Abstract

Free-space optical interconnection offers a promising solution to the data bottleneck problems encountered in current computing machinery. Hybrid systems which use electronic processing and optical data connections could alleviate future restrictions on bandwidth. Spatial light modulators (SLMs) are a key part of the input interface between electronic and optical data transfer.

The ferroelectric liquid crystal over silicon SLM is a reflective amplitude- or phase-modulation device that has been pioneered at Edinburgh University. The fabrication of such devices still involves delicate manual steps, in part due to the poorly-understood nature of the FLC material. The exacting tolerances involved lead to characteristic error patterns in finished devices, which affect performance in quantifiable ways.

One of the most effective ways that the liquid crystal SLM can be used as a data switch is by phase holography, whereby a binary phase pattern on the SLM causes a convolution of input light with the Fourier transform of the binary pattern upon passage through a lens. SLM faults can affect this operation in subtle ways by causing phase errors and contrast loss in the modulation. This is shown using a matrix treatment of birefringence and the output of modified hologram design software.

Moreover, the design algorithm can compensate somewhat for the faults in a given SLM by adjusting the phase mask to suit. In a process analogous to adaptive optical techniques, some of the losses can be recovered. This may allow systems to function tolerably even with imperfect modulator devices. Results are presented which show that it is possible to design phase masks to minimize losses due to typical SLM defects. Cell thickness profiles are captured and used to create fault-compensated phase masks.

Acknowledgements

I would like to thank my PhD. supervisors, David Vass, Will Hossack and Ian Underwood, for their support and encouragement throughout this work. Also Sergei Samus, whose thesis was invaluable and whose C code seemed to be partly in Russian; and Iain Rankin, who collaborated on SLM production and designed the PC interface (and knew the details). I want to thank, too, the peerless Andy Garrie for efficient technical support. My funding was from the EPSRC, for which opportunity I am grateful.

Throughout my time in the Applied Optics group, much discussion and coffee was usually forthcoming from my colleagues James Gourlay, Rhys Williams, Georg Bodammer, Mark Begbie and Matthew Hart. Most of the conversation had some sort of point, and the coffee was handy. I wish them all well.

Personally, I owe a debt of sanity to Gavin Inglis, who understands the worth of games, and to Uta Staiger, Dilip Sequeira and Sophi Marriage who gave me reasons to go home occasionally. Lastly, I want for once to thank my parents, who were always finding ways to be extremely useful.

Fault-Compensated Free-Space Holographic Interconnects Using FLC Spatial Light Modulators

Patrick Smith

Submitted for the degree of PhD.
Department of Physics & Astronomy
The University of Edinburgh

April 15, 1999

Contents

Chapter 1

Free Space Optical Interconnection

> *wanting connections, we found connections —*
> *always, everywhere, and between everything*
> Umberto Eco, *Foucault's Pendulum*

1.1 Overview

This thesis seeks to address problems in the field of free-space optical data interconnections that use:

1. a Fourier-transformed phase hologram to route optical data channels
2. a ferroelectric liquid crystal spatial light modulator (FLC SLM) as a dynamic optical switching device

In doing so it will illustrate some of the issues associated with the holographic strategy of optical data switching, as well as the fabrication of FLC SLMs. I will discuss the possibilities for correcting FLC SLM fabrication defects by exploiting the methods of Fourier hologram design, and present experimental results demonstrating this technique.

In **chapter 1** I will briefly put the argument for the development of optical interconnections and describe some of the proposed techniques for their implementation. **chapter 2** describes in some detail the principle and operation of FLC SLMs, with some practical aspects of their manufacture. **chapter 3** is a discussion of the common defects associated with these devices and their relative severity for coherent optical processing, illustrated by reference to an optical system using 512^2 resolution FLC SLMs. **chapter 4** addresses interconnection strategies, especially an iterative algorithm for generating discrete-phase-level Fourier holograms. Finally in **chapter 5** there is an extended treatment of the use of iterative techniques to optimize holographic masks for imperfect devices, backed up by findings from adapted hologram-design software.

1.2 Optical Computing

The extent to which optical devices can profitably be used as an information processing technology in conjunction with — or in place of — the ubiquitous electronics of current hardware is still an open question. It is widely recognised that we will reach some important limits of electronic computing technology within a decade, as the 'VLSI–chip package–circuit board' paradigm will be working to capacity. Optical components, particularly for data switching, have some seemingly key strengths if they can be allied to electronic data processing. Potentially, an optical approach could alleviate various data transfer bottlenecks. To date, though, this has not been achieved in widely available machines. At a research level, nascent practical implementations of optoelectronic computing are suggesting the complementary roles to be played by these different technologies in future computing equipment.

Electronic computing has increased in capability at a literally exponential rate. The so-called Moore's Law, from a remark made by Intel president Gordon Moore in 1965, holds that the number of transistors being put on silicon chips *doubles*

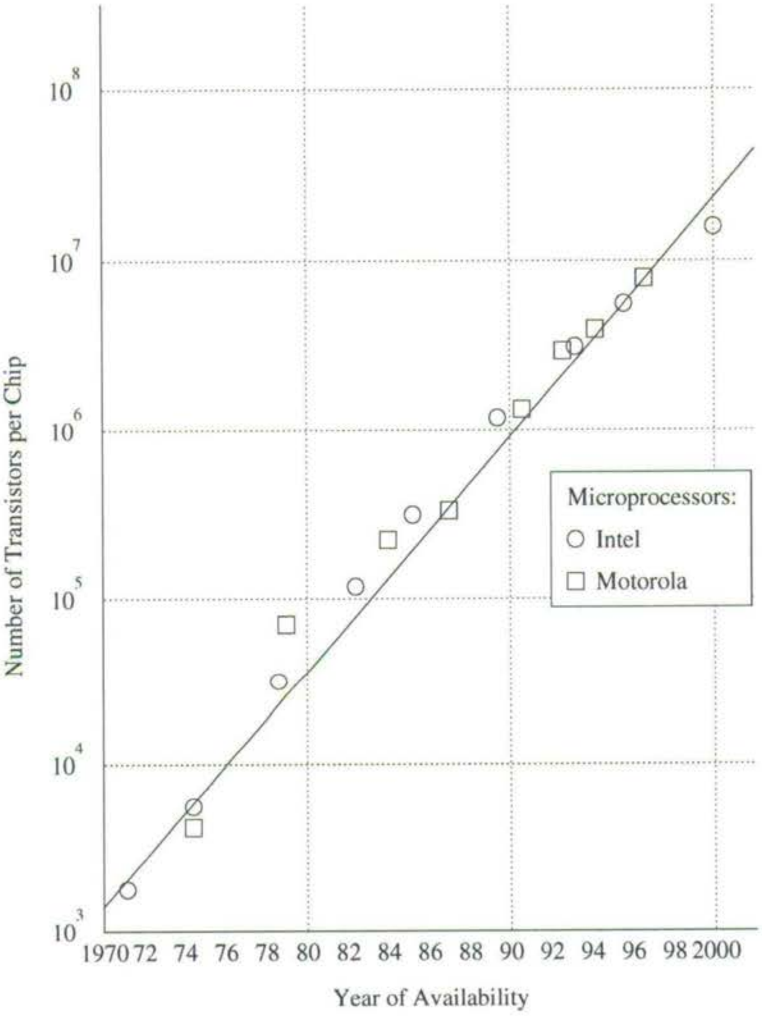

Figure 1.1: 'Moore's Law' — that the complexity of available computer chips grows exponentially with time — has so far held true; after [1].

every 18 months, as can be seen from figure 1.1.[1] This 'law' is inevitably doomed, but when it will break down is uncertain. A recent Scientific American issue on semiconductor technology[1] quotes Gordon Moore today:

> Moore predicts that tweaking the lenses, robots and ultraviolet lasers used to etch circuits onto silicon will extract perhaps two more generations of processors, with features 0.18, then 0.13 micron across, from

---

[1] Actually, Moore originally said every *year*.

> current optical techniques. 'Beyond that, life gets very interesting,' he notes.

The *physical* limits of purely electronic computing machines are quantifiable, and in sight[2], while the practical implementation of optical and optoelectronic computing strategies is beginning to catch up with the wealth of theoretical models. I wish to draw a distinction here between the *optical computer* which uses optical *logic* elements, and the encroachment of optical solutions in 'conventional' hardware.

Computing *with* light, which is ostensibly a provably good idea, has a chequered reputation. By the term *optical computing*[3] I imply using optical effects themselves to process (as well as transmit) information, in a manner analogous to electronic computers. Optical computing has been much theorised, but some of the proposed strategies have come to be regarded as misconceived, or at least esoteric. The projected optical or electro-optical machine architectures described a decade ago have not become commonplace. In particular, as research effort has moved into the experimental regime, a generic set of optoelectronic components has been slow to emerge — the fundamental building blocks of future optical computing. Emphasis has shifted, too, from building optical analogues of electronic computers, to overcoming specific limitations of electronics.

