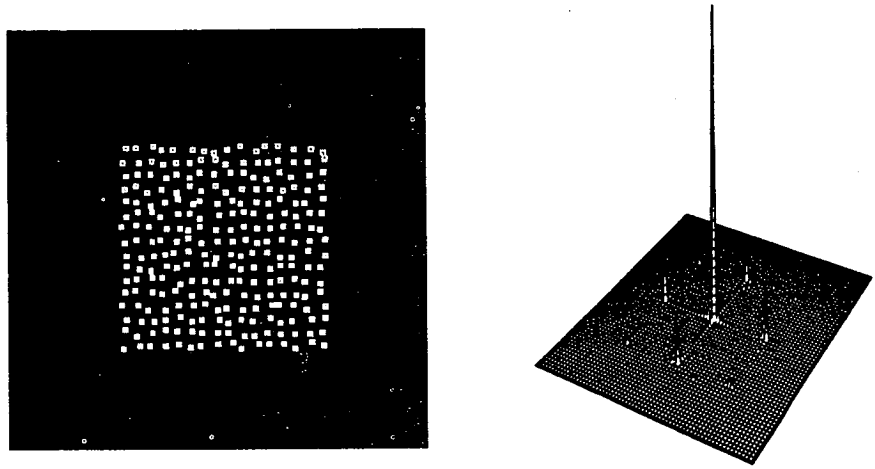

(a) (b)

Figure 4. (a) Uniformly random 16×16 array, with $\beta = \alpha/2$. (b) Power spectrum, $|t_2(x, y)|^2$, of (a).

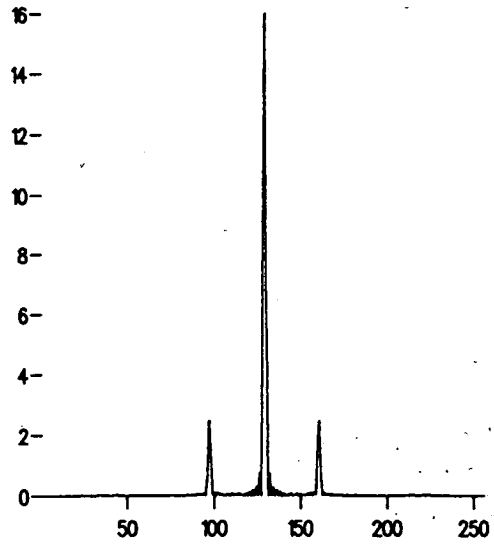

Figure 5. Power spectrum along the x axis, $|t_2(x, 0)|^2$, for uniformly random array, with $\beta = \alpha/2$.

even spectral orders. This is not necessarily the optimum pixel size to attenuate the odd spectral orders as shown later in the table.

Figure 5 shows that the analytic solution for the ensemble averaged power spectrum of the array represents an individual member of the ensemble well, with the higher even orders missing, the odd orders attenuated with respect to that from the regular array, and the addition of a low-level diffuse background. Although the odd orders are smaller than before, they will only be eliminated using this randomization scheme if $\beta = \alpha$, which is unrealizable, or if $\beta = 0$, which is of no practical use.

5.2. *Discrete position random distribution*

The second approach is to choose explicitly a $\rho(x)$ which has a zero at $x=1/\alpha$, the position of the first spectral order, and preferably zeros at the higher odd spectral orders as well, using $P(x, y)$ to eliminate the higher even orders. This $\rho(x)$ can then be inversely Fourier transformed to establish the probability distribution function and discover if it is realizable, given the non-overlapping condition.

A simple function which is suitable is $\rho(x)=\cos((\pi\alpha/2)x)$, which has zeros at the centres of all the odd spectral orders. Inverse transforming $\rho(x)=\cos(\pi\alpha x/2)$ gives $P_{\phi_{n,m}}(\phi)=\frac{1}{2}\delta(\phi-\alpha/4)+\frac{1}{2}\delta(\phi+\alpha/4)$. With $\rho(y)$ specified similarly, i.e. $\rho(y)=\cos((\pi\alpha/2)y)$, the joint probability distribution is

$$P_{\phi_{n,m},\psi_{n,m}}(\phi,\psi)=\frac{1}{4}\delta\left(\phi-\frac{\alpha}{4},\psi-\frac{\alpha}{4}\right)+\frac{1}{4}\delta\left(\phi-\frac{\alpha}{4},\psi+\frac{\alpha}{4}\right)$$
$$+\frac{1}{4}\delta\left(\phi+\frac{\alpha}{4},\psi-\frac{\alpha}{4}\right)+\frac{1}{4}\delta\left(\phi+\frac{\alpha}{4},\psi+\frac{\alpha}{4}\right). \tag{30}$$

Thus there are four discrete positions that the pixel can lie at, within the α by α cell that is allowed under the non-overlapping condition. Also, given this condition and probability distribution function, the maximum size the pixel can be is $\beta=\alpha/2$.

A member of the ensemble of possible arrays is shown in figure 6, along with the resulting power spectrum.

Setting $y=0$, this probability distribution function gives the form of the power spectrum along the x axis as

$$\langle|t_2(x,0)|^2\rangle=Q^2\beta^4\,\mathrm{sinc}^2\left(\frac{\alpha}{2}x\right)\left[1-\cos^2\left(\frac{\pi\alpha}{2}x\right)+\cos^2\left(\frac{\pi\alpha}{2}x\right)\frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)}\right]$$
$$=\underbrace{Q^4\beta^4\,\mathrm{sinc}^2(\pi Q\alpha x)}_{\text{zero order}}+\underbrace{Q^2\beta^4\frac{\sin^4((\pi\alpha/2)x)}{((\pi\alpha/2)x)^2}}_{\text{low level background}}, \tag{31}$$

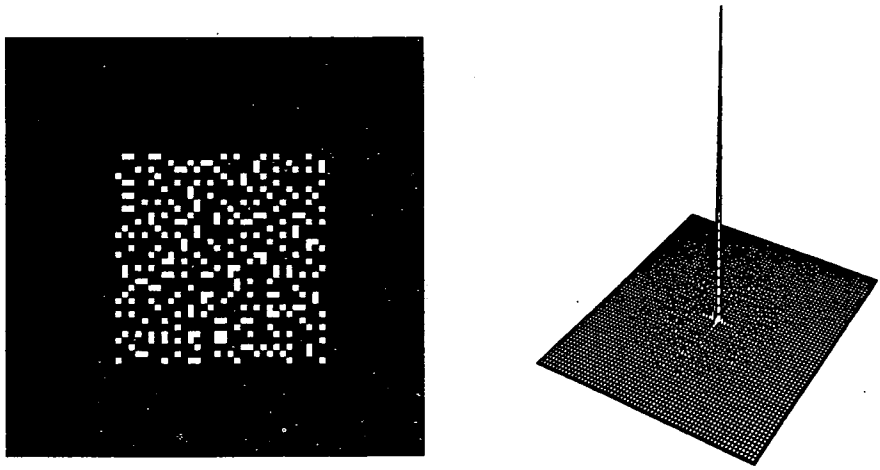

(*a*) (*b*)

Figure 6. (*a*) Discrete positions random 16×16 array, with $\beta=\alpha/2$. (*b*) Power spectrum, $|t_2(x,y)|^2$, of (*a*).

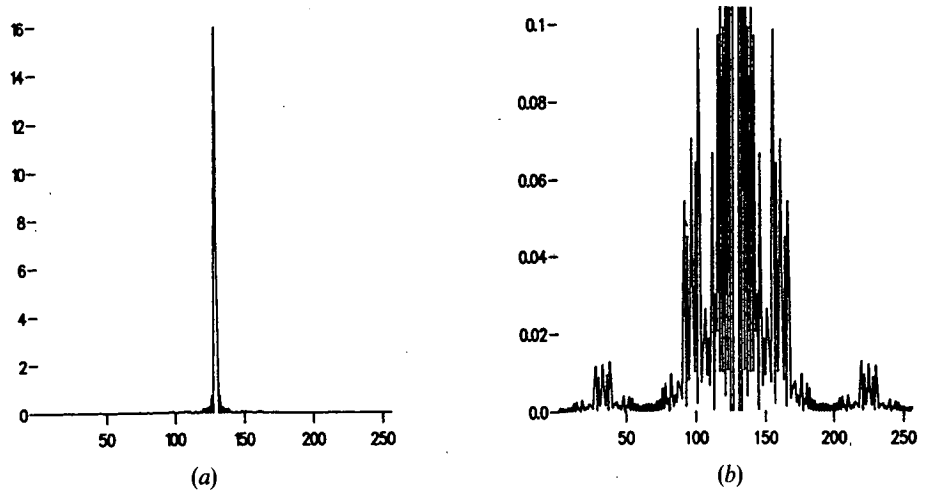

Figure 7. Power spectrum along the x axis, $|t_2(x,0)|^2$, for discrete positions random array, with $\beta=\alpha/2$. (*a*) Range of intensity axis same as in figures 3 and 5. (*b*) Range of intensity axis reduced to show low level detail.

which shows the presence of only the zero order and a background illumination of the order of Q^2 times smaller than the zero order and of the order Q^δ times broader. The cross-section along the x axis of the power spectrum of the array shown in figure 6 again matches the analytic solution for the ensemble averaged power spectrum very well. The cross-section is shown in figure 7 (*a*), and again in figure 7 (*b*) with the intensity axis rescaled to show the detail of the low level background, with the rapid oscillations of the zero order's $\mathrm{sinc}^2(Q\alpha x)$ function added to the low level background term.

This randomization scheme demonstrates how the clearly defined replicas can be eliminated, and how the energy contained in the higher spectral orders is diffused into low level illumination constituting the light depletion zone. The maximum value of this light depletion zone decreases relative to the zero order as the number of pixels in the SLM increases, and its distribution becomes more diffuse.

A similar discrete position randomization scheme can be employed for pixels with $\beta \geqslant \alpha/2$. The non-overlapping condition then requires that the random displacements are correspondingly smaller, with the joint probability distribution defined as

$$P_{\phi_{n,m},\psi_{n,m}}(\phi,\psi)=\frac{1}{4}\delta\left(\phi-\frac{\alpha-\beta}{2},\psi-\frac{\alpha-\beta}{2}\right)+\frac{1}{4}\delta\left(\phi-\frac{\alpha-\beta}{2},\psi+\frac{\alpha-\beta}{2}\right)$$
$$+\frac{1}{4}\delta\left(\phi+\frac{\alpha-\beta}{2},\psi-\frac{\alpha-\beta}{2}\right)+\frac{1}{4}\delta\left(\phi+\frac{\alpha-\beta}{2},\psi+\frac{\alpha-\beta}{2}\right), \quad (32)$$

and $\rho(x)=\cos^2[\pi(\alpha-\beta)x]$, $\rho(y)=\cos^2[\pi(\alpha-\beta)y]$. While this will not eliminate the first order unless $\beta=\alpha/2$, the first order will still be attenuated compared to that of the regularly spaced array with the same shape and size of pixel. This is shown in the table.

Comparison of first-order intensities for regular and randomized arrays, relative to normalized zero order.

	Relative intensities of first orders		
β/α	Regular case[1]	Uniform random distbn.[2]	Discrete random distbn.[3]
0·95	0·0082	0·0081	0·0079
0·90	0·0119	0·0115	0·0108
0·80	0·0547	0·0479	0·0358
0·70	0·1353	0·0997	0·0467
0·60	0·2546	0·1458	0·0243
0·50	0·4052	0·1642	0·0000

[1] Function attenuating spectral orders is $\text{sinc}^2(\beta x)$.
[2] Function attenuating spectral orders is $\text{sinc}^2(\beta x)\,\text{sinc}^2((\alpha-\beta)x)$.
[3] Function attentuating spectral orders is $\text{sinc}^2(\beta x)\cos^2(\pi(\alpha-\beta)x)$.

6. Comparison of power spectra of regular and randomized arrays

The table shows a direct comparison between the power spectra of the regular and randomized arrays described, by tabulating the intensity of the spectral order nearest the origin for various ratios of $\beta:\alpha$.

The table shows how both the randomization schemes described reduce the magnitude of the first spectral order compared to that of the regular array. As intimated earlier, the choice of $\beta=\alpha/2$ to eliminate the higher even spectral orders is not the optimum choice for attenuating the odd orders when using the uniformly randomized array of Section 4.1, and in actual fact gives the least possible attentuation for that randomization scheme.