The key components in a hybrid optoelectronic technology are, naturally, the ones which interface between the two regimes:

electrically addressed emitters and modulators
optically addressed detectors

The spatial light modulator (SLM), which can write two-dimensional information into a light beam, is one such device. Of the various SLM technologies[4], the Applied Optics Group at the University of Edinburgh has pioneered the development of ferroelectric liquid crystal (FLC) over silicon devices [5, 6, 7].

Other technologies proposed for the electro-optic interface include Vertical Cavity Surface-Emitting Lasers, Micro-Cavity LEDs, Self-Electrooptic Effect (quantum well) Devices, Fabry-Perot Modulators, integrated waveguides and optical Micro-Electro-Mechanical Systems[8, 9].

The key problems for an efficient optical detector technology are *integrating* the detectors with the electronics which uses them as inputs, and then the *alignment* of incoming optical channels with those detectors. In effect, this amounts to the creation of optical 'pins' on-chip, that is, input channels to the VLSI in the same way as the traditional metal pins of a chip package. SEEDs, for example, are capable of acting as detectors as well as modulators, and with the high bandwidth required for low-level integration. SLMs are proposed as a means to address such an array of 'optical pins'.

Practicable solutions have and are being developed which combine electronics and optics in systems which play to the strengths of each. Specifically, the field of optically-interconnected machines has received attention due to current communications bottlenecks at the chip-to-chip and board-to-board scales. Light is a high-bandwidth, low interaction medium. VLSI is a powerful technology for processing data, but essentially two-dimensional. Ideally we would use both.

1.3 Optical Interconnection

1.3.1 Scale

The argument for free-space optical interconnections is dependent on the *range* of the data connection: we can provisionally distinguish various interconnection regimes distributed on a range scale from VLSI interconnect layers to inter-machine networks. Is there a regime where a free-space optical connection offers an advantage over electronics? There are certainly identifiable data bottlenecks in electronic systems which present an interconnection challenge.

1. At the shortest ($\sim\mu$m, mm) ranges are on-chip connections, joining logic gates or functionally different areas of an integrated circuit. VLSI interconnect is a problem for increasingly complex integrated circuits, as the *proportion* of processor area required for interconnect lines increases as a function of the number of gates; on a 10^7 transistor chip, over half of the die area is interconnect[10]. Moreover, the physical limits of the VLSI lithography determine both minimum feature size (due to focusable wavelengths) and maximum number of interconnect layers (due to the depth of focus). Proposed optical solutions at this range have used waveguides in so-called *planar optics*[11] configurations: emitters and detectors integrated either into the silicon VLSI or in a parallel GaAs circuit communicate through an overlaid plane of waveguide, as shown in figure 1.2. There are no tracks defined in the waveguide; the light pulse is either aimed at the receiver or broadcast.

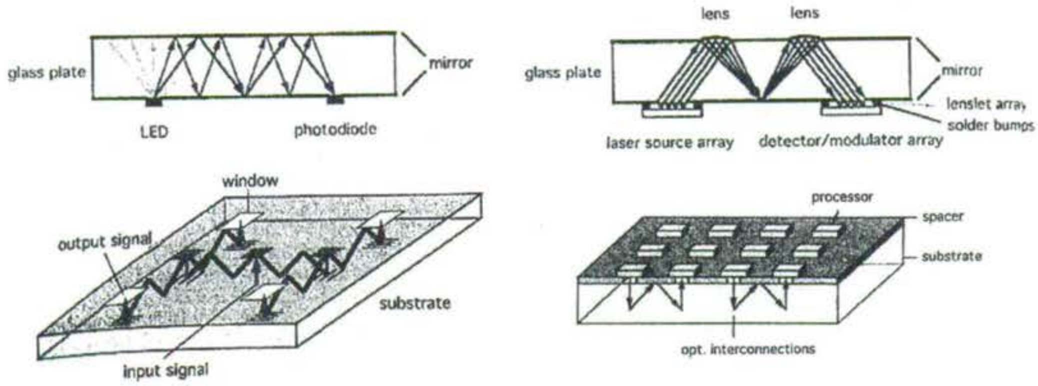

Figure 1.2: Planar optical interconnect configurations from [11].

2. At the $\sim$cm range, chip-to-chip connections within a single multi-chip-module (MCM) typically join the components of a semi-autonomous processing module. As the amount of processing possible on a single chip increases, devices run into packaging problems. Data transfer on and off chip, and the number of generalised 'pins' required in the package, becomes a limiting factor. At this 'range free-space interconnections become viable[12].

'Optical pins', such as VCSEL/SEED pairings, offer an integrated solution to inter-chip communications.

3. At the ~10cm scale, MCMs are connected at the board-to-board range, typically to build some form of server — be it a desktop PC or network router. This is a problematic regime for high-bandwidth applications requiring complex, possibly dynamic, connection geometries. The ability to *reconfigure* a connection would be an attractive quality, for example, for a parallel computer. A parallel machine whose topology can be altered remotely — and possibly during the course of operation — needs an electrically-controllable connection.

   Under consideration here are high-performance centralised computing facilities such as massively-parallel servers, video-on-demand servers and telecommunications switches. Crucially in these cases, the connection scheme is being reconfigured at significantly less than the data rate in any channel. As noted, this is an important scale for massively-parallel systems, which potentially require a high-bandwidth bus between parallel processing modules[10]. Path lengths, at ~0.1m, favour a free-space optical system, and accordingly, free-space interconnections have been proposed at this range for distributed clock signals[13] and an inter-processor optical bus[14].

   It is significant that today's most powerful computers use off-the-shelf processors in custom bus architectures: processing power is readily available, but parallelism requires high-performance connectivity. Optical fibre-based solutions — as have been implemented[15] — are not easily reconfigurable. For telecoms switching, currently available all-electrical and opto-mechanical solutions will be supplanted by electro-optical devices[16]. Over the board-to-board connection range, the advantage of free-space is its flexibility.

4. There is a further intra-machine data connection range at ~1m for the high bandwidth (GBit/s) link between servers and peripherals, for which fibre optics is a possibility. A consideration for optical links which are

plugged and unplugged manually is that mechanical connectors which allow for repeatable alignment need to be available. Since fibres may have a 1μm core, this is not a trivial requirement.

5. Beyond these intra-machine ranges, local-area and wide-area networks include data connections over optical fibre of all larger ranges. On these scales there is a requirement for efficient routing of high-bandwidth optical network channels. Free-space optical network routers are envisaged, to interface between arrays of fibres. Free-space optical switching can provide complex reconfigurable single-step connections, without conversion of data to and from electronics en route.

1.3.2 Encoding

Various properties of light are appropriate for encoding information. This adds depth to optical interconnection design, especially since the non-interacting medium makes it easy to multiplex information channels. Free-space light data channels can be routed *through* each other without crosstalk. This fundamental lack of interaction is both a basic advantage of optics over electronics for interconnects, and a limitation for data processing operations. WDM is an example: different wavelengths can be carried by a single fibre, but wavelength itself is not primarily used as the data-representing property.

The obvious counterpart to electrical signal strength is intensity. Data is encoded by a modulated source, and the nonlinear conversion from analogue intensity to a binary signal is made either at a thresholding detector or by using pairs of differential detectors.

Wavelength 'encoding', as noted, is used in WDM to discriminate multiplexed beams in the same path, and can serve to separate read/write or input/output beams. This is important in optical memory applications e.g. bacteriorhodopsin[17] and holographic memory[18].

Polarisation encoding is used in some optical neural network and interconnect designs[19]; arguably, the beams in the optical systems in Chapter 5 are 'polarisation-encoded' to permit optical multiplexing. Specifically *optical* processing steps are possible by taking advantage of polarisation effects. Future 'optical logic' may have to use such effects as its set of primitive logic operations.

Phase is a further ephemeral optical property, and is almost exclusively relevant to light-light interactions. For this reason, phase encoding is the preserve of wholly optical processing. Phase-modulated light, as an information-laden beam, is obliged to perform some all-optical processing step whose results can be detected electro-optically. It is also the most fragile of these encoding strategies; phase is only meaningful over a single wavelength, thus in system terms we are dealing with a ~10nm–100nm length scale. This imposes strict tolerances on the alignment and stability of a system, the reward for which is a unique set of primitive processing operations.

Compared to electrical encoding of data, the effects of this depth in optical information carriage are twofold. There is the possibility of simultaneous multiplexing without crosstalk, which electrical multiplexing is not so well suited for. There is less need, for example, to use a time-division multiplexing scheme. Secondly, the fundamental operations of optical processing are different from those of digital electronics: Fourier transforms are easy in optics, logic gates more difficult. Optical correlators exemplify the advantage in encoding data (a target pattern) as the appropriate optical quantity (phase) in order to perform an operation that would be computationally intensive.

1.3.3 Routing

For any optical interconnect, light channels need to be opened between arbitrary source and detector locations. In addressing the problem of steering light from one to the other, there is, first of all, a strategic choice between free space and waveguide routing. Some proposed methods — either with fibres or integrated waveguide layers — have been touched on above. Free space routing strategies further subdivide into smart-pixel and all-optical systems.

An array of so-called 'smart pixels' — detectors/emitter pairs with built-in logic — can be used as a free-space router. In order to route an array of light channels into an arbitrary new arrangement, the pixel array backplane can dynamically 'wire' an array of detectors to an (interleaving) array of emitters. This is, in itself, potentially a data processing step, as smart pixels can perform nearest-neighbour operations in parallel[20].

All-optical routing is attractive, as it bypasses the delay and hardware overhead in an optics–electronics–optics routing operation. We would like an optical element which can steer an array of beams into an array of detectors in any arrangement — a full crossbar switch. An array of diffraction gratings has been used for this purpose[21]. The diffractive elements dealt with in this thesis are used to *replicate* their optical input at multiple output locations; this operation has also been proposed as an optical crossbar[22, 23]. Other all-optical routing mechanisms include tilt-mirrors[9] and acousto-optic switches.

This thesis is concerned with the holographic phase mask as a one-to-many interconnection element, envisaged as the broadcast operation in a reconfigurable multi-module server.

1.4 Miniature Display Technology

The development of SLM technology for electronics-to-optics interfaces has an overlap with the development of SLMs as miniature displays. Display elements, such as the LC modulator from a video projector, have been pressed into service by optical correlator designers — just as optical MEMS devices are now being developed as dedicated displays.

The Edinburgh 512×512 pixel DRAM SLM is intended chiefly as a miniature display, and has been driven as a 24-bit colour display by Iain Rankin[24]. As such it compares favourably with other miniature displays being developed around the world for commercial applications. Table 1.1 summarises a number of current miniature display devices. As well as liquid crystal, prototype miniature displays have been made using micro-electro-mechanical systems (MEMS) (Texas Instruments, Silicon Light Machines) and electroluminescent materials (Planar) as the modulation mechanism.

Miniature displays will start to be incorporated in consumer products in 1998, and are expected to become ubiquitous.

Organisation	Silicon Light Machines	Microdisplay	Kopin
Resolution	$\sim 1000^2$ pixels	640x480	320x240
Array Size	25mm diagonal	5mm diagonal	4.8x3.6mm
Modulator	MEMS grating	AM nematic	AM nematic
Modulation	Binary colour	24-bit colour	8-bit greyscale
Contrast	200/1	N/A	80/1
Speed	20ns pixel switch	1ms pixel switch	60fps

Organisation	Spatial Light	Spatial Light	Displaytech	SLiMDis
Resolution	704x512	1024x768	640x480	1024x768
Array Size	14x10mm	20x15mm	8x6mm	12x9mm
Modulator	AM nematic	AM nematic	FLC	FLC
Modulation	8-bit greyscale	8-bit greyscale	Binary	Binary
Contrast	>128/1	N/A	N/A	100/1
Speed	50fps	N/A	9-bit colour 60fps	8-bit colour 50fps

Organisation	Texas Instruments	Planar	Planar
Resolution	848x600	640x480	640x480
Array Size	14x10mm	15.5x10.6mm	15.5x10.6mm
Modulator	Tilting mirrors	Electroluminescent	Electroluminescent
Modulation	Binary	Binary mono	Binary colour
Contrast	N/A (but high)	100/1	100/1
Speed	<1ms pixel switch	60fps	60fps

Table 1.1: A sample of worldwide miniature display technology, as of late 1998. (*fps*: frames per second, when displayed at the quoted modulation depth)

Chapter 2

FLC Spatial Light Modulators

Hence, differences observed 'twixt light and light
Derive from this, and not from rare and dense;
This formal principle makes dim and bright,

Agreeably to its own excellence.
Dante, *Paradise* trans. Dorothy Sayers

The *ferroelectric liquid crystal over silicon spatial light modulator* (FLCoS SLM) is a pixellated, phase-modulating device that uses the electric dipole and birefringence properties of FLC to produce what is effectively a switchable half-wave plate. The device is shown schematically in figure 2.1, and a packaged device can be seen in figure 2.2.

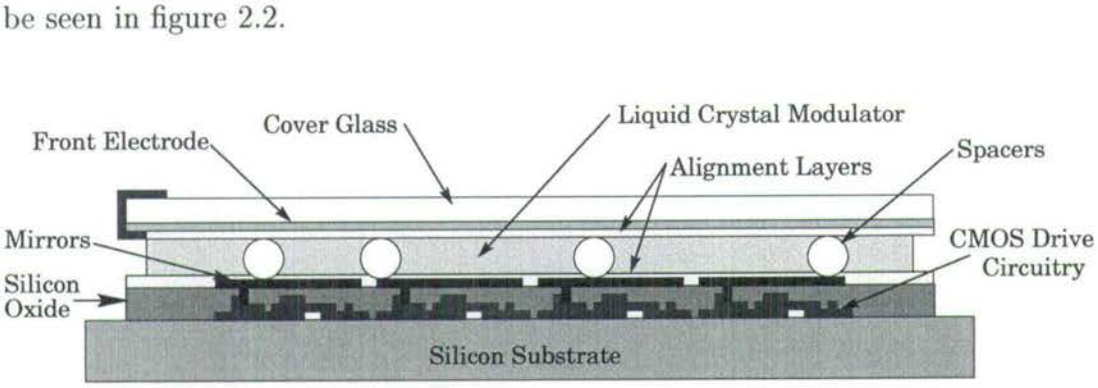

Figure 2.1: Cross-section of a planarised FLCoS SLM.

As the diagram shows, the liquid crystal layer is sandwiched between the addressing circuitry and the cover glass, with very little between the top CMOS metal

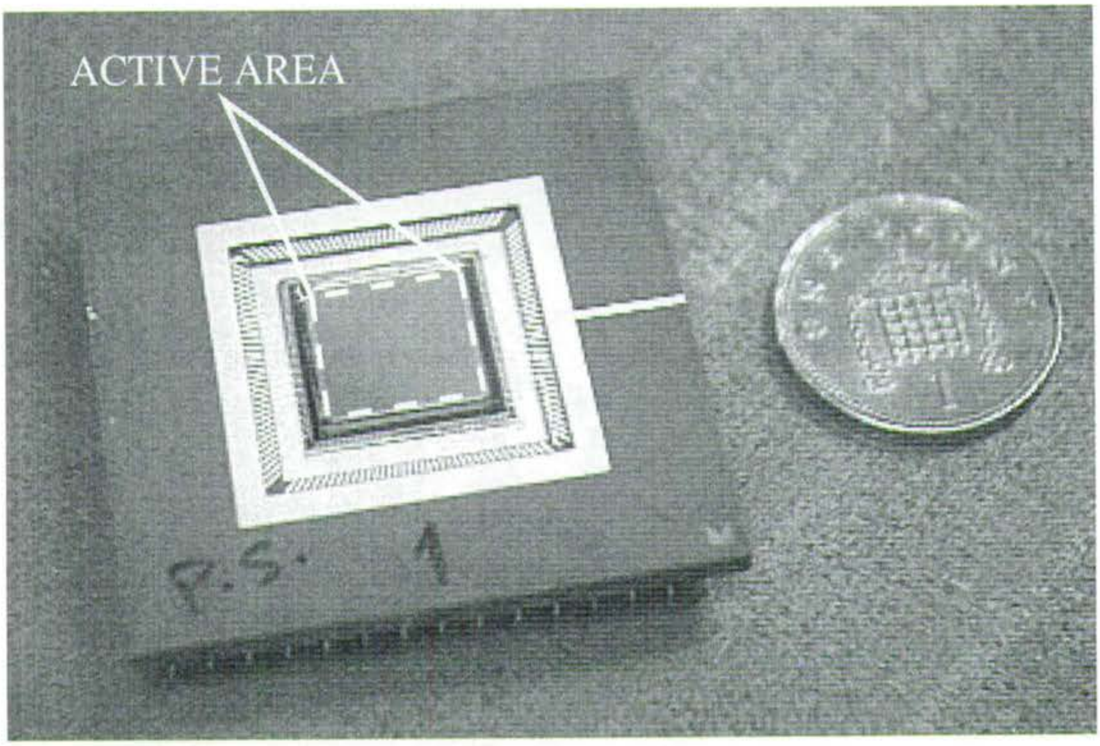

Figure 2.2: A packaged FLCoS SLM.

layer and the crystal. Separation between silicon and glass is maintained by glass spacer particles, either in the adhesive or distributed over the surface. There is a transparent conductive layer of indium tin oxide (ITO) on the underside of the glass to serve as an electrode, and additional coatings to promote crystal alignment on either inner surface.