In every case the discrete case randomization scheme attentuates the first order the most, with all the higher orders being eliminated for $\beta=\alpha/2$. If $\beta<\alpha/2$ the odd orders can still be eliminated if the allowed position for the pixel centres is as for $\beta=\alpha/2$, i.e. $\phi_{n,m}, \psi_{n,m}=\alpha/4$ or $-\alpha/4$. This is important if the area that the pixel is allowed to lie in is slightly less than $\alpha\times\alpha$ (as is necessary to prevent electrical short-circuits when using a matrix-addressed SLM) thus preventing the size of the pixel from being as large as $\alpha/2$. If ϵ is the separation necessary to prevent short-circuits between adjacent pixels then the transformed pixel function will become $P(x)=(\alpha/2-\epsilon)\,\text{sinc}\,(\alpha/2-\epsilon)x$ and the higher even orders will be eliminated; however, so long as ϵ is small compared with α, the effect of any spectral orders higher than the first should be minimal due to the decreasing $P(x)$ being small-valued at their positions. Even for the extremely low resolution case of a 16×16 array, the effect of the light depletion zones which arise in the random array's spectrum is not intrusive, becoming apparent only on closer inspection of the spectra (figure 7 (*b*)). As noted the magnitudes of the zero orders of all three cases are the same, for the same total transmissive area. The analysis also shows that as the magnitude of the light depletion zone is zero at the origin, and rises slowly with respect to the width of the zero order in the two random examples described, the zero order will effectively be unchanged, leading to negligible loss of sharpness in the central image. This is illustrated in figure 8 which shows the results of optical simulations carried out on a

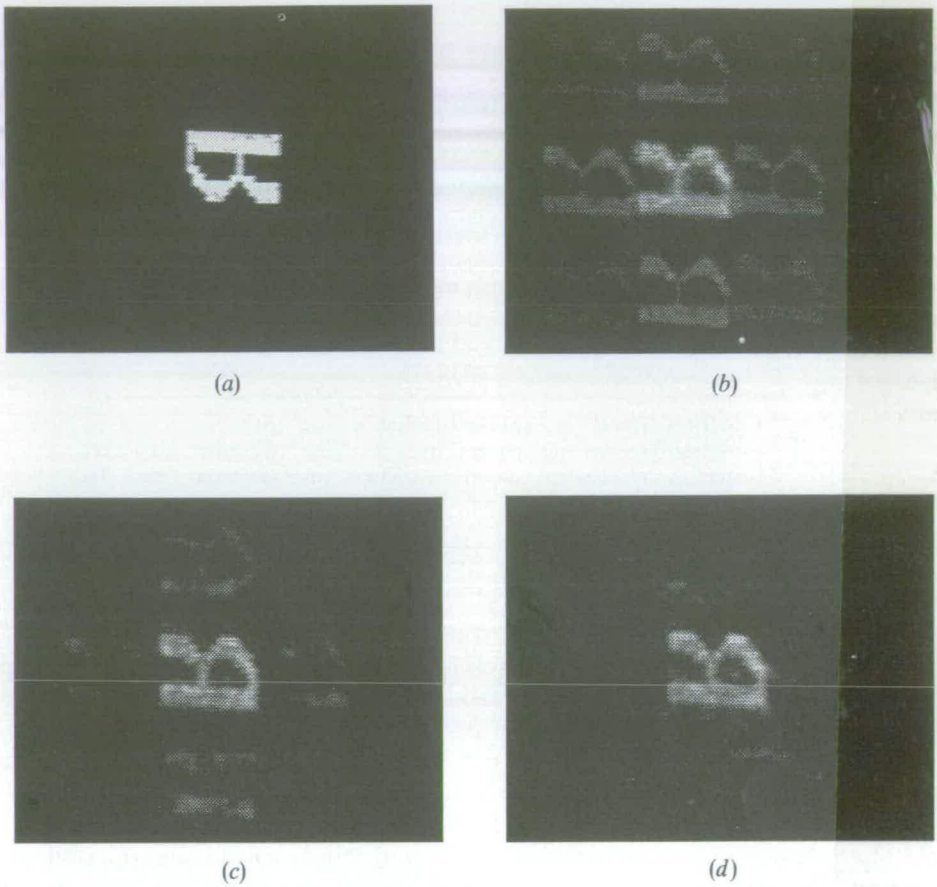

Figure 8. Optical results of imaging through the filter arrays: (*a*) Input scene: letter 'R'. (*b*) Output using array of figure 2 as Fourier plane filter. (*c*) Output using array of figure 4 as Fourier plane filter. (*d*) Output using array of figure 6 as Fourier plane filter.

Fourier optical processor of the kind shown in figure 1. The input to the processor is a transparency, opaque except for a letter 'R' defined on a 15×15 grid of 'pixels' (shown in figure 8 (*a*)). The Fourier spectrum of the input is sampled at the Nyquist frequency by each of the three arrays shown in figures 2 (*a*), 4 (*a*), 6 (*a*) in turn, and the corresponding outputs shown in figure 8 (*b–d*). The filter arrays are in the form of aluminized optical flats with the transmissive pixel areas defined lithographically. The images in figure 8 were captured on a framestore using a 512×512 pixel array camera with a manual gain facility and gamma equal to 1. The illumination was attenuated to prevent the camera saturating. The optical results show that the randomization has been effective in eliminating any distinct replicas.

7. Conclusions

We have shown how the distinct spectral orders arising from the underlying regular array structure of a pixellated SLM lead to replication, and the consequent problem of false correlations when using a low resolution SLM matched to a small

target in a large input scene has been described. We have proposed a solution—randomization of the pixel positions—which has been shown to be effective in removing unwanted spectral orders, and redistributing their energy into a diffuse low-level background illumination which falls to zero around the zero order; the zero order itself is unchanged in magnitude for a given pixel size. The power spectra of randomized arrays has been analysed and used to predict a 'discrete position' randomization scheme, which results in the removal of the first spectral orders for pixel sizes half the pixel pitch or less. For larger pixel sizes the same randomization scheme is effective in attenuating the spectral orders with respect to the corresponding regular array.

Following the experiments using metallized masks, SLMs based on the array in figure 6 (*a*) have been constructed. These are configured as amplitude and phase-only filters, and will be used in correlation experiments.

Acknowledgments

The idea that the intensity of the image replicas might be reduced by randomizing the pixel positions was posed to us as a question by Dr D. G. Vass; its validity thereafter demonstrated by two Final Honours students, Messrs. A. M. McIntosh and W. Yuill.

We would like to thank Dr D. G. Vass, Dr W. J. Hossack and Mr D. J. Potter, for useful discussions during the preparation of this paper, and particularly Mr I. W. Chisholm, whose software was used in the simulations. S. Heddle acknowledges the financial support of the SERC.

References

[1] PSALTIS, D., PAEK, E. G., and VENKATESH, S. S., 1984, *Opt. Engng*, **23**, 698.
[2] HORNER, J. L., and LEGER, J., 1985, *Appl. Optics*, **24**, 609.
[3] COTTRELL, D. M., LILLY, R. A., DAVIS, J. A., and DAY, T., 1987, *Appl. Optics*, **26**, 3755.
[4] ROSS, W. E., PSALTIS, D., and ANDERSON, R. H., 1983, *Opt. Engng*, **22**, 485.
[5] UNDERWOOD, I., WILLSON, P. H., SILLITTO, R. M., and VASS, D. G., 1987, *Proc. Soc. Photo Opt. Instrum. Eng.*, **860**, 2.
[6] MCKNIGHT, D. J., VASS, D. G., and SILLITTO, R. M., 1989, *Appl. Optics*, **23**, 4757.
[7] POTTER, D. J., RANSHAW, M. J., AL-CHALABI, A. O., FANCEY, N. E., SILLITTO, R. M., and VASS, D. G., 1991, *Proc. Soc. Photo-Opt. Instrum. Eng.*, **1564**, 363.
[8] JOHNSON, K. M., HANDSCHY, M. A., and PAGANO-STAUFFER, L. A., 1987, *Opt. Engng*, **26**, 385.
[9] GASKILL, J., 1978, *Linear Systems, Fourier Transforms and Optics* (New York: Wiley), pp. 62–63.
[10] YAU, H. F., TONG, Y. P., LEE, C. M., CHANG, C. C., and CHANG, M. W., 1992, *Jap. J. appl. Phys.*, **31**, 49.
[11] GHEEN, G., WASHWELL, E., and ARMITAGE, D., 1990, *Technical Digest on Spatial Light Modulators and Applications* (14) (Washington, DC.: Optical Society of America), pp. 161–164.
[12] HOSSACK, W. J., VASS, D. G., and UNDERWOOD, I., 1991, *Proc. Soc. Photo-Opt. Instrum. Eng.*, **1564**, 697.
[13] DEBYE, P., 1913, *Verh. dt. phys. Ges.*, **15**, 678, 738, 857.
[14] STARK, H., 1977, *J. opt. Soc. Am.*, **67**, 700.
[15] STARK, H., 1978, *J. opt. Soc. Am.*, **68**, 1785.
[16] AIME, C., 1978, *Astron. Astrophys.*, **67**, 1.
[17] MARTIN, F., and AIME, C., 1979, *J. opt. Soc. Am.*, **69**, 1319.
[18] BRACEWELL, R. N., 1986, *The Fourier Transform and its Applications*, second edition, revised (Singapore: McGraw-Hill), p. 157.

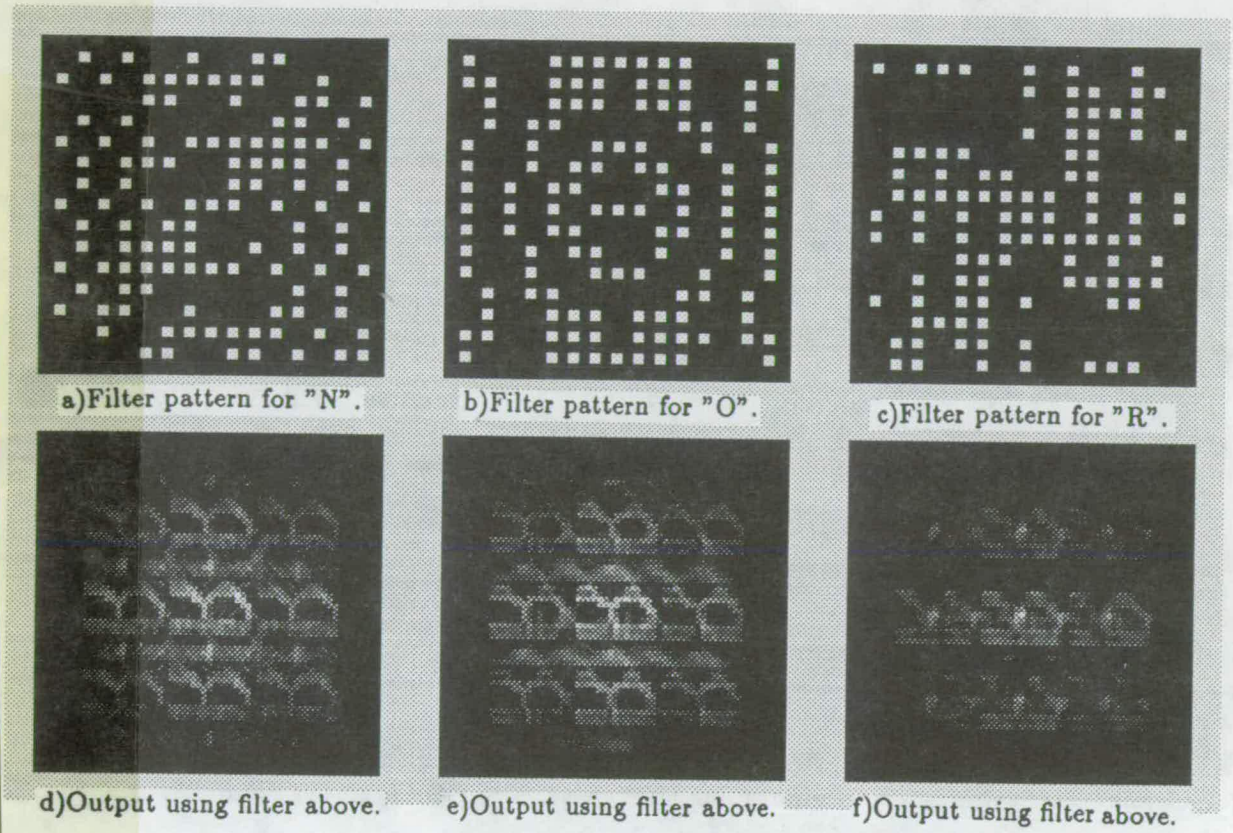

Figure 8: Filters for regular pixel position case, and correlation outputs with "R" as input

Input	Regular filter for...			Randomised filter for...		
	N	O	R	N	O	R
N	0.402	0.088	0.071	0.379	0.068	0.045
O	0.073	0.360	0.083	0.107	0.407	0.070
R	0.071	0.058	0.160	0.101	0.125	0.183

Table 1: Table of unnormalised correlation results for single letter input

replication is not noticeable, with only one sharp correlation peak for the "R" imaged through its corresponding filter. The coordinate reversed input scene is also evident in the output, as expected with an AEBPOF.[6]

5.2 Multiple letters as input

When multiple letters as shown in Figure 7 are used as the input, we have the case where the SBP of the input is greater than that of the filter, which is matched to each letter but not to the total input scene containing the three letters. Thus replication will occur within the region of the correlation plane where we wish to detect the correlation signals, when a regular array is used. Further, as the letters are adjacent, their first order replicas will coincide, potentially leading to large, false correlation signals. Using the filters with the randomised pixel positions we would expect that the absence of distinct replicas would allow the position of each letter to be still be distinguisable.

The correlation results for the regular and randomised cases are shown in Figures 10 and 11 respectively. Again each ouput has been scaled to a maximum of 255. Nevertheless, the correlation peaks for each letter should be obvious, with the "N" to the right, the "O" in the centre, and the "R" to the left. The figures show the central 64x64 point region of the correlation plane.