Used as a 2D modulation component, the FLCoS SLM promises a number of advantages compared to other technologies — such as those summarised in table 1.1:

- $\sim 10^5/\text{cm}^2$ pixel resolution
- $100\mu s$ pixel switching time
- amplitude-only or phase-only modulation modes
- conventional CMOS backplane technology

- 80% pixel fill ratio

The FLCoS SLM has a faster switching time than acive-matrix nematic devices, with a contrast ratio (in amplitude modulation mode) typically at 100 : 1. It can also modulate purely in phase, which current micro-electro-mechanical SLMs (although faster and higher contrast) cannot.

The origins of these SLM properties will be discussed in more detail, as a lead-in to showing how the performance of real SLMs can fall short of that expected from the design.

2.1 Technology

2.1.1 Ferroelectric Liquid Crystal

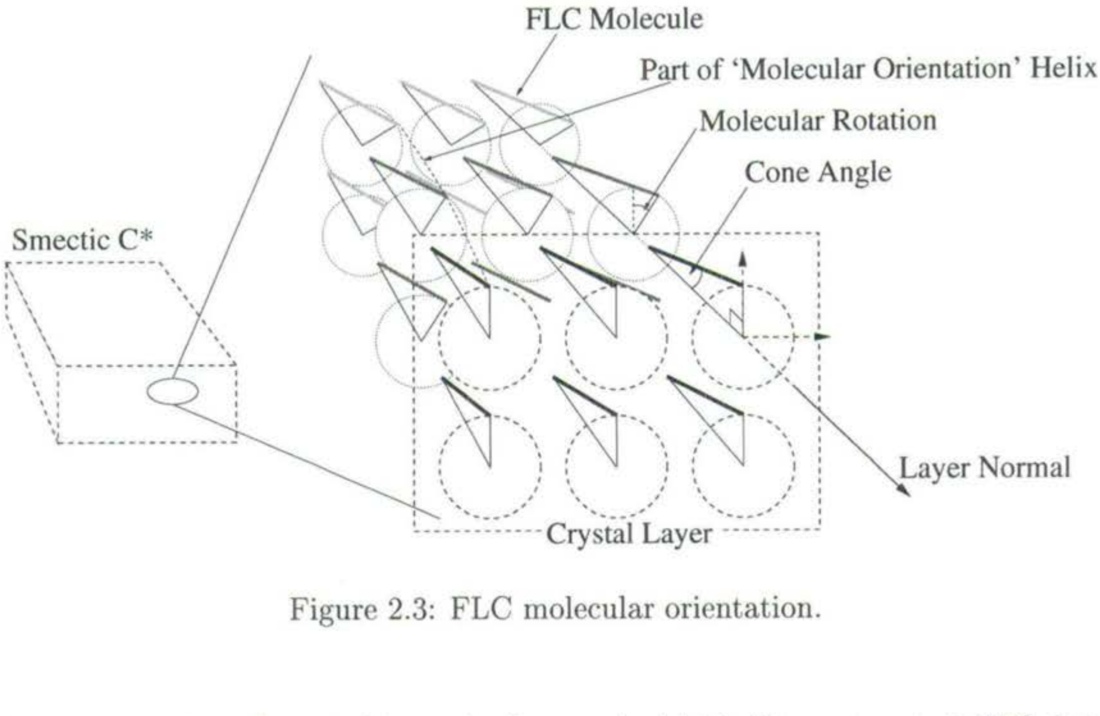

Figure 2.3: FLC molecular orientation.

Ferroelectric liquid crystal is, optically, a uniaxial birefringent material[25]; it is also ferroelectric because the rod-shaped organic molecule of which it is composed possesses a permanent electric dipole. We speak of the molecule having an 'axis' along its length, with the dipole approximately at right angles to this axis. These properties can be exploited most dramatically in the smectic C* liquid crystal phase, in which molecules exhibit an unusual positional and orientational order. As shown in figure 2.3, in this phase the molecules lie in layers with their long axis at a slight *cone angle* to the layer normal. Furthermore, each successive layer has the molecular orientation rotated slightly *around* the layer normal, producing a chirality in the molecular direction on progression through the layers.

It is possible to influence the bulk material with an applied electric field, since a molecule will experience a force so as to align its internal electric dipole with the external field. In particular, a thin 'slice' of FLC, held between two surfaces parallel to the layer normal, has just two stable molecular orientations as shown in figure 2.4. This is the so-called *surface-stabilised* (SSFLC) orientation[26]: the chirality of the layers is lost, and all molecules lie parallel to the surfaces of the FLC slice, so that the angle between the molecular axis and the layer normal provides only two symmetrical orientations. With no applied electric field, molecules lie in either orientation in arbitrary domains. The molecular dipole, however, is at right angles to the long molecular axis, so that a field applied across the FLC slice will cause molecules to (effectively) rotate around the layer normal in an attempt to align with the field. This will drive the entire FLC into the same molecular orientation. Reversing the field polarity then drives the crystal into the other orientation, with the molecules rotated 180° around the layer normal and thus angled away from the normal in the opposite direction. Hence the SSFLC is bistable: it can be driven into two symmetrical alignments.

The optic axis of the FLC lies along the molecule long axis, so the above electric switching behaviour constitutes a switchable wave plate with two optic axis directions.

For use as an optical modulator, it is necessary to consider how the effect of the FLC on a transmitted beam will differ for either state. Given an incident polarised beam, it can be shown that as long as the incoming polarisation plane lies somewhere between the optic axis directions, the two possible resultant polarisa-

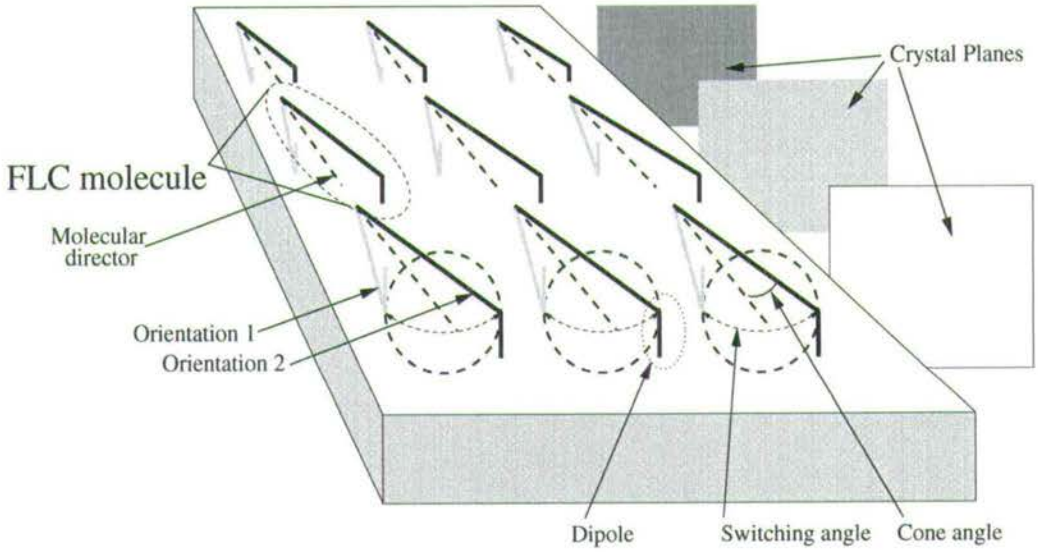

Figure 2.4: FLC in the surface-stabilised orientation.

tion states are exactly π out of phase with one another[27]. This is a consequence of the sense in which polarisation is rotated during passage through a uniaxial medium, and will be derived in the Jones matrix treatment of the situation, in chapter 4.

Hence, the FLC cell serves as a binary phase modulator; however, in general the output states are not plane polarised and have differing amplitudes. In order for the FLC to be 'just' a phase modulator, we impose the conditions that the output states should be of equal amplitude, and polarised in the same plane.

To provide equal amplitude outputs, the incoming polarisation plane must bisect the FLC switching angle. This then gives symmetrical conditions for either axis direction, whatever happens to the incident light.

To achieve plane-polarised light upon exit from the modulator, the FLC layer needs to be of half-wave plate thickness. For a half-wave plate, the plane of polarisation is 'reflected' in the optic axis. This will not in general result in the two states being polarised in the same plane, however, unless the incident plane of polarisation bisects a 90° switching angle. This means that in general, if the output light is analysed at right angles to the input light polarisation, only a component of each output will be passed.

Thus the output states can be of equal amplitude and polarised in the same plane, but relatively phase modulated by 0 or π depending on FLC orientation. That is, if the FLC slice has the appropriate thickness for a given light wavelength, it becomes a half wave plate with two bistable optic axis directions — which are, moreover, selectable electrically.

For this purpose, then, the most optically useful angle between the FLC molecular directions would be 90°. Unfortunately, this would represent a larger cone angle for the FLC molecule than has been achieved to date. The switching angle, as noted, is twice the FLC cone angle, and a 45° switching angle is currently achievable with FLC materials.

A 45° switching angle is, however, convenient for an amplitude modulator, since plane-polarised light incident on a half-wave plate, and polarised at 45° to its optic axis, will be rotated into a crossed polarisation plane. Therefore, if the incoming light to an FLC wave plate is polarised parallel to one of the FLC optic axis directions, the possible output states are orthogonally polarised. Adding an analyser then turns the FLC into a binary amplitude modulator. Note that this mode of operation and the phase modulation mode are separated simply by a rotation of the input polarisation, relative to the FLC axes, of half the switching angle, and that they are special cases of a phase/polarisation modulation behaviour. Again, this will be treated later by a Jones matrix description.

The use a of thin SSFLC cell as a reflective phase modulator (as in a FLCoS SLM) is shown schematically in figure 2.5. The cell is illuminated with polarised light whose plane of polarisation lies halfway between the optic axis directions, in this example at 22.5° from each. Having passed through a half-wave plate (the cell is traversed twice, and therefore is actually required to be the thickness of a *quarter-wave* plate), the light emerges with its polarisation plane effectively reflected in the optic axis, i.e. at ±45° relative to the incoming light. The output light is then analysed to produce two output states of amplitude 0.7 relative to the incident light, and identical save for a π phase difference.

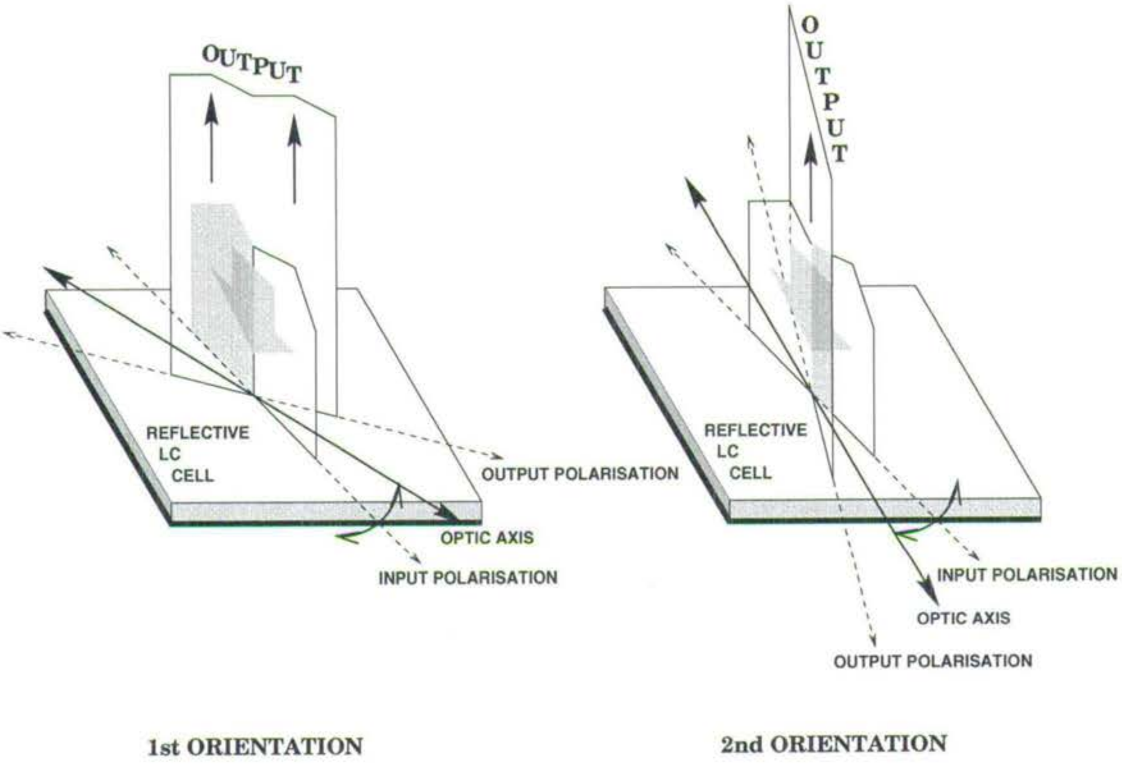

Figure 2.5: FLC as a switchable half-wave plate.

Commercial FLC products are mixtures of FLC and other liquid crystal materials designed so that the mixture is in the smectic C* phase at the operating temperature range. Of the SLMs used, the 256×256 pixel SRAM SLM has BDH SCE13[1], and the 512×512 pixel DRAM SLM uses Chisso CS1031[2]. Both have an effective refractive index of approximately 1.5, and SCE13 has a birefringence $(n_e - n_o)$ of 0.15 while CS1031 has a birefringence of 0.17, where n_e and n_o are the extraordinary and ordinary refractive indices, respectively.

---

[1] from Hoechst AG, Germany

[2] from Chisso Corporation, Japan

2.1.2 Silicon VLSI Backplane

The FLC-over-silicon SLM uses the electro-optic modulation characteristics of FLC in concert with an electronic device capable of storing binary pixel planes. It is effectively a CMOS memory chip, with a metal mirror connected to each bit-storage location. The CMOS is overlaid directly with FLC, and a charge on a pixel mirror switches the area of FLC between that pixel and the cover glass. The glass has a transparent conductive indium tin oxide layer to provide a front electrode — see figure 2.1.

This switching mechanism, therefore, resembles a conventional (nematic) LCD; indeed, the DRAM SLM has some similarities with a TFT nematic display — a logic gate at each pixel — on a miniature scale. However, the FLCoS SLM has several advantages. From the modulating material:

- The switching speed of FLC improves by an order of magnitude over nematic LC. Fast nematic mixtures are reported with ~1ms switching[3], but FLC will switch in 26μs[28].
- FLC is bistable on device time-scales (~ms). It is driven into both of its modulation states, thus has no relaxation delay during operation.
- A ±5V DC field is required for FLC switching, as opposed to the constant 1kHz AC, plus DC switching field required for nematic.

From the device:

The space-bandwidth product, defined as

$$\text{SBP} = [\text{pixel spatial frequency}] \times [\text{diameter of active area}],$$

is a useful measure for the capability of a coherent optical processing system. Miniature, high resolution modulator devices afford a greater degree of phase

---

[3] e.g. Microdisplay Inc. use such nematic mixtures — see table 1.1

control over a coherent beam. As an optical device driven directly from CMOS logic, the FLCoS SLM affords a large space-bandwidth product, with $\sim 10^4$ or $\sim 10^5$ pixels in a square array of side ~10 mm.

This high SBP makes holographic applications feasible; the size, and hence quality, of holographic output relies inversely on the pixel size. Several repetitions of the hologram unit pattern are needed. The high resolution of a FLCoS SLM gives it great versatility in coherent applications.

Of the several devices developed at Edinburgh, the 256^2 S-RAM SLM and the 512^2 D-RAM SLM were used in my experimental work. The practical differences between the use of these devices stem from the demands of static and dynamic memory addressing.

Resolution	256^2	512^2
Memory	S-RAM	D-RAM
Liquid Crystal	SCE13	CS1031
Die Size	14mm^2	14mm^2
Array Size	10.24mm^2	10.24mm^2
Pixel Pitch	$40\mu\text{m}^2$	$20\mu\text{m}^2$
Mirror Size	$37\mu\text{m}^2$	$18.4\mu\text{m}^2$
Fill Factor	85%	85%
Frame Rate	5kHz	8kHz

Table 2.1: Details of the 256^2 and 512^2 SLM

Both devices have a 14mm^2 silicon die with a 10mm^2 active area. For the 256^2 SLM, see [6]; for the 512^2 SLM, see [29]. Some relevant data for the devices are shown in table 2.1. Both devices have a *planarised* pixel array[30]. This is a silicon post-processing technique whereby the CMOS pixel array is overlaid with a thick silicon oxide layer which is chemical-mechanically polished flat. Metal mirrors are deposited on the flat oxide surface and each is joined to its corresponding pixel circuitry by a CMOS via. The end result is a flat mirror array with a very high fill factor, together with protection from light-induced charge leakage for the underlying CMOS.

2.1.3 Crystal Alignment

The smectic C* liquid crystal phase, as described above, is formed as the FLC mixture cools into its operating temperature range — which includes room temperature for SCE13 and Chisso 1031. While cooling it passes from isotropic through the nematic and smectic A phases, becoming progressively more ordered. During the nematic–smectic A–smectic C* phase transitions the crystal forms domains, within which molecular alignment is consistent. However, in an SLM we want just one domain over the device area. Moreover, we want to suppress variations in the molecular director angle over the device area, to produce a homogeneous modulation. Some method must be found to induce the FLC to align in a given direction without forming crystal defects[31]/domain walls.