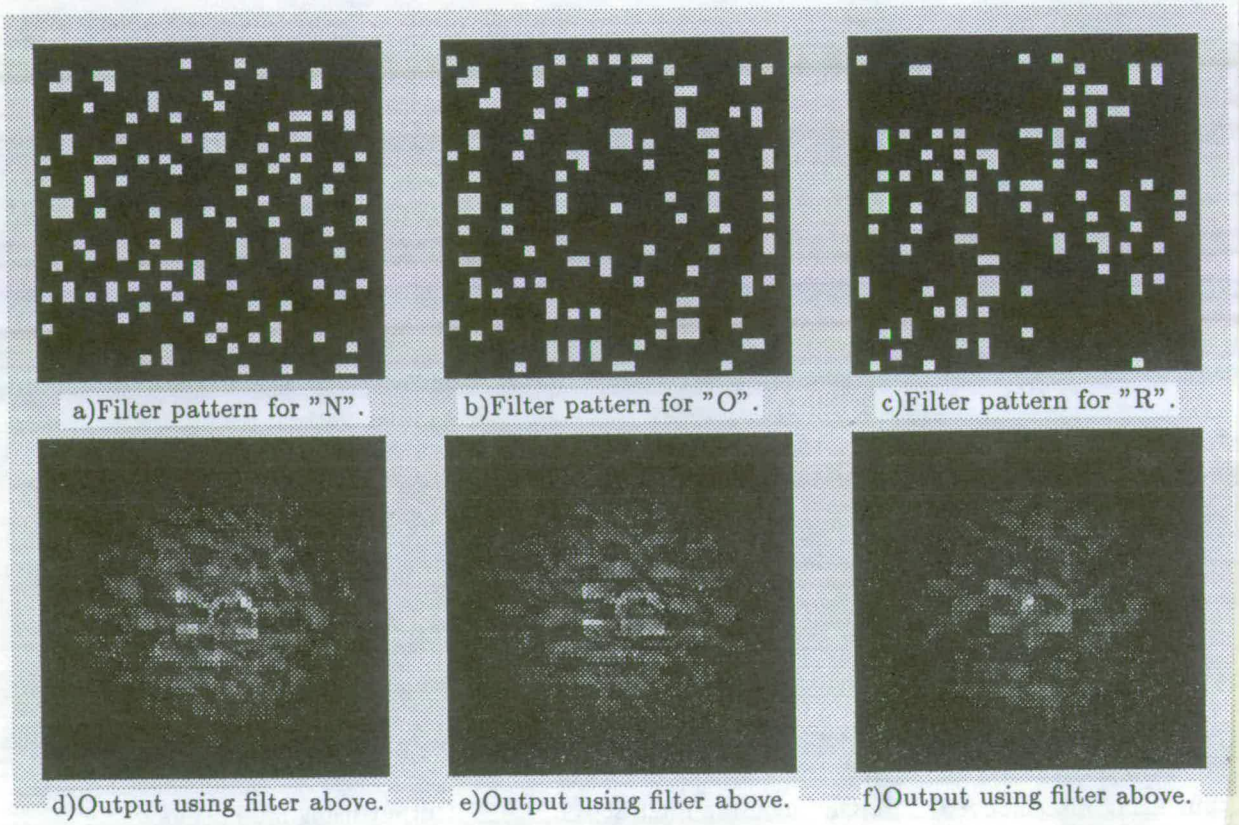

Figure 9: Filters for randomised pixel position case, and correlation outputs with "R" as input
Note comment about scaling of correlation outputs.

The correlation outputs of Figure 10 illustrate the replication problem well, with particularly large,sharp, false peaks occuring using the filters for "N" and "O", where the filters are as in Figure 8. In the case of the output when the filter for the "R" is used, the correlation peak, though sharp, is lower in intensity than the noise in the vicinity of the "O".

The results of the simulations using the randomised filter show a distinct improvement, with unambiguous sharp correlation peaks occuring for the letters that each filter sets out to detect, with the correlation peak being the highest intensity point in the region of interest.

Table 2 shows the unnormalised correlation values for each filter when the multiple letters of figure 7 are used as the input. The correlation values in the table are simple measures of maximum intensity in the region of each letter in the output and need not be related to a sharp peak, though for the true correlations they are.

Input	Regular filter for...			Randomised filter for...		
	N	O	R	N	O	R
N	0.502	0.278	0.159	0.396	0.138	0.101
O	0.294	0.499	0.201	0.228	0.586	0.147
R	0.124	0.179	0.167	0.140	0.141	0.174

Table 2: Table of unnormalised correlation results for multiple letter input

This should be noted when comparing the filter results for detecting the "N"– the figures in the table suggest that the regular filter outperforms the random filter, while taking no account of the fact that the maximum false correlation value is much more sharply peaked for the regular filter and more likely to be mistaken for a true correlation. For the filters to detect the "O" and the "R", the filters with the randomly positioned pixels clearly

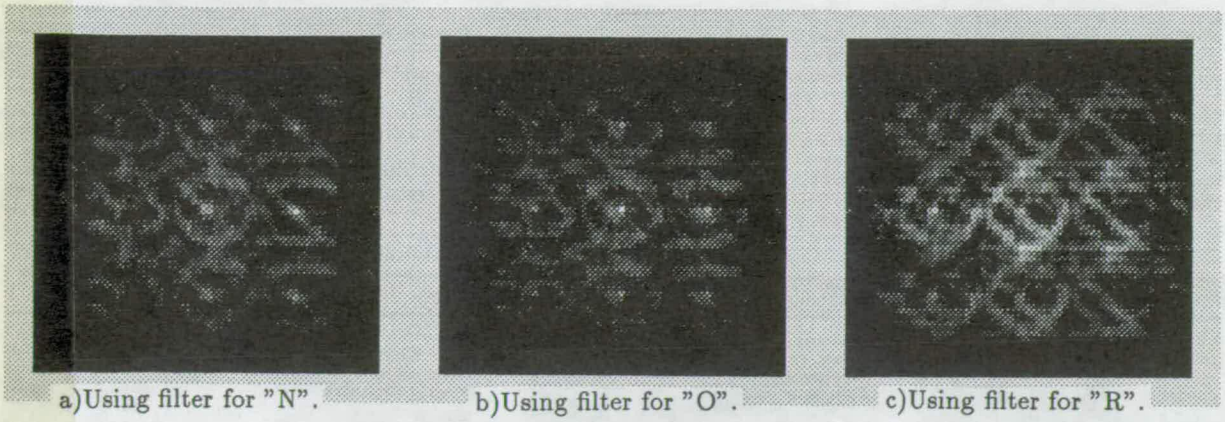

a)Using filter for "N". b)Using filter for "O". c)Using filter for "R".

Figure 10: Correlation outputs for regular case with input as Figure 7

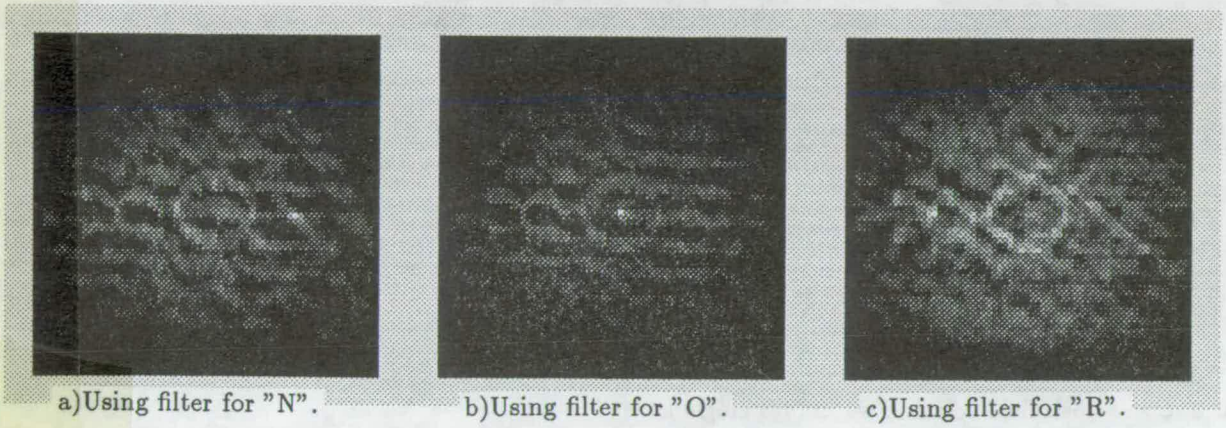

a)Using filter for "N". b)Using filter for "O". c)Using filter for "R".

Figure 11: Correlation outputs for random case with input as Figure 7

perform best, with higher actual correlation values and lower, broader, "false" correlation values.

6 OPTICAL CORRELATION RESULTS

The SLM described earlier was used to verify the results of the simulations, by being included as the Fourier plane filter in an optical processor of the type shown in Figure 1.

Input to the processor was provided by transparencies of the letters represented in the computational simulations, individually and side by side as in the simulations. The transparencies were mounted in the input plane in a liquid gate, to remove phase variations introduced in the transmissive regions of the transparency by the photographic film.

The letters were required to be scaled so that their Fourier transforms were matched to the SLM. Given that the underlying pitch of the pixels was 500 μ, the laser wavelength 633nm and the focal lengths of the transforming lenses were 1m, this resulted in the letters needing to be $(\lambda f)/\alpha = 1.26$mm high. Once made, the letters in the transparencies were actually (1.23 ± 0.02)mm high, resulting in a slight mismatch of the Fourier transform to the SLM.

After careful alignment of the SLM, necessary because of the adverse effect of rotational misalignment on the filters' performance, the filter patterns shown in Figure 9 were written to the SLM, and the intensities in the output plane recorded using a calibrated 512x512 pixel array camera and framestore. The camera controller was set with gamma=1 and on manual gain to prevent erroneous scaling of intensities.

Optical correlation outputs are shown in Figures 12 and 13. The outputs are again displayed here with

gamma=1.5 to retain detail, and for comparison with the earlier simulations. These figures have not had their maximum values scaled to 255.

Figure 12 shows the output with the random pixel position SLM displaying the filter patterns for detecting "N", "O" and "R" in turn with a single "R" as input, and is directly comparable with figure 9. Similarly, Table 3 can be compared with Table 1, where the unnormalised correlation values in Table 3 are read from the framestore.

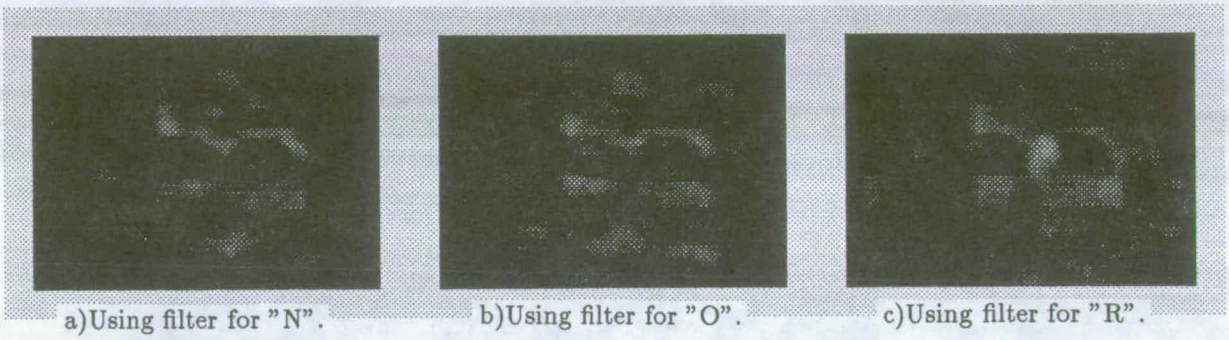

a)Using filter for "N". b)Using filter for "O". c)Using filter for "R".

Figure 12: Optical correlation outputs using randomised SLM with single "R" as input

Input	Randomised filter for...		
	N	O	R
N	188	59	56
O	69	151	74
R	93	85	118

Table 3: Table of optical correlation results for single letter input

Figure 13 shows the optical correlation results for the multiple letter input, and Table 4 the correlation value read from the framestore.

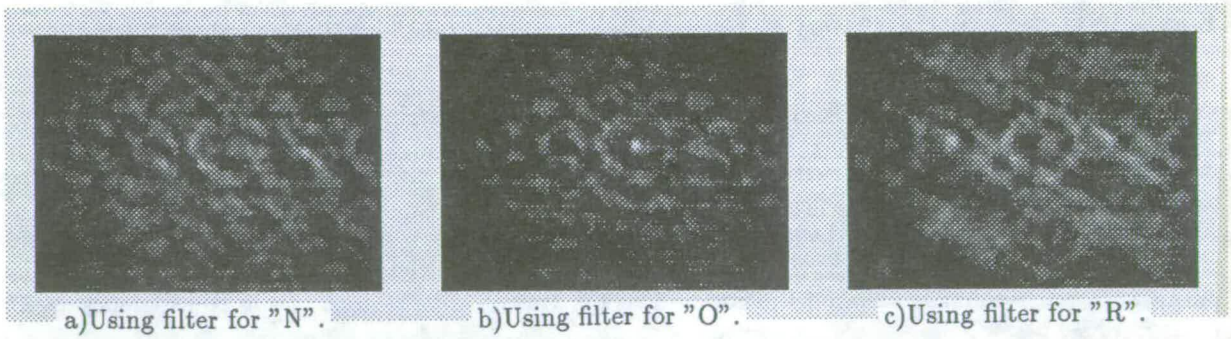

a)Using filter for "N". b)Using filter for "O". c)Using filter for "R".

Figure 13: Optical correlation outputs using randomised SLM with multiple letters as input

The optical correlation results are in very good visual agreement with the simulations. Numerically, the intensities of the true correlation signals are smaller relative to the noise than predicted by the simulations, but still show good correlation performance, with the correct letter detected in each case. The performance of the SLM could be improved if inevitable phase variations over its area could be quantified and compensated for in the filter design. The small size discrepancy of the input objects should be noted, as well as the fact that the pixels on the SLM are slightly less than half the size of the underlying pitch, resulting in very faint second order replicas.

- Underlying pixel pitch = 500μm
- Contrast ratio > 20:1
- Frame rate Not known– limited by BBC microcomputer.