The completed SLM device is a sandwich of FLC between a pixel mirror array built up on a silicon wafer, and a piece of cover glass. To apply a voltage across the FLC, the inner surface of the glass is coated with indium tin oxide, with an edge contact joined by bond wire to the backplane. SSFLC does not spontaneously align well, so the inner surfaces of the cell cavity are coated with an *alignment layer* to induce alignment in one direction across the array. The mechanism of FLC alignment is not well understood[28], and various materials are advocated for use as the alignment layer. All the SLMs used in this work had obliquely-deposited thin (200nm) films of silicon oxide over the mirror and ITO surfaces.

2.1.4 Cell Spacing

The thickness of the LC cell is set and maintained by the inclusion of glass spacer particles — either mixed in the glue which holds on the cover glass, or distributed across the mirror array before assembly. The cell is spaced to be a half-wave plate in reflection at some operating wavelength λ. That is, thickness $d = \frac{m\lambda}{n_e - n_o} + \frac{\lambda}{4(n_e - n_o)}$. Using CS1031 in He-Ne (633.28nm) laser illumination, $d = 1.86\mu$m for a $\frac{5\lambda}{4}$ thickness. However, the 512^2 SLM is intended as a display and is spaced (for green illumination) at 1.2μm. A $\frac{5\lambda}{4}$ cell is more easily spaced in this way than a $\frac{\lambda}{4}$ cell, and less likely to short the electrical contacts.

The inclusion of these spacer particles within the cell can cause seeding of FLC defects. There is a tradeoff between thickness uniformity and crystal quality, governed by the number of spacer particles per area.

2.1.5 Interfacing

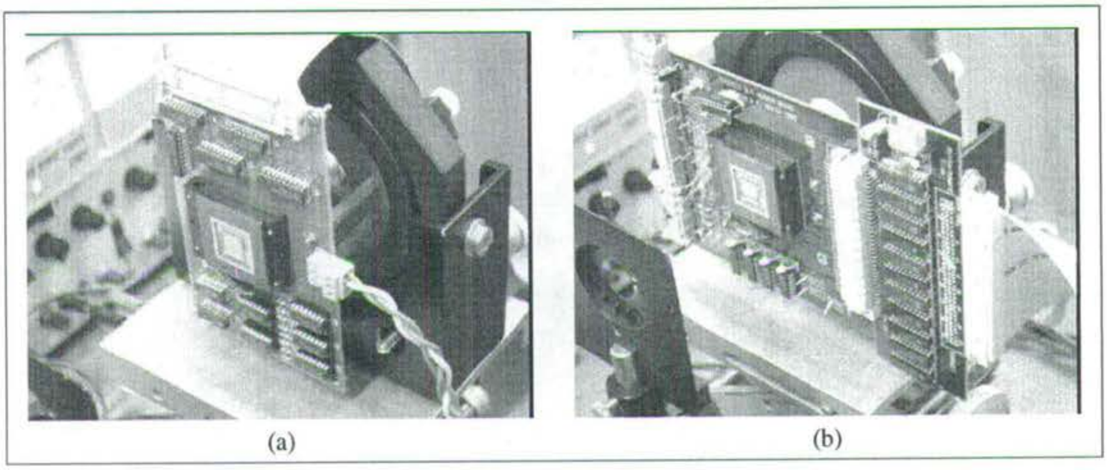
(a) (b)

Figure 2.6: (a) 256^2 and (b) 512^2 SLMs mounted in an experimental system.

The static/dynamic RAM distinction causes differences in the interfacing to these devices. The 256^2 SLM has effectively a single-frame memory which can be written once, without intervening frame memory, from a PC. Accordingly the 256^2 header PCB is driven from a PC IO card. The header PCBs for both devices can be seen in use in figure 2.6.

The 512^2 SLM must be constantly supplied with frames from off-chip. A moment's consideration shows that this is going to be a high data-rate exchange. If the binary frame rate is 1000Hz and each frame is 512^2 bits i.e. 32kbit, the connection is clearly going to run at 10-100Mbit/s. In order to accommodate this interfacing requirement a frame-serving interface board designed by Iain Rankin[32] is used to decouple the PC from the real-time SLM driving task. This main interface board contains an FPGA chip loaded with the interfacing logic, and two flash RAM framestores which alternately dump frames to the SLM's header board through a ribbon cable link.

Ironically, the interfacing requirements of the 512^2 SLM illustrate data bottlenecks such as those discussed in chapter 1. This is a serious issue for VGA+ resolution miniature displays, which must continually move at least $640 \times 400 \times 8$ bits of image onto the display at (at the very least) 20-50Hz. I will expand upon this below in a comparison of current miniature display technologies.

The 512^2 SLM header PCB is my implementation of a design by Iain Rankin. It can be seen driving an SLM from a ribbon cable link in figure 2.6(b). The board includes output signals to drive RGB LEDs in order to illuminate the SLM in synchronisation with the display of colour planes. For a coherent system I use a single frame, and use the appropriate LED signal to shutter a laser beam.

2.2 SLM Fabrication

2.2.1 Cell Assembly

Many of the problems associated with the devices arise during cell fabrication. The 512^2 SLMs used here were assembled and filled in the Applied Optics clean room, using a variety of techniques. The fabrication stages are:

1. Deposition of alignment layer material on backplane and cover glass. Silicon oxide is evaporated and deposits ballistically on the surfaces, which are held at an angle. This yields a 'ploughed field' surface by fractal growth[28].

2. Addition of spacers — mixed with glue or distributed across the backplane. 1.2μm diameter glass rods are used. Distribution is achieved by blowing an amount of spacers into a closed container placed over an exposed backplane, and allowing them to settle.

3. Gluing of cover glass to backplane. UV curing glue is deposited by hand at

the edges of the glass. After some trials, an automated glue-writing system was found to be less reliable than depositing the glue by hand. This is a key weak stage in the process; future work is planned in the group to address this.

4. Vacuum packing of cell. The cover glass is placed over the backplane, the cell is moved to a transparent plastic bag and the bag is sealed and evacuated. Atmospheric pressure forms an even pressure across both surfaces of the cell, with the object of producing an even cell gap, set by the spacers. A later version of this apparatus uses two rubber membranes with a perimeter seal: see [33]. The cell is held between the membranes, and then this sealed gap is evacuated.
5. Curing of glue under UV source. This is done while the cell is held in the evacuated pocket. The aim is to fix the cell while it is held at the uniform correct spacing.
6. Packaging and bonding of cell in proprietary ceramic PGA chip package. The glued cell is packaged and bonded by staff in the Edinburgh microfabrication facility.
7. Filling of cell under vacuum with FLC material in isotropic phase. Various filling ideas and pieces of equipment have been tried for SLMs. I found that a simple technique worked at least as effectively and reliably, given two provisos. See below.

Cell Filling

This is a large and complex subject, treated at length by Georg Bodammer[28], to whom I am indebted for advice in this area. The technique used to fill a thin gap with FLC is to evacuate the cell and then heat the FLC material until it changes phase (through smectic A and nematic) to isotropic, at which point it reduces in viscosity and will flow into the cell by capillary action. The cell is then removed from vacuum and cooled slowly, so that the crystal condenses back through its phases to smectic C* i.e. FLC. The final alignment of the FLC is found to depend critically on the *direction* and *speed of flow* of isotropic material into the cell; the optimum is a fast flow which maintains a flat filling front parallel to the edge of the cell. Additionally, the FLC must first be 'outgassed' (have its dissolved gases expelled) in vacuum to improve its optical properties, so it is typically left in vacuum for several hours before heating to isotropic.

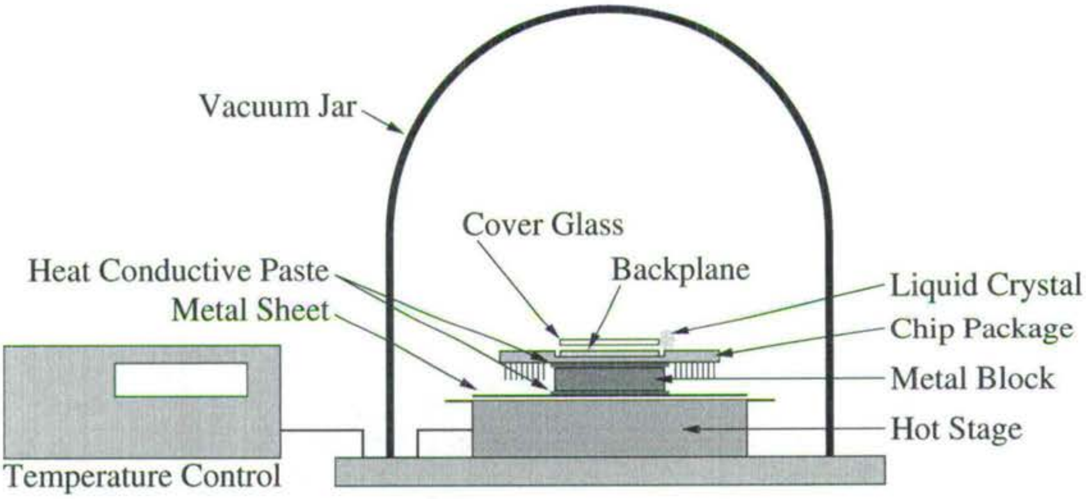

Figure 2.7: Schematic of the cell filling apparatus

The unfilled SLM has FLC added along one side of the glass/backplane gap. Due to the large number of bond wires on the 512^2 package, FLC can only be added in the middle of one side. The SLM is placed on a heat stage in a vacuum jar, as shown in figure 2.7. A metal block and thermal conducting paste provide good thermal conduction between a metal sheet on the heat stage, and the SLM package across the area of the silicon backplane. However, when the jar is evacuated there is very poor thermal conduction between the stage and the sheet. Although not originally the intention, this is a serendipitous effect for creating the filling conditions for optimum alignment, as will be seen.

The jar is evacuated and the FLC left to outgas. It was found that not only was over 24 hours required before bubbling gas was no longer visible in the FLC, but that any disturbance in previously-outgassed FLC resulted in fresh gas bubbles.

It was concluded that outgassing only occurs in a *thin surface layer* of the bulk FLC, since dissolved gas could be released from the deeper FLC only once it was exposed. For use in SLM filling, a reservoir of FLC was kept outgassing in the vacuum jar throughout the filling runs, and the surface layer drawn from.

After the outgassing period the stage is heated past the isotropic phase transition temperature of the FLC material. Chisso 1031 becomes isotropic at 97°C. The stage was heated to at least 115°C. Due to the poor thermal conduction to the SLM package, the FLC on the SLM does not change phase or flow into the cell. Air is let back into the chamber, at which point the package quickly heats to above the isotropic transition temperature across the area of the SLM array. The crystal changes phase to isotropic and flows rapidly across the hot surface. The filled cell is then left to cool in atmosphere.

In practice it was found that the the crucial variable in this simple filling strategy was a high temperature of backplane. The stage has to be hot enough to heat the SLM package to above the isotropic transition for long enough for the FLC material to flow across the whole array, and the hotter the cell the faster and more uniform the flow. The second proviso is that the FLC be outgassed throughout, and especially immediately before filling while on the SLM. Both of these provisions add time to the process.

2.2.2 Non-functioning Backplanes: Investigation Using SEM

The 512^2 SLM dies from one post-processed wafer were rendered inoperable by insufficient etching of the vias through the planarisation layer. Figure 2.8 shows a scanning electron microscope (SEM) image of a mechanically-prepared cross-section of a shallow via, taken from one such faulty SLM backplane. Once the planarised wafer has been chemical-mechanically polished, the precise thickness of the insulating planarisation layer is not known. This means that the subsequent etching step, whereby vias are put through to the CMOS circuitry, can be miscalculated. In the worst case there is no path for the contact to the metal pixel

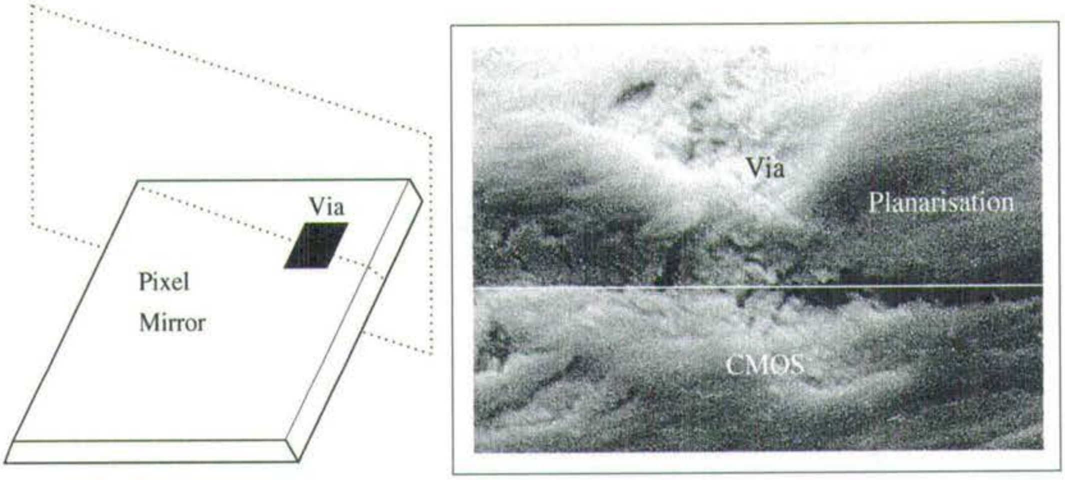

Figure 2.8: A cross-section through a pixel mirror, showing a via which has been insufficiently etched so that it does not supply a path through the insulating oxide layer to the underlying circuitry. There will be no voltage applied to the mirror, and hence no control over the FLC. This was a problem which caused the rejection of a whole 512^2 SLM wafer.

mirror, which is seen to be the case in the picture. Unfortunately this applies to the entire wafer of SLM dies, and in this instance was not diagnosed until after cell fabrication and filling using some of the dies. If there is no contact between the pixel mirrors and the CMOS, the SLM does not function, and is irreparable. This proved to be a major setback in obtaining 512^2 SLM devices for use. The solution to this problem is to characterise the chemical-mechanical polishing and etching steps in SLM wafer processing, and in fact this is now a research topic in the group.

2.3 Conclusions

FLCoS SLMs, then, involve a wide range of disciplines in their manufacture and operation: classical optics, complex materials, VLSI design, process and electrical engineering. The device is, undeniably, a hybrid of technologies. Some of the necessary science is new and/or poorly understood — for example, cell filling and its effect on crystal defects. In many areas there is a need for formal process engineering, especially to characterise and standardise the fabrication steps, and in fact this is currently taking place.

It is not enough to think of the FLCoS SLM purely as an optical element, as a silicon chip or as an exotic liquid crystal layer. Ignoring any one of these aspects has proved disastrous. It would indeed be a surprise if the interaction of these disparate components did not produce some difficulties, and in the next chapter the task is to describe and quantify them.

Chapter 3

FLCoS SLM Defects

> *Murphy's Law appears to be an ineluctable feature of our universe*
> Robert Matthews, in [34]

FLCoS SLMs are prone to certain characteristic defects due to their construction and operation. Much development effort has been invested in engineering solutions to improve the quality of manufactured devices. There are still reasons, however, for a post-facto approach where we accept that optical systems will need to function using imperfect components.

Assuming that the specific problems which affect FLCoS devices will change in severity, but not in nature, we can usefully ask how strict the tolerances need be for likely optical processing systems. More generally, there are some defect modes common to all reflective SLM technologies which can be considered, stemming in particular from the use of large-area integrated circuits. Also, as high volumes of SLMs are likely to be available as display elements, we can ask how the lower tolerances required for display applications will match up to the more demanding constraints of a coherent optical system.

This chapter seeks to describe and quantify the practical effects of FLC SLM defects in a phase modulation application. Figure 3.1 shows the classes of error-producing phenomena I will consider:

1. VLSI defects; pixel drop-out
2. cell defects:
    (a) backplane warp
    (b) thickness variation
3. liquid crystal misalignment

The hybrid nature of the SLM, which is both a VLSI device and an optical component, together with the caprices of a complex liquid crystal phase, gives rise to these three distinct areas of performance-reducing defects, which I will treat separately. This treatment is aimed primarily at coherent optical processing applications — such as Fourier transform holography, which will be covered in chapter 4. Eventually, we seek to minimise the performance loss caused by SLM defects in such applications.