Note that the pixel size is less than half that of the underlying pitch– this means that the higher even spectral orders will not be eliminated completely, though the odd orders should still be eliminated by the randomisation.

5 CORRELATION SIMULATIONS FOR AEBPOFS

Computational simulations were carried out, using a 256x256 point FFT routine, to simulate correlation using an Amplitude Encoded Binary Phase Only Filter.[6]

In the simulations the input scenes were provided by binary representations of the letters N,O, and R, individually and side by side as shown in Figure 7. The input letters were each defined on a 15x15 point grid in the 256x256 point input plane, with the filter filling the filter plane and the filter pixels each occupying 8x8

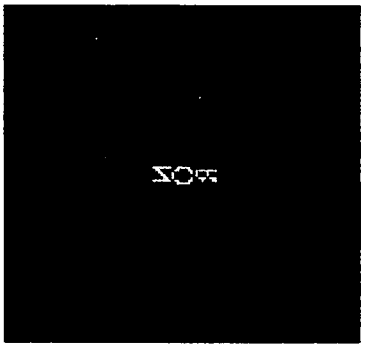

Figure 7: Input of letters N,O and R, side by side

points. The underlying pixel pitch was 16 points, to sample the transform of the input letter(s) at the Nyquist frequency. A filter function to detect each letter was calculated by summing its Fourier transform over each pixel and binarising the pixel value to 1 if the real part of the Fourier transform was positive or zero, and 0 if it was negative. This was done for the array shown in Figure 6, and for the corresponding regularly pixelated case.

The correlation outputs show intensities, and have been displayed using the package "xv" with the gamma value of the image set at 1.5. The effect of this is to emphasise the lower intensities, and prevent detail being lost when this paper is reproduced.

5.1 Single letter as input

With a single letter as input the SBP's of the input and filter are matched, and we can note if the randomisation degrades the fundamental performance of the correlator. With the letter centred in the input plane this corresponds to the case where the region of interest in the correlation plane can be restricted to a 15 x 15 region which excludes the replicas.

Figure 8 shows the filters for the regular case, with the corresponding correlation results for a single R as input, underneath. Figure 9 is the equivalent figure for the randomised case. Note that in both these figures the correlation outputs have been normalised by the plotting package such that the highest value is scaled to 255. The unnormalised peak correlation values are tabulated in Table 1. Note also that the correlation outputs show only the central 64x64 points of the correlation plane.

The table of results helps show that the correlation performance of the regular and randomised pixel filters is not significantly different, in the case of a single centred target constituting the entire input scene– for each letter the maximum correlation value came from its corresponding filter, and for each filter the maximum correlation came from the corresponding letter, for both the regular and randomised cases. The replicated outputs are clearly visible in the regular case, and show additional correlation peaks for the filter for the "R". In the random case

the transmissive conducting electrodes which sandwich the liquid crystal layer, and which define the switchable pixel regions by their overlap. Some details of the SLM's construction and operation follow.

4.1 The electrode pattern

The electrodes are patterned onto glass using a "lift-off" process whereby the non-conducting regions are covered with photoresist, which is stripped off after the glass has been sputtered with Indium Tin Oxide. The randomised array described in the previous section cannot be fabricated directly, as the positioning of neighbouring pixels would lead to unwanted open circuits being defined by the photomask, as shown in Figure 5.

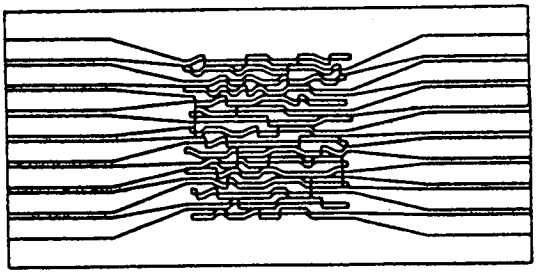

Figure 5: Photomask pattern showing breaks in electrodes

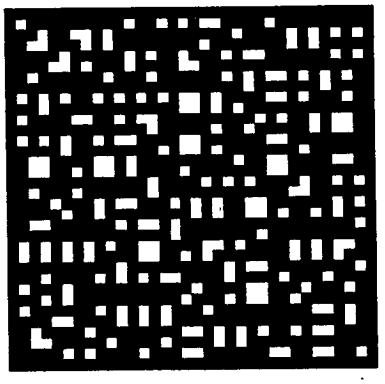

Figure 6: Pixel array used in the SLM

Instead the array shown in Figure 6 was used, which has the same allowed pixel positions as before, but with the choice of position subject to an additional nearest neighbour condition, weighted to help ensure that each possible position is equally probable. The diffraction pattern of this array is as before, but the problem of open circuits is eliminated.

4.2 Liquid crystal layer

As the device was to be used as an amplitude modulator, a 90° twisted nematic configuration was chosen with polarisers either side of the SLM oriented so that minimum transmission occured when the SLM was not driven. The LC layer was set to be 12μm thick by polyester spacer strips placed between the glass surfaces patterned with the electrodes. The 90° twist in director orientation was effected using rubbed PVA alignment layers. The liquid crystal used was BDH E300+0.1% CB15 chiral dopant, added to ensure only one handedness of twist.

4.3 Connection and interfacing

Electrical connection to the SLM is made using push on PCB edge connectors on a 0.1" pitch. The multiplexed drive scheme is implemented using a Hitachi LCD driver IC, which generates the requisite waveforms to maintain DC balancing, and accepts frame information written to it from a BBC microcomputer. SLM contrast can be optimised by applying a negative DC voltage to one pin of the IC, which varies the amplitude of the drive waveforms.

4.4 Device characteristics

- Number of pixels =16x16 = 256
- Pixel dimensions =210μm x 210μm

$$= \underbrace{Q^4\beta^4 \frac{\sin^2(Q\alpha x)}{(\pi Q\alpha x)^2}}_{\text{zero order}} + \underbrace{Q^2\beta^4 \frac{\sin^4(\frac{\pi\alpha}{2}x)}{(\frac{\pi\alpha}{2}x)^2}}_{\text{low level background}} . \tag{12}$$

Figure 3 shows the power spectra of the randomised array described, and the corresponding regular array with $\beta = \alpha/2$, and demonstrates how successful the randomisation has been in removing the higher spectral orders.

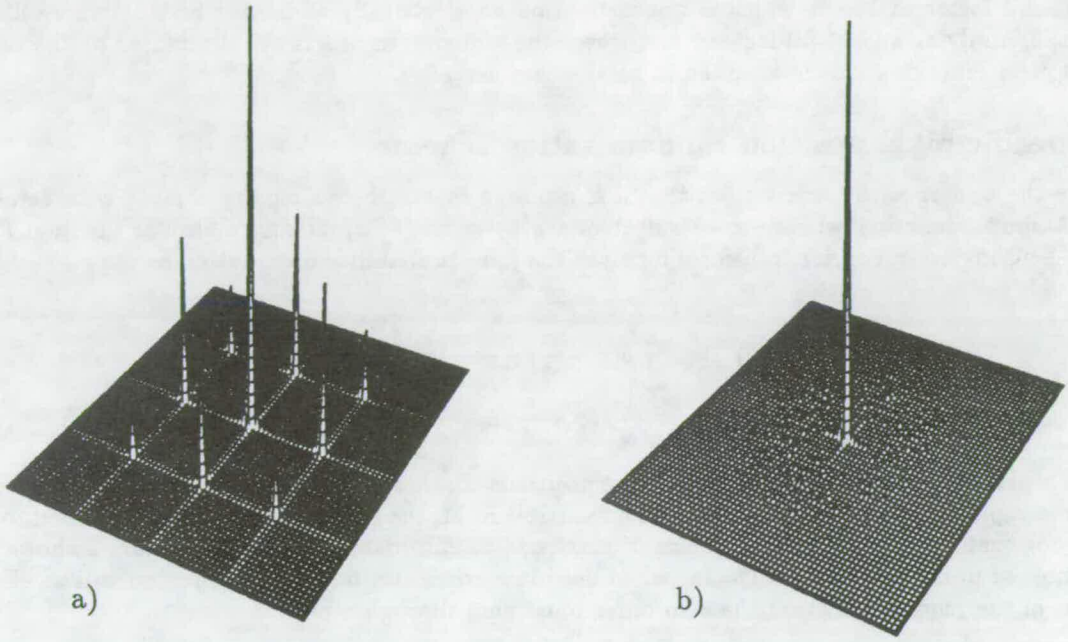

Figure 3: Power spectra of a) Regular pixel array b) Random pixel array

4 A RANDOMISED PIXEL SPATIAL LIGHT MODULATOR

A transmissive 16x16 pixel liquid crystal SLM using a multiplexed matrix addressed drive scheme was fabricated and is shown in figure 4, with all pixels "on", and displaying filter patterns used later.

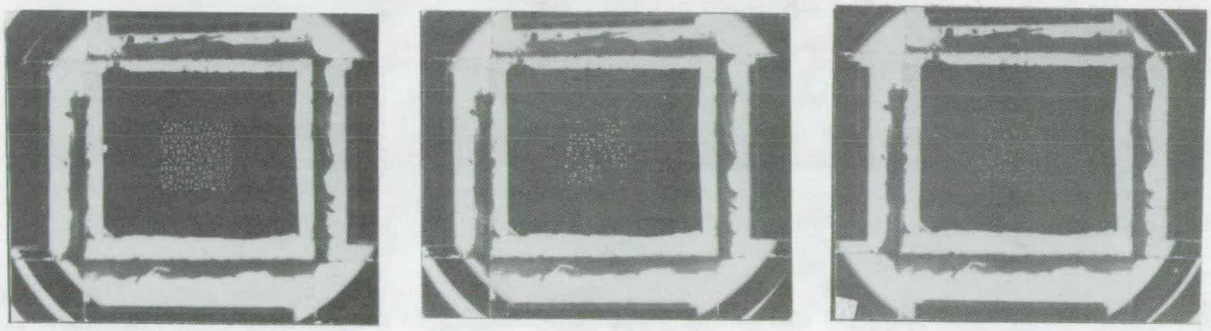

a)All pixels on. b)Filter pattern for "R". c)Filter pattern for "O".

Figure 4: SLM with all pixels on, and displaying filter patterns

This configuration was chosen because of its flexibility, allowing arrays of different shape, size, or position to be fabricated by simply changing a photomask. This mask is used to define the non-conducting regions between

3.2 Effect of the pixel function, $p(x,y)$

The pixel function also plays an important role in attenuating the spectral orders. For an amplitude modulator with square pixels of side β and transmittance of 0 outside the pixel and 1 inside, the pixel function Fourier transforms to

$$P(x,y) = \beta^2 \frac{\sin(\pi\beta x)}{\pi\beta x} \frac{\sin(\pi\beta y)}{\pi\beta y}. \tag{10}$$

If β was equal to α then $P(x,y)$ would eliminate all spectral orders other than the zero order by its zeroes– this implies a fill factor of 100 % which is impractical on an electrically addressed SLM. A more likely case is where $\beta = \alpha/2$, implying a 25 % fill factor,[4] and where the higher even orders are eliminated by the zeroes of the transformed pixel function. This is assumed to be the case hereafter.

3.3 A specific pixel position randomisation scheme

To eliminate the odd spectral orders from the SLM's power spectrum, we require a $\rho(...)$ with zeroes at their positions. A simple function which proves suitable is $\rho(x) = \cos(\frac{\pi\alpha}{2}x)$. Using a similar function for $\rho(y)$,i.e. $\rho(y) = cos(\frac{\pi\alpha}{2}y)$, inversely Fourier transforming gives the joint probability distribution as

$$\begin{aligned} P_{\phi_{n,m},\psi_{n,m}}(\phi,\psi) &= \frac{1}{4}\,\delta(\phi - \frac{\alpha}{4}, \psi - \frac{\alpha}{4}) + \frac{1}{4}\,\delta(\phi - \frac{\alpha}{4}, \psi + \frac{\alpha}{4}) \\ &\quad + \frac{1}{4}\,\delta(\phi + \frac{\alpha}{4}, \psi - \frac{\alpha}{4}) + \frac{1}{4}\,\delta(\phi + \frac{\alpha}{4}, \psi + \frac{\alpha}{4}). \end{aligned} \tag{11}$$

Thus each pixel can lie centred on one of four positions displaced from its regularly positioned centre. To prevent adjacent pixels overlapping, necessary in a realisible SLM, the size of the pixel is constrained to be $\leq \alpha/2$, which does not conflict with the earlier choice of pixel size to eliminate even orders. Figure 2 shows a member of the ensemble of possible arrays for Q=16: when used in a correlator only a 15x15 region will be effectively be used because of the requirement that the zero order must pass through a pixel.

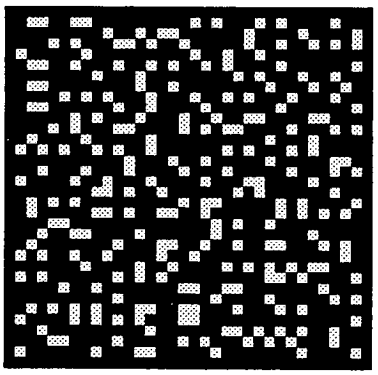

Figure 2: 16x16 random pixel array

Substituting the expression for $\rho(x)$ and setting $y = 0$ gives the form of the power spectrum along the x-axis as

$$\langle |\, t_2(x,0)\,|^2 \rangle = Q^2\beta^4 \frac{\sin^2(\frac{\alpha}{2}x)}{(\pi\alpha x/2)^2}[1 - \cos^2(\frac{\pi\alpha}{2}x) + \cos^2(\frac{\pi\alpha}{2}x)\frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)}]$$

3.1 Analysis

The two-dimensional transmission function for a Fourier plane filter with a randomised array of pixels, all "on", is

$$T_2(\xi,\eta) = p(\xi,\eta) * \sum_{n=-\frac{Q-1}{2}}^{\frac{Q-1}{2}} \sum_{m=\frac{Q-1}{2}}^{\frac{Q-1}{2}} \delta(\xi - n\alpha - \phi_{n,m}, \eta - m\alpha - \psi_{n,m}) \tag{4}$$

where $p(\xi,\eta)$ is the transmission function for a single pixel, and $\phi_{n,m}$ and $\psi_{n,m}$ are $\xi-$ and $\eta-$ components of the displacement of the centre of the pixel labelled by n and m from the position it would have occupied were it regularly positioned with pitch α in each direction. Q is chosen to be odd, to ensure that a pixel is centred on the origin.

Calculating the Fourier transform using specific values for the random displacements does not yield a useful general expression, so we calculate an ensemble average using all members of the set of possible sets of random displacements. As we detect intensity the averaging is performed over

$$\begin{aligned}\langle \mid t_2(x,y) \mid^2 \rangle &= \langle \mid P(x,y) \mid^2 \rangle \\ &\times \Big\langle \sum_{n=-\frac{Q-1}{2}}^{\frac{Q-1}{2}} \sum_{m=-\frac{Q-1}{2}}^{\frac{Q-1}{2}} \sum_{n'=-\frac{Q-1}{2}}^{\frac{Q-1}{2}} \sum_{m'=-\frac{Q-1}{2}}^{\frac{Q-1}{2}} \Big[e^{-i2\pi(n-n')\alpha x} e^{-i2\pi(m-m')\alpha x} \\ &\times e^{-i2\pi\phi_{n,m}x} e^{i2\pi\phi_{n',m'}x} e^{-i2\pi\psi_{n,m}y} e^{i2\pi\psi_{n',m'}y} \Big] \Big\rangle \,. \end{aligned} \tag{5}$$

This is not easily separable as the random displacements are specified for each pixel by labels relating to both the ξ and η directions: however the expression is simplified greatly by removing the terms where $n = n'$ and $m = m'$ from the summation, leaving only mutually independent random variables. This gives

$$\begin{aligned}\langle \mid t_2(x,y) \mid^2 \rangle &= \mid P(x,y) \mid^2 \times \Big[Q^2 + \Big(\frac{sin^2(\pi Q\alpha x)}{sin^2(\pi\alpha x)} \frac{sin^2(\pi Q\alpha y)}{sin^2(\pi\alpha y)} - Q^2 \Big) \\ &\times \langle e^{-i2\pi\phi_{n,m}x} \rangle \langle e^{i2\pi\phi_{n',m'}x} \rangle \langle e^{-i2\pi\psi_{n,m}y} \rangle \langle e^{i2\pi\psi_{n',m'}y} \rangle \Big] \,. \end{aligned} \tag{6}$$

Using

$$\langle e^{-i2\pi\phi_n x} \rangle = \int_{-\infty}^{\infty} P_{\phi_n}(\phi) e^{-i2\pi\phi x} \, d\phi \,, \tag{7}$$

where $P_{\phi_n}(\phi)$ is the real, positive probability of finding ϕ_n with value ϕ, and is identical to $P_{\phi_{n'}}(\phi)$, we write

$$\langle e^{-i2\pi\phi_n x} \rangle \langle e^{i2\pi\phi_{n'} x} \rangle = \mid \rho(x) \mid^2 \,, \tag{8}$$

where $\rho(x)$ is the Fourier transform of the probability distribution function, $P_{\phi_n}(\phi)$, to give a general expression for for a QxQ randomised pixel array:

$$\begin{aligned}\langle \mid t_2(x,y) \mid^2 \rangle &= \mid P(x,y) \mid^2 \times Q^2 \Big[1 - \mid \rho(x) \mid^2 \mid \rho(y) \mid^2 \\ &+ \frac{1}{Q^2} \frac{\sin^2(\pi Q\alpha x)}{\sin^2(\pi\alpha x)} \frac{\sin^2(\pi Q\alpha y)}{\sin^2(\pi\alpha y)} \mid \rho(x) \mid^2 \mid \rho(y) \mid^2 \Big] \,. \end{aligned} \tag{9}$$

Note how the $\mid \rho(...) \mid^2$ terms attenuate the spectral orders, redistributing their energy into a background illumination of magnitude of the order Q^2 less than that of the zero order, which falls to zero at the zero order.

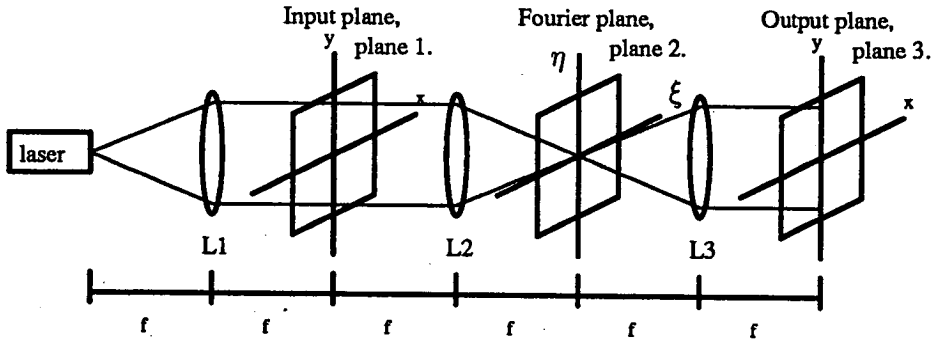

Figure 1: Fourier optical processor

$$t_2(x,y) = P(x,y)\frac{\sin(\pi Q\alpha x)}{\sin(\pi\alpha x)}\frac{\sin(\pi Q\alpha y)}{\sin(\pi\alpha y)} \tag{2}$$

where $P(x,y)$ is the Fourier transform of a single pixel. This function exhibits spectral orders of spatial frequency at

$$\left(x_n = \frac{m}{\alpha}, y_m = \frac{n}{\alpha}\right) \text{ for } m,n = 0,\pm1,\pm2,\ldots, \tag{3}$$

which will result in replicas when convolved with an input object.

Replicated correlations become problematic when the filter SLM has a lower space bandwidth product (SBP, equal to the number of pixels across the SLM, Q) than that of the input, and is best considered using a specific example. We consider the case of input to the processor via a binary SLM of SBP Q_1, SLM1, and filtering using an SLM of SBP Q_2, SLM2.

To prevent the replicas overlapping, the Fourier transform of the input must be sampled at a frequency greater than or equal to the Nyquist frequency. If the pixel pitch of SLM2 is α this will restrict the maximum size of the object to be $\lambda f/\alpha$. If this object then occupies Q_1xQ_1 pixels on SLM1 then the lowest spatial period that can be represented is $(2/Q_1)(\lambda f/\alpha)$. This in turn will transform to a point $Q_1\alpha/2$ distant from the position of the zero spatial frequency component of the transform, which passes through the centre of the central pixel of SLM2, requiring that SLM2's SBP be at least as large as that of the target in order to represent accurately its transform or complex conjugate, and also ensure that the replicas lie outwith the region of interest in the output plane, which will be restricted to the area defined by the geometrical image of the perimeter of SLM1. If however the target or targets each occupy a 16x16 pixel region in a 256x256 input SLM for example, then a 16x16 pixel SLM in the Fourier plane, though adequate to represent the target's Fourier transform, will produce replicas lying within the region of interest in the output plane which could give rise to peaks which may be misinterpreted as signifying correlation.

As it is advantageous to use a low resolution SLM in the Fourier plane because of fast frame rates, compact size, modest memory size for the filter bank,[2] and lower cost, and because many correlation tasks look for a small target in a large input (e.g. tracking aircraft), we consider how the spectral orders which give rise to the replicas may be reduced or removed, by randomisation of Fourier plane pixel positions.

3 RANDOMISATION OF FOURIER PLANE FILTER PIXEL POSITIONS

The analysis of this section is carried out in more detail elsewhere,[7] but is briefly summarised here. We wish to determine the form of the Fourier transform of a randomised pixel array.

Reduction of aliasing in correlation using a pixelated Spatial Light Modulator

S. Heddle, D.G. Vass & R.M. Sillitto

Applied Optics Group, Department of Physics,
The University of Edinburgh
James Clerk Maxwell Building
King's Buildings, Mayfield Road
Edinburgh EH9 3JZ, Scotland

ABSTRACT

Replication in the output plane of an optical correlator, due to pixelation of the Fourier plane filter, can lead to false correlation signals. This paper suggests randomisation of the pixel positions of the Fourier plane filter as a solution, and demonstrates its effectiveness through computational simulations and optical correlation experiments using a custom made SLM.

1 INTRODUCTION

The use of spatial light modulators (SLMs) as input devices and Fourier plane filters in optical processors has received much attention. In particular the use of SLMs in the Fourier plane to implement pattern recognition or correlation has attracted much effort, with the development of a wide range of modulation schemes and algorithms,[1] often binary in nature, in line with the capabilities of current SLMs. If the filter SLM is pixelated, as is usual to enable convenient writing of the filter pattern, then the regular filter structure will result in the correlator output being replicated at intervals across the image plane. In the correlator the region of interest where correlator signals are to be detected is defined by the geometrical image of the periphery of the input scene– clearly if the replicas fall within this region of interest false correlation signals may result.[2,3,5] The following section shows briefly how and why this may occur. The remainder of the paper proposes a solution based on randomising the pixel positions of the Fourier plane filter, and presents comparitive simulations for both regular and randomised pixel Amplitude Encoded Binary Phase Only Filters (AEBPOFs), for a carefully chosen pixel position randomisation scheme. A randomly pixelated spatial light modulator is introduced and actual optical correlation results presented.

2 ORIGIN OF REPLICATION PROBLEM

We consider the correlation operation with regard to the optical processor and coordinates shown in figure 1.

Defining the transmittances of the input plane and Fourier plane filters by $t_1(x,y)$ and $T_2(\xi,\eta)$ respectively, the output can be written as

$$u_3(x,y) = t_1(-x,-y) * t_2(x,y) \tag{1}$$

where $t_2(x,y)$ is the Fourier transform of $T_2(\xi,\eta)$ and $*$ denotes convolution. The ξ and η are scaled spatial frequency coordinates equal to $u/\lambda f$ and $v/\lambda f$ respectively, where u and v are spatial frequency coordinates and f is the focal length of the transforming lenses L2 and L3. Thus the output is a coordinate reversed version of the input convolved with the Fourier transform of the filter, $t_2(x,y)$. If the filter is a regular Q by Q array of square pixels, all "on", of side β on a pitch α[8] then

0-8194-0945-6/92/$

PROCEEDINGS REPRINT

SPIE—The International Society for Optical Engineering

Reprinted from

Optical Information Processing Systems and Architectures IV

20–21 July 1992
San Diego, California

Volume 1772

Input	Randomised filter for...		
	N	O	R
N	176	101	142
O	124	249	135
R	115	100	165

Table 4: Table of optical correlation results for multiple letter input

7 CONCLUSION

The problem of false correlations due to replicated outputs has been noted and a solution proposed– randomisation of pixel positions in the Fourier plane filter. An analysis of randomised pixel arrays has been carried out, and a particular pixel position probability distribution chosen which removes the higher spectral orders in the SLM's power spectrum which cause replication. Simulations have been performed which show that a Fourier plane filter based on such a randomised array offers improved correlation performance when the space bandwidth product of the input exceeds that of the Fourier plane filter. This suggests that a Fourier plane SLM with appropriately randomised pixel positions need only have an SBP equal to that of its target, rather than of the entire input scene, to perform successful correlation. A binary 16x16 randomised pixel SLM has been constructed, and demonstrated optical correlation performance in good agreement with the computational simulations.

8 ACKNOWLEDGEMENTS

The SLM was produced with technical assistance from Eric Davidson, Andrew Garrie, and the photomasks and input transparencies prepared with the assistance of Peter Tuffy. We wish to acknowledge Will Hossack and Duncan Potter for many helpful discussions, and particularly Ian Chisholm whose software was used extensively in the simulations. Steven Heddle wishes to thank Tom Stevenson, Alan Gundlach and Tony Snell of the University of Edinburgh's Electrical Engineering Department for much help with lithography and the use of their department's facilities. Steven Heddle acknowledges the financial assistance of the SERC, and is greatly indebted to the Orkney Islands Council, Elf Enterprise Caledonia Ltd, Institute of Physics Optical Group, his Physics Department, the IOP C.R.Barber Trust Fund, and Digital Equipment Scotland, for conference funding.

9 REFERENCES

1] B.J.Pernick. Phase-only and binary phase-only spatial filters for optical correlators:a survey. *Optics & Laser Technology*, 23:273, 1991.

2] G.Gheen, E.Washwell and D.Armitage. The effect of filter pixelation on optical correlation. *OSA Technical Digest on Spatial Light Modulators and Applications*, 14:161, 1990.

3] H.F.Yau, Y.P.Tong, C.M.Lee, C.C.Chang and M.W.Chang. Pattern recognition with circular harmonic filtering and phase-only circular harmonic filtering using a Vander Lugt filter. *Jpn. J. Appl. Phys.*, 31:49, 1992.

4] I.Underwood, P.H.Willson, R.M.Sillitto, and D.G.Vass. Real time image processing: concepts and technologies. *Proc. Soc. Photo-Opt. Instrum. Eng.*, 860:2, 1986.

5] J.A.Davis, E.A.Merrill and D.M.Cottrell. Effects of sampling and binarization in the output of the joint Fourier transform correlator. *Opt. Eng.*, 29:1094, 1990.

6] M.A.Flavin and J.L.Horner. Amplitude encoded phase-only filters. *Applied Optics*, 28:1692, 1989.

7] S.Heddle and R.M.Sillitto. Reduction and removal of replication in an optical processor by randomisation of pixel positions in the Fourier plane filter. Submitted to *Journal of Modern Optics.*

8] W.J.Hossack, D.G.Vass and I.Underwood. Fourier processing with binary spatial light modulators. *Proc. Soc. Photo-Opt. Instrum. Eng.*, 1564:697, 1991.

Specification of an optical correlator for road sign recognition

S.Heddle, W.J.Hossack, D.G.Vass

Applied Optics Group, Department of Physics, University of Edinburgh,
The Kings Buildings, Mayfield Road, Edinburgh, Scotland. EH9 3JZ
B: heddle@castle.ed.ac.uk

Optical correlation of road signs is considered, with particular reference to space-bandwidth product (SBP), input to the correlator, and efficient readout in the correlation plane.

Optical correlation of EC standard signs[1] from a moving vehicle demands a highly specified correlator. The sign may appear in a wide range of positions, in sizes varying over a range of 400%, and distorted from the idealised square-on pose.

The wide field of view of the input suggests a high input SBP ($\sim$ 1500), achievable using an optically addressed SLM. With an input frame rate of $\gtrsim$50Hz we show that the scale invariance necessary in the filter may be reduced to <2%, making such an input practical for use with a filter generated by an optimisation process[2] to account for distortions. Conventionally a pixellated Fourier plane filter should have the same SBP as the entire input to prevent aliasing effects through undersampling– impractical here both in terms of SBP and required frame rate. We present a novel phase only filter SLM with pixel positions randomly selected from a carefully defined set of allowed positions, which need be matched only to the much lower SBP of the target signs ($\sim$ 50), without introducing aliasing of the correlation signals[3,4], and allowing rapid scanning of the filter library. The high SBP of the input is retained in the output– correlation peaks can be detected rapidly by imaging the output onto crossed diagonal linear arrays of photodetectors.

1.The Traffic Signs Manual, Chapter 1 (1982) (HMSO, London).
2.T.D.Wilkinson, D.C.O'Brien and R.J.Mears, submitted to Opt. Lett., Feb. 1993.
3.S.Heddle and R.M.Sillitto, J. Mod. Opt., accepted 14th Oct. 1992.
4.S.Heddle, D.G.Vass and R.M.Sillitto, Proc. SPIE **1772**, 116 (1992).

1 The basic problem

The problem that we are required to address is that of road sign recognition, using an optical correlator which is capable of being mounted in a car. The task is apparently well-defined, but we must still ask some fundamental questions before any progress can be made, *viz.*

- What are we looking for?
- Where might we expect to find it?
- How hard do you have to look?
- Is it alone?
- What prior knowledge about the target do we have?

We know what we are looking for– road signs– but what size are they, what shape are they, are they standardised at all? We can make progress on these additional questions because we do have the prior knowledge that the signs are standardised, they fall into four fundamental shape classes (equilateral triangle, inverted equilateral triangle, circle and octogon), and range in size from 450 - 1500 mm in height, dependent on the road[1]. We also know that the signs need not be alone, with two signs often appearing on the same post, or identical signs either side of a carriageway.

The remaining questions of where and how hard to look are of crucial importance when specifying the input stage of the correlator. This is considered in the next section.

2 The input

Figure 1 shows a schematic of the typical geometry of the problem. Some assumptions have necessarily been made, such as assuming signs may be on either side of the road (valid if the diagram is taken as depicting a dual carriageway with signs in the central reservation), and that each lane of the carriageway is 4m wide.

Placement of the sign is less well defined than the sign's shape and dimensions. Guidelines exist in the UK which suggest that the sign should be mounted with its lowest edge between 900 and 1500 mm above the height of the highest part of the carriageway, and between 450 and 1200 mm from the edge of the carriageway– in practice signs will be found mounted on overhead bridges, and in any practible place in an urban setting.

Referring to the figure, and taking $\theta_{1_{max}} = \theta_{2_{max}}$ gives an estimate of the width of the field of view in which we would expect to find signs as approximately 13m, of which the sign may occupy as little as 0.45m.

We must now estimate how hard we need to look for the sign, i.e. how few pixels we can use to characterise the signs while still retaining enough information to discriminate between them. A useful guideline for this can be

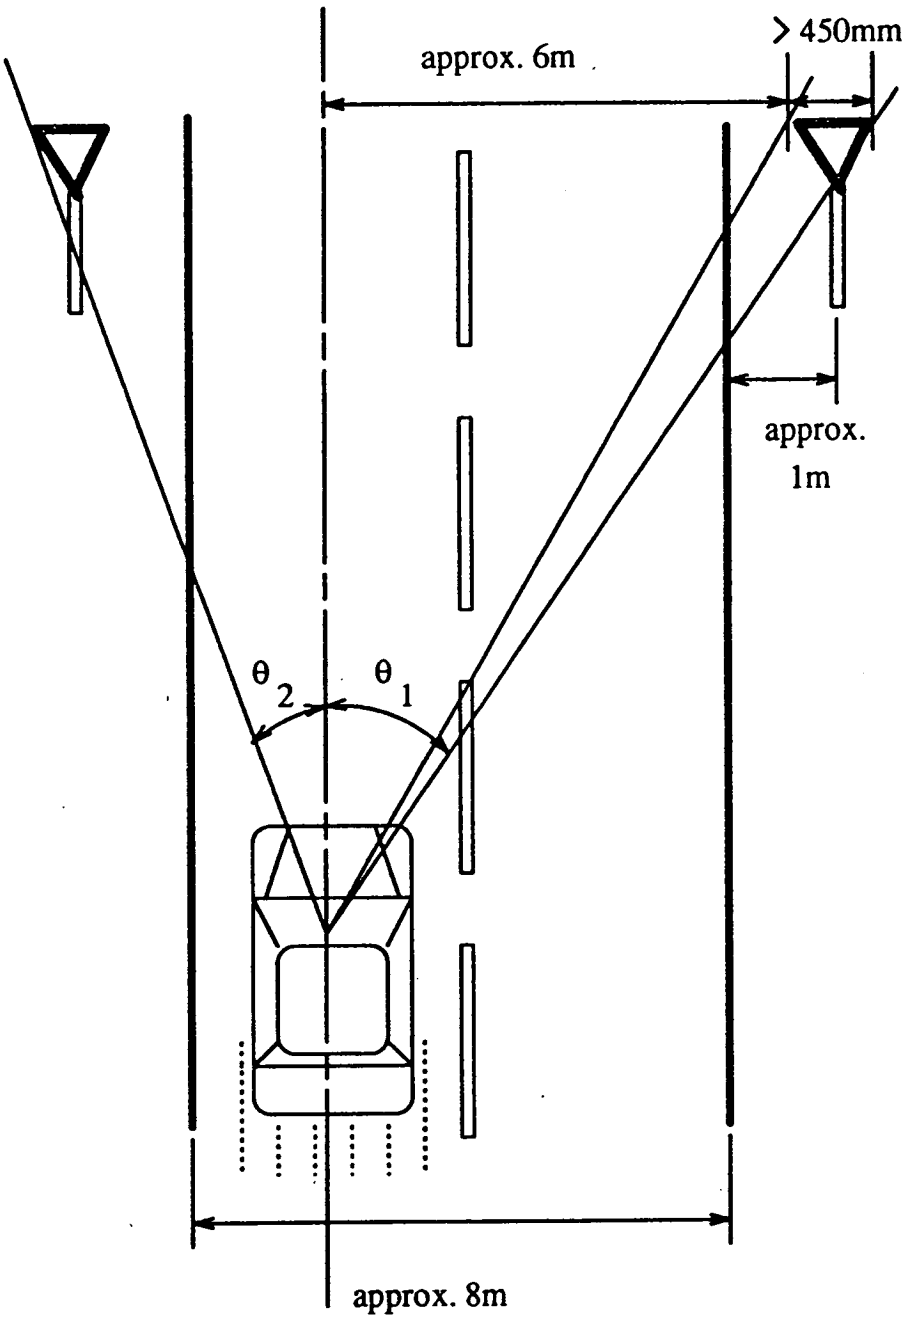

Figure 1: Geometry of the road sign recognition problem

found by considering reconfigurable signs as manufactured by Schott[2]. These signs are made up of resolution elements consisting of bundles of optical fibres, and are typically 50 resolution limits wide. This estimate has also been validated by scanning pictures of road signs into a computer and reducing the resolution of the picture using the display package **xv**, until the sign became hard to identify visually. Using a 'Give Way' as a reference, 50 by 50 pixels was indeed sufficient to allow the sign to be identified visually without difficulty.

Using this level of resolution for the sign further implies that the entire field of view requires (13/0.45 x 50)= 1444 pixels across its width, and approximately half that over its height. This is a formidable space bandwidth product, and certainly beyond the capabilities of current electrically addressed SLMs. Nevertheless such resolution is available on optically addressed SLMs, and in CCD arrays, so several posssible input architectures exist:

1. Observe the entire field of view with a high resolution CCD camera and display a part of it (e.g a 256 by 256 window) on an electrically addressed SLM, with this window of interest determined by a low level digital pattern recognition algorithm which detects the presence of a sign.

2. Observe the entire field of view with a high resolution CCD camera and display all of it on an electrically addressed SLM as a series of overlapping segments (see Figure 2).

3. Have a comparatively low resolution CCD camera lock onto a sign and track it, displaying the output on an electrically addressed SLM with the same space-bandwidth product as the camera.

4. Image the entire field of view onto an optically addressed SLM (OASLM) at the input to the correlator.

Option 1 is certainly feasible, and may well be attractive if some level of preprocessing of the input is already necessary, such as edge-enhancement to improve discrimination and correlation peak sharpness, and reduce the dynamic range of the input. However, by doing partial target recognition digitally makes it seem logical to complete the task digitally, by template matching or some other means. The windowing operation must also be capable of following multiple windows in the case of several signs of the same size but differing placements appearing in the input.

Option 2 removes the need for any digital pattern recognition, but will reduce the maximum frame rate for the overall input scene as the set of input segments are presented sequentially to the electrically addressed SLM. The importance of the overall input frame rate when considering tolerance of scale variance in the signs, and scale differences between successive frames is considered later– at this stage we assume that the sign is presented for correlation occupying an area 50 pixels wide. An overlap with adjacent segments of at least 25 pixels is thus necessary to ensure that the sign is entirely contained by a segment. This results in the need for 18 segments, as shown in Figure 2.

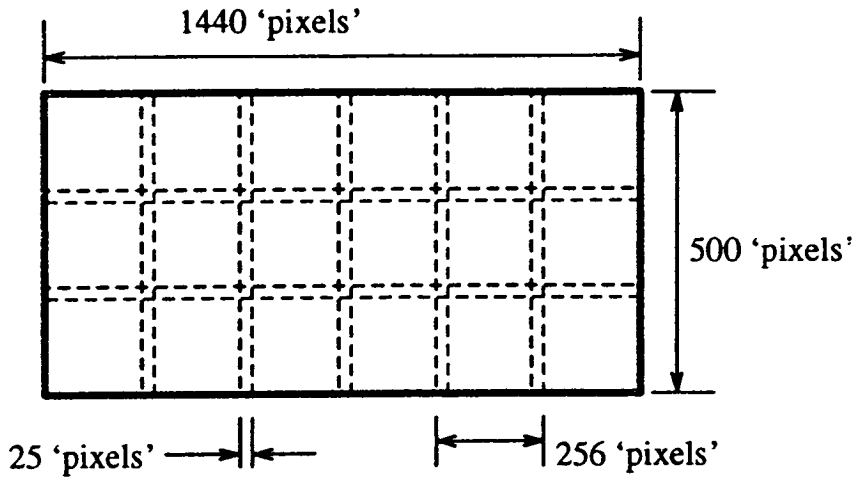

Figure 2: Input divided into overlapping segments

Option 3 can be dismissed as less practical than the others as it requires a significant amount of digital processing and mechanical design to arrive at a system which would be capable of even following one sign at a time.

The fourth option is appealing as it comes closest to an 'all optical' implementation of the input. As it observes the entire field of view there is no need for complicated tracking or windowing, and any number of signs can be presented at the input of the correlator simultaneously. Edge enhancement of the input is also possible using a non-linear response of the OASLM, with the advantages described earlier, though this depends critically on the intensity of the light incident on the write side of the OASLM. Some means of exposure control is thus desirable, which could be achieved through a simple shuttering mechanism. As the correlation process will probably be based on comparing a set of references with the input, it will be necessary in any case to sample and hold the input for the length of time it takes to cycle through the references. The huge space-bandwidth product of the input normally would require an electrically addressed filter SLM of equal space-bandwidth product when using a matched filter correlator, or twice the space-bandwidth product if a joint transform correlator is used. Using a filter SLM which implements a scheme whereby the active pixels are randomly positioned according to a carefully chosen probability distribution function, it is possible to design a matched filter which is matched to the target of interest rather than the entire input scene, reulting in a filter which in this case needs a space bandwith product an order of magnitude less than that of the input.

This architecture will be pursued here. First, the optical properties of the OASLM and the input imaging system must be considered.

3 Using an OASLM as the input device

This is considered with respect to the GEC-Marconi twisted nematic OASLM (model Y-35-3222), which I believe is no longer available– this should not matter as it is reasonable to assume that newer OASLMs should be better specified. For example ferroelectric OASLMs will offer higher speed, and the sample and hold operation should be trivial using the LC's bistability.

The GEC-Marconi SLM has the following basic specifications:

Specification	
Writing wavelength	620 nm[1]
Read wavelength	514 nm
Open aperture	40 x 40 mm
Resolution at 50% MTF	35 lp/mm
Read power handling	1 $mWcm^{-2}$
Sensitivity 10% (Sat)	40 μWcm^{-2}
Sensitivity 90% (Sat)	150 μWcm^{-2}
Switching time	<20 ms
Electrical drive	approx. 5 V square wave at 1000 Hz

The SLM's sensitivity to the entire visible range is necessary, and direct imaging of the input scene seems possible– intensity measurements made on a dull day in Edinburgh gave a value of 700 μWcm^{-2} when the detector was pointed at the sky, and 42 μWcm^{-2} when pointed at a grey office wall. This suggests that the sensitivity is at least in the right ball park, though the situation will be influenced by the magnification and aperture of the imaging system. Some measurements of light intensity falling on the focal plane of a camera will be made to gain a more accurate idea of the required sensitivity.

Using the resolution specification suggests that an input scene 1440 pixels wide (equivalent to 720 line pairs) should be imaged onto an area of the SLM at least 20.6 mm wide, if the MTF is to remain above 50%.

4 The imaging system

The imaging system should be free from distortions to simplify the correlation task, suggesting proprietry corrected optical systems such as photographic or CCTV camera lenses to reduce expense. CCTV lenses are available for image sizes up to 12.7 mm wide, beyond which vignetting and field curvature will become problematic. Lenses for 35 mm cameras are appropriate, with an image size of 24 mm × 36 mm. Therefore the approach is to either take 50% MTF as a benchmark and choose a lens which will give an image 20.6 mm wide, or to choose a lens of a common focal length giving an image more than 20.6 mmm across.

The imaging problem requires more consideration. If we take the constraint that the driver should not have to look more than 10° to the side from the direction of travel to see a sign (as suggested in UK traffic signing guidelines),

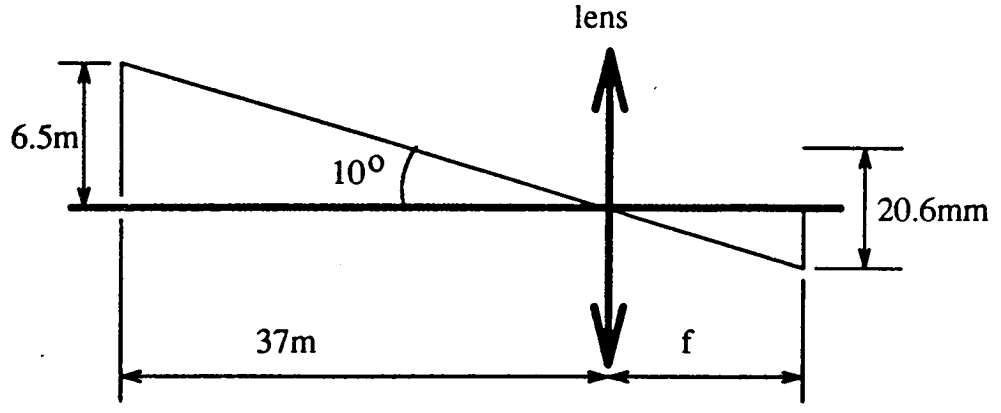

Figure 3: Determination of focal length of input lens

setting $\theta_{1_{max}} = \theta_{2_{max}} = 10°$, and look at a sign 450 mm wide 6 m to the left (or right) of the centre of the car on the edge of the field of view, then this sign is at a distance of $x = \frac{6.5}{\tan 10°} \simeq 37$ m. To reduce the number and complexity of the filters, the imaging system should be capable of forming sharp images of the larger versions of the sign, which occur at distances giving an image the same size as that given by the 450 mm sign at 37 m. In effect the depth of focus should be from 37 m to $(\frac{1500}{450} \times 37) \simeq 123$ m.

The minimum focal length of the lens is apparent from Figure 3 and is 58 mm. Taking the front focal plane to be at 80 m, the midpoint of the range of distances which are to be in focus, we can determine the $F_{No.}$ from the Strehl limit of the lens. This gives

$$\begin{aligned} F_{No.} &= \sqrt{\frac{\Delta z_0}{2\left(\frac{z_0}{f}\right)^2 \lambda}} \\ &= 4.5 \end{aligned} \tag{1}$$

where $f = 58$ mm, the depth of focus extends from $f - \Delta z_0 = 37$ m to $f + \Delta z_0 = 123$ m, $z_0 = 80$ m, and λ is taken to be 550 nm.

A further calculation allows us to estimate the resolution limit of the lens i.e. the smallest angle that can be resolved is

$$\begin{aligned} \Delta\theta &= \frac{1.22\lambda F_{No.}}{f} \\ &= 5.2 \times 10^{-5} \text{radians.} \end{aligned} \tag{2}$$

which is smaller than the required $\Delta\theta_{max} = 2.4 \times 10^{-4}$ radians (see Figure 4). Thus such an imaging system is feasible, and the comparatively wide maximum aperture corresponding to $F_{No.} = 4$ offers the possibility that stopping down the lens may be used for exposure control in addition to, or in place of some shutter system. Stopping down the aperture may prove undesirable if edge enhancement is used, as a greater depth of field will lead to more sharp edges cluttering the input.

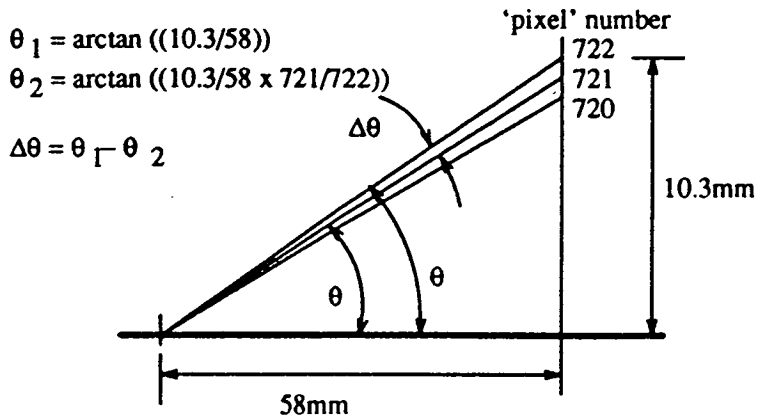

Figure 4: Determination of $\Delta\theta_{max}$

5 The requirement to temporally sample the input

It has been suggested that it is necessary to sample the input– this is to enable the series of reference objects (or filters generated from the refences) to be compared with the same input scene, giving a meaning full set of correlation responses. Otherwise the case may arise where a false correlation from a sign that is of the appropriate scale at some instant may be larger than the 'true' correlation from a subsequent sign which is mismatched in scale with the reference or filter.

The sampling rate will be decided by two competing factors:

- the amount of scale invariance which the correlator can cope with
- the size of the bank of references/filters.

Consider, for example, a correlator which can cope with a variation in scale of ±2% (this corresponds to a mismatch of ±1 pixel for the idealised 50 pixel wide input sign), mounted in a car travelling at 60mph (approx. 100kph or 26ms^{-1}). A 450mm sign that is 50 pixels wide in the input will be at 37m: if this is imaged onto the OASLM as a frame, the subsequent frame must be at $0.98 \times 37 = 36.26$m, corresponding to 29 ms later *or less* in order for the scale difference between frames to be less than 2%.

This implies an input frame rate of 34.5 Hz, which is certainly not inconcievable, with a filter frame rate of (34.5 × number of references or filters)[2]. This latter figure is more problematic– if we take a set of 100 signs and simply use the correlator to compare each sign with the input then a filter frame rate of 3.45 kHz is implied. This is certainly pushing the technology, and it may well be advantageous to perform the correlation heirarchically, searching for a sign's shape, then comparing the possible signs of that shape, or possible combinations of numerals or features on signs once one numeral or feature has been

[2]Filter frame rate here should be taken to mean the rate at which the reference patterns or Fourier plane filters appear– references at the input in the case of the JTC, or actual filters in the case of the matched filter correlator

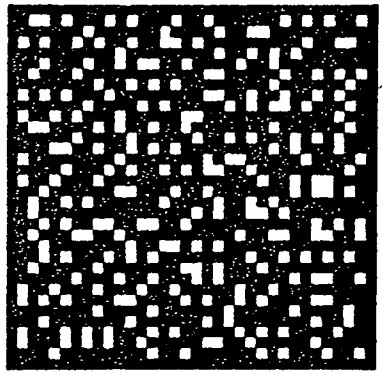

Figure 5: Array of randomly positioned pixels for non-aliasing filtering

identified etc.. Such an approach is possible for both matched filter and joint transform correlators, and should reduce the filter frame rate by a third.

A further way in which the filter frame rate could be reduced is to build greater scale variance into the filter, reducing the necessary input frame rate and thus the filter frame rate. This is certainly possible with the matched filter correlator, as techniques exist which incorporate various types of distortion invariance into the filter, such as simulated annealing[3]. For the JTC more invariance requires a larger bank of input images representing the distorted versions of the reference, so the filter frame rate effectively remains constant. The JTC also requires that the SLM in the filter plane have twice the SBP of the input, thus requiring either a further (fast) OASLM, or a fast EASLM used in conjunction with a fast CCD acting as a square law detector, with a windowing scheme at the input. The latter EASLM, operating at frame rates of the order of 1 kHz with a SBP of the order of 512, will be extremely hard to realise without some sophisticated data management.

The remainder of this consideration is restricted to a matched filter based correlator, using lower SBP filters designed to avoid aliasing problems.

6 The filter SLM

Conventionally the filter SLM has to have at least the same space-bandwidth product as the input scene[4], or aliasing of the correlation output in the correlation plane may lead to misleading correlation signals. Recently we have demonstrated that the aliasing can be suppressed by means of randomising the pixel positions according to a carefully chosen probability distribution function[5] allowing the constraint on the filters SBP to be relaxed to being at least that of the target, rather than of the larger input scene. This has been demonstrated optically using custom made 16×16 filter SLMs to implement Amplitude Encoded Binary Phase Only Filters (AEBPOFs)[6] and BPOFs[7].

The random pixel arrays used thus far have been as shown in Figure 5, with each pixel $< \alpha/2$ square placed with equal probability in one of four positions in an $\alpha \times \alpha$ cell defined by the underlying pixel pitch, α. A paper is currently being prepared which shows that the aliasing can be suppressed even more efficiently

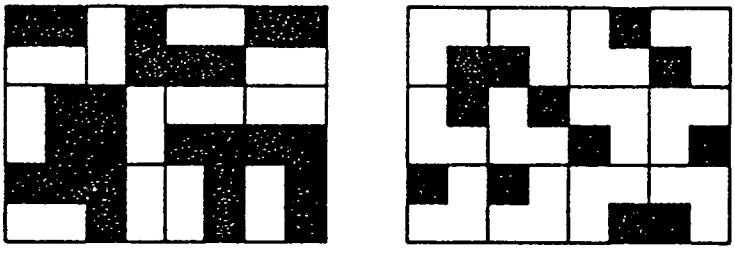

2 from 4 allowed positions. 3 from 4 allowed positions

Figure 6: Non-aliasing randomised pixel positioning schemes

by using similar randomisation schemes with the pixels occupying two and/or three of the four allowed positions (see Figure 6), with a resulting increase in the optical efficiency of the filter, and that intensity weighting of the spatial frequency components incident on the Fourier plane filter can be carried out by choice of either one, two or three transmitting regions per cell.

Two possible implementations of these randomised pixel schemes exist. The first is to build a custom SLM with pixels sized, positioned and shaped to conform to the chosen randomisation scheme, such as the 16×16 transmissive SLMs described above, or using planarised silicon backplane devices with the pixel defined on a second (or third) metal layer.

The alternative is to take an existing device with pixels spaced on a regular grid, and regard its pixels as subpixels of a randomly pixellated SLM– i.e. each 'pixel' of the randomly pixellated SLM is made up of a 2×2 grid of subpixels, with one, two or three of the subpixels permanently 'off', and the remainder switchable between all 'off' or all 'on'. This appears wasteful of the SBP of the SLM, as it is apparently halved– realise though that the resulting filter may be used to detect objects which appear in input scenes with a SBP more than that of the filter itself. More importantly it allows the the actual randomisation scheme itself to be freed as a parameter for optimisation with respect to the correlation response, by simulated annealing or a genetic algorithm. This is being investigated at Edinburgh by Ian Chisholm and Steven Heddle.

A hindrance of the second, non-custom, option is that the fill factor of the pixel cannot be optimised to suppress the even orders in the transform of the Fourier plane filter, with which the correlation output is convolved. In the case of a 176×176 HPSLM2 the resulting second order replicas' intensities would be a significant fraction of that of the zero order[3]: a solution to this problem would be possible by use of a custom made planarised HPSLM2, with a randomised pixel pattern defined on the top whose pixel size can be optimised to suppress

[3]Taking the pixel pitch as 30μm and pixel sizes as 30 × 14.5μm, second order replicas with intensities of ($\mathrm{sinc}^2(14.5/30) = 0.43$) times that of the zero order may result

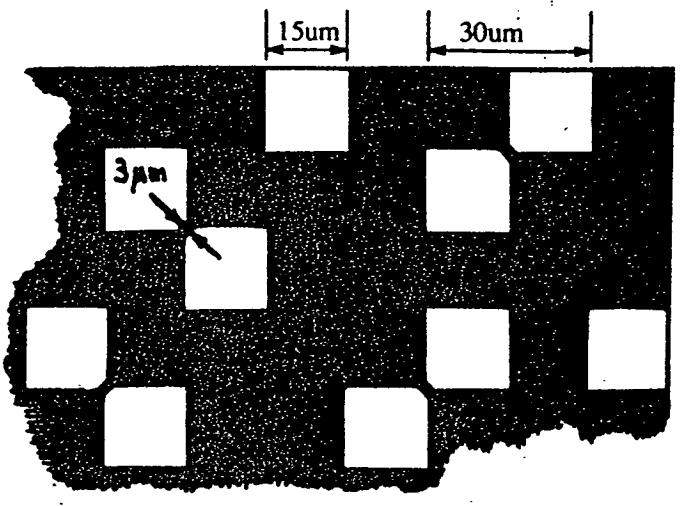

3um gap specified by design rules

Figure 7: Random array for silicon backplane with optimum fill factor

the even orders, subject to the presence of neighbouring pixels. Such a pattern is shown schematically in Figure 7.

This would give us a 176×176 pixel array, which is overkill in terms of SBP, given that speed of addressing is an important factor. Nevertheless, by restricting attention to a 50×50 region of this hypothetical SLM, we can proceed to specify an optical correlator and get an indication of its possible size.

7 Dimensions of the optical correlator

In the classical matched filter correlator, the focal length of the Fourier transforming lenses is specified by the wavelength of the illumination, the pixel pitch of the Fourier plane filter and the width of the input scene, according to

$$f = \frac{x_f x_0}{\lambda} \times \frac{50}{1440} \quad (3)$$

$$= \frac{30 \times 10^{-6} \times 20.6 \times 10^{-3} \times 50}{514 \times 10^{-9} \times 1440} \quad (4)$$

$$= 4.17\text{cm} \quad (5)$$

where x_f is the pixel pitch, x_0 is the input plane width, λ= 514.5nm, and the 50/1440 factor arises due to the matching to the filter of the target rather than the entire input scene.

Using this focal length a correlator of the dimensions shown in Figure 8 can be constructed. The figure shows the use of $F2$ lenses– this should be more than adequate. A smaller correlator should be possible using the geometry described by Cottrell et al.[8], which uses additional negative lenses to compact the system with a positive lens to compensate for phase curvature of the Fourier plane.

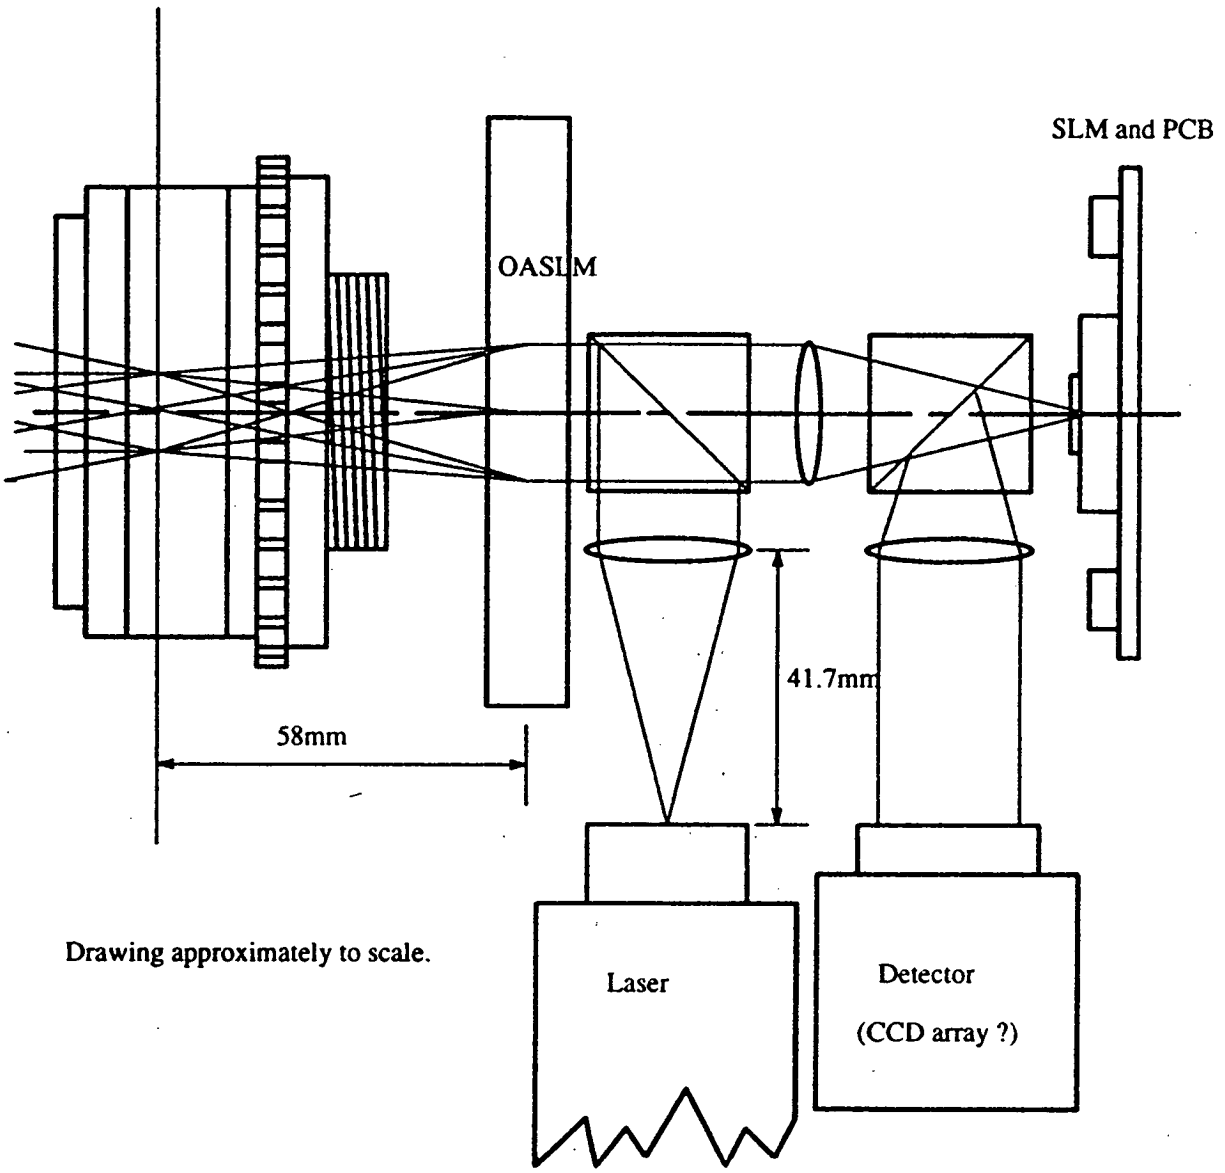

Figure 8: Geometry of optical correlator

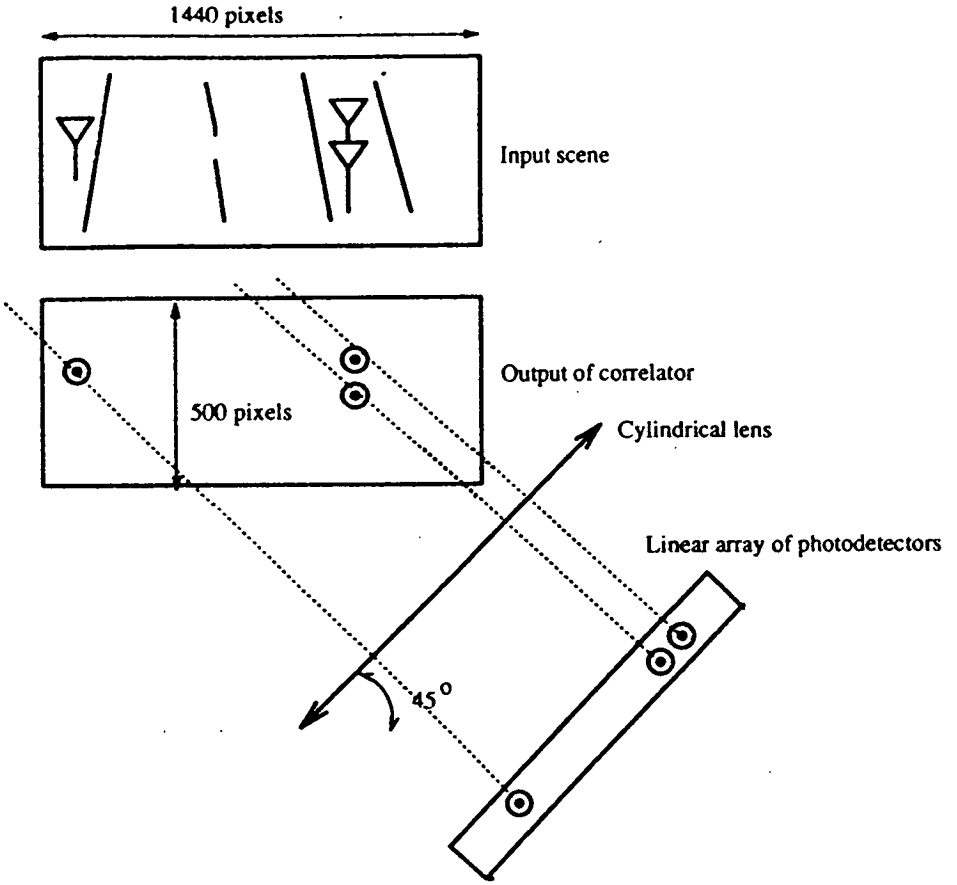

Figure 9: Proposed fast readout system

8 Readout

The system shown in Figure 8 has a CCD camera at the output– this is probably not practical, as the SBP of the output is the same as that of the input, and has to be read out at the frame rate of the filter SLM. A potential solution could be provided by using a fast readout linear array of photodetectors, onto which the correlation plane could be imaged using a cylindrical lens. As signs can be found vertically above each other on the same signpost, and at the same height as each other on either side of the road, the linear array should be at 45^0 to the horizontal so that each sign can be resolved– this is shown schematically in Figure 9. Thus each pixel on the linear array registers an intensity which is the modulus of the coherent sum along a slice across the correlation plane at 45^0– if the input is edge enhanced so that cross-correlations are much smaller than autocorrelations, and correlation peaks are narrow, then the correlation peaks should show unambiguously.

If this is a practical solution at all, it can be developed by imaging the correlation plane onto two orthogonal linear arrays at 45^0 to the horizontal, which offers a degree of redundancy in the detection operation and allows the positions of the correlations in the output plane to be calculated.

9 Comments and Conclusions

Some practical issues relating to optical correlator design with particular reference to the task of road sign recognition have been addressed, and a correlator system described. This system follows the basic Vander Lugt classical matched

filter architecture, to reduce the demands on the Fourier plane filter in terms of frame rate, space bandwidth product and distortion invariance– with this architecture some trade-off between frame rate and distortion invariance is possible, which is not apparent with the JTC. Using an optically addressed SLM to input the entire field of view removes the need for complicated windowing systems, but would normally demand a Fourier plane SLM with the same (or greater) SBP– a low resolution filter avoiding the aliasing problems is proposed, with the attendant advantages of reduced complexity and higher frame rate. The SBP of the input is retained in the output, leading to problems reading out the huge amount of data involved– imaging the output plane onto a linear array of photodiodes, placed at 45^0 to the horizontal to spatially resolve signs at the same height or on the same signpost, is suggested as a possible solution. Physically the resulting correlator is very compact.

Many questions remain:

- Which filtering scheme to use?
- How much distortion of the input signs is likely/tolerable?
- How to cope with a high dynamic range in the input (e.g. bright sun with the sign in the shade)?
- How to cope with rapidly changing light levels, or non uniformly illuminated targets?
- Is edge enhancement of the input necessary? (Tentative answer: yes, to reduce the dynamic range of the input, reduce cross correlation, and sharpen correlation peaks)
- What is the best way to readout the data?
- How to cope with multiple targets– can the autocorrelations be normalised?

Undoubtedly other important considerations have been missed– it is hoped however that this paper may be useful in future consideration of the general problem.

References

[1] *Traffic Signs Manual*, Chapter 1 (1982), Chapter 3 (1986), Chapter 4 (1986). HMSO, London. Chapters are published separately.

[2] *Schott information*, No.4/1985 English. Schott Glaswerke, Mainz, Germany.

[3] *Scale invariant binary phase-only matched filter using a ferroelectric liquid crystal spatial light modulator*, T.D. Wilkinson, D.C. O'Brien and R.J. Mears, submitted to Opt. Lett., Feb. 1993.

[4] *The effect of filter pixelation on optical correlation*, G. Gheen, E. Washwell and D. Armitage, OSA Technical Digest on Spatial Light Modulators and Applications **14**, p161 (1990).

[5] *Reduction and Removal of replication in an optical processor by randomisation of pixel positions in the Fourier plane filter*, S. Heddle and R.M. Sillitto, accepted by J. Mod. Opt., 14th Oct. 1992.

[6] *Reduction of aliasing in correlation using a pixelated Spatial Light Modulator*, S. Heddle, D.G. Vass and R.M. Sillitto, Proc. SPIE **1772**, p116 (1992).

[7] *Optical correlation using pixellated spatial light modulators*, S. Heddle, PhD thesis, University of Edinburgh (1993).

[8] *Compact optical correlator design*, J.A. Davis, M.A. Waring, G.W. Bach, R.A. Lilly and D.M. Cottrell, App. Opt. **28**(1), p10 (1990).