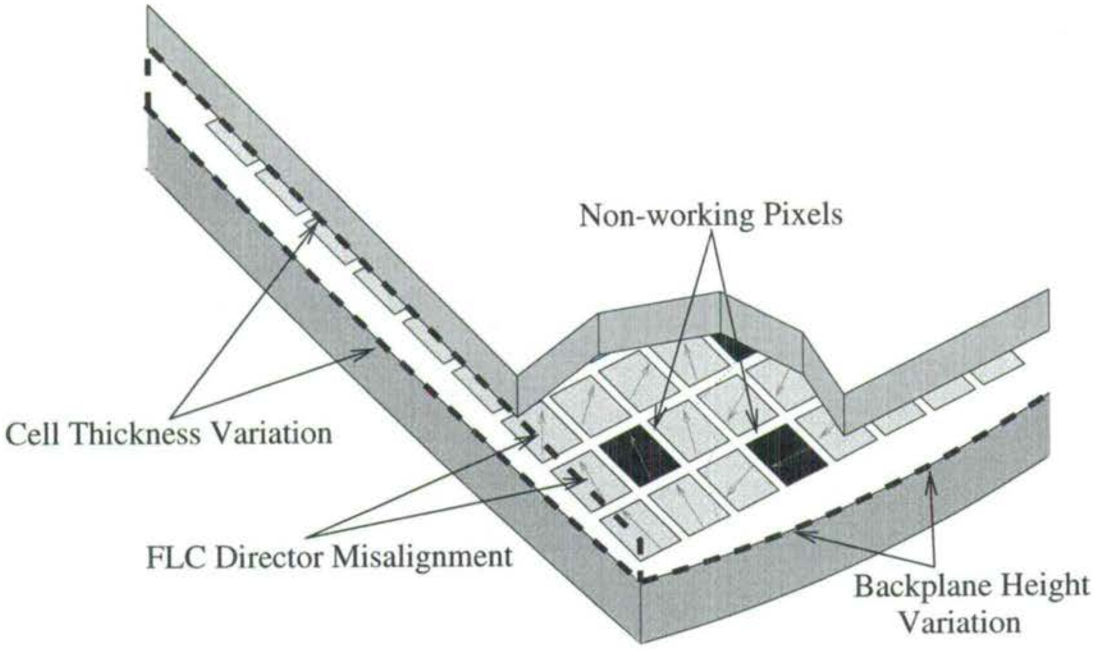

Figure 3.1: Features of a FLCoS SLM which reduce performance.

3.1 A Menagerie of SLM Defects

3.1.1 VLSI Backplane

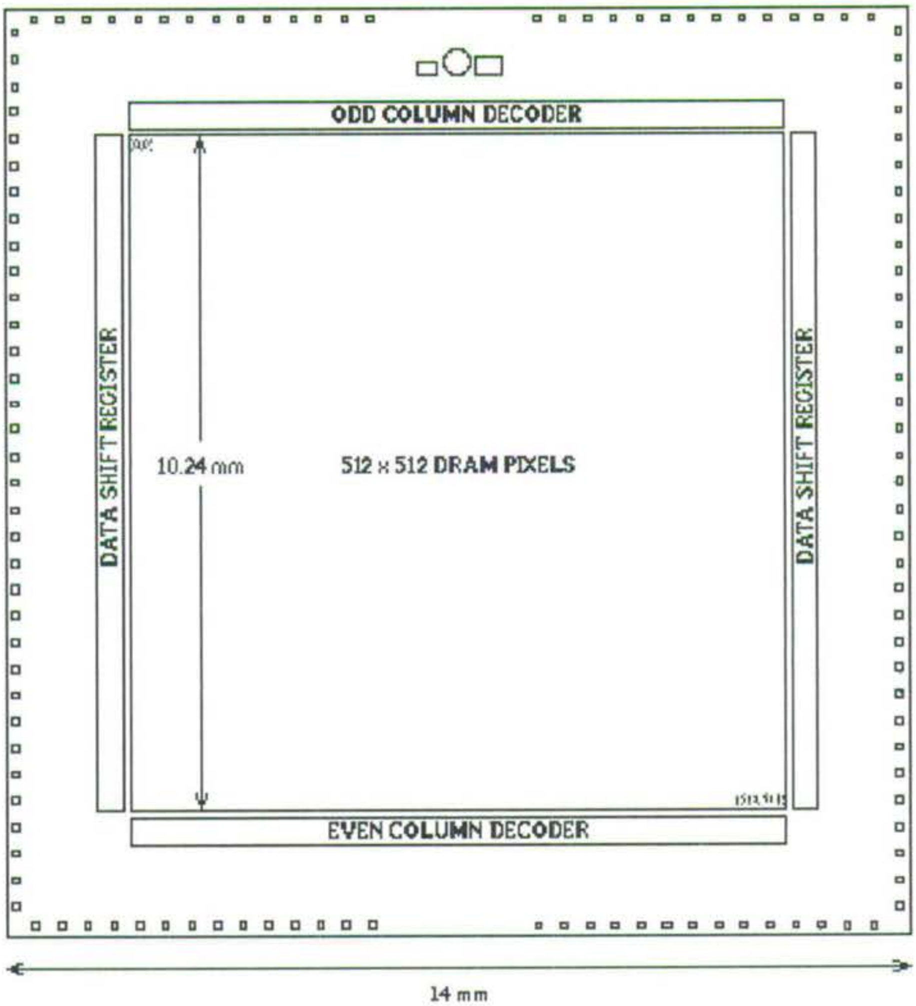

Figure 3.2: Diagram from 512^2 SLM documentation[29] showing the backplane addressing scheme

The VLSI backplane of a FLC SLM is effectively a large-area DRAM or SRAM array. In practice this post-processed 10mm^2 area of VLSI is prone to some of the established, characteristic faults associated with such memory chips, whereby an area of the memory array with a common addressing line can fail as a unit. Drabik[35] has considered the potential for optically-interconnected wafer-scale systems to tolerate fabrication defects, through an optical connection between operational areas. This, then, might be an application for a reconfigurable holographic interconnect. However, where VLSI elements such as FLCoS SLMs are themselves used to control optical switching, how is this very switching function affected by those same fabrication defects? The first point to note, perhaps, is that the holographic interconnection is inherently robust to failures of part of the device area, due to the non-localised encoding of information.

The effect of VLSI faults can, hence, be compounded by the addressing scheme of an SLM[29], whereby the pixel array is divided into 8-line blocks which are addressed by a series of 8-bit shift registers. In the 256^2 SLM there are 32 shift registers along one side of the array; in the 512^2 SLM this is repeated on the opposite side, and alternate blocks addressed from either side. Bus lines drive the shift registers: the 512^2 has bus line pads at each end of the two shift register strips. Within an 8-line block, pixel lines are addressed in sequence, and cross-referenced — to reach the individual pixel level — by sequentially addressing the SLM columns through 8-bit decoders. The 512^2 needs 2 decoders, which address odd and even columns respectively.

From the above, we might expect various examples of electrical failure to manifest themselves in finished devices:

1. Failure of a decoder removes control from 256 columns — the entire 256^2 device, or every other column of the 512^2.
2. Failure of a bus line removes control from the shift registers — the entire 256^2 device, or *alternate 8-line blocks over half of the* 512^2 *device.*
3. Failure of a shift register will remove electrical control from a rectangular stripe of the SLM, over which the FLC will relax into an arbitrary alignment and, hence, the SLM into an arbitrary modulation state. The visible effect

is a non-switchable band at some contrast to the rest of the array; this is not acceptable for a practicable device.

4. Line or column drop-out — resulting perhaps from a CMOS track discontinuity — is a reasonable possibility given that the SLM is extensively post-processed and, moreover, subjected to cell fabrication work before being packaged. This would be an unacceptable fault for a display device, but phase holography should tolerate a number of missing lines in a 512 × 512 resolution device.

5. Lastly, failure of some amount of the individual pixels should also be tolerable. Further, as this last effect is a graceful degradation it is possible to quantify the proportion of randomly-distributed pixels that must drop out of operation to seriously affect the operation of the SLM as a holographic fan-out element, and this will be shown in chapter 5.

3.1.2 Cell Uniformity

As a reflective phase-modulating optical element, the SLM exhibits errors associated with optical components. These arise due to two complementary features: height variation of the mirror array, and thickness variation of the birefringent material, which are ultimately due to the difficulty in preventing a 10mm^2 area of CMOS from showing height variations of the order of $0.1\mu\text{m}$. For display (and other incoherent) applications, it is minimally necessary that the mirror height variation does not alter the cell thickness to create a full-wave plate — and thus a blank — at any point. However, colour dispersion effects will be visible for lower variations, and contrast will suffer with a $\cos^2$ dependence upon the birefringence error. A reasonable limit for tolerable operation is half of this, with an allowable birefringence depth variation of $\frac{\lambda}{4}$. Hence this implies a maximum cell thickness variation of $\frac{\lambda}{8(n_e-n_o)}$ due to the double-pass in the cell. For HeNe illumination and CS1031, this condition equates to a maximum allowable cell thickness variation of $0.47\mu\text{m}$.

For coherent (including phase-modulation) applications such a '$\lambda/4$' condition certainly implies a relevant amount of *phase* error, as well as the loss of contrast; we expect modulation performance to degrade progressively until a full-wave plate prevents any throughput. The full implications of this condition for modulation efficiency will be demonstrated below.

Backplane Warp

Curvature of the SLM backplane is due to stresses induced in the silicon wafer during manufacture which are released when the wafer is cut into SLM dies. An individual die is typically warped with negative curvature[33]. Since the SLM is a reflective phase-modulating element, this adds an unwanted phase distortion to the outgoing wavefront. Moreover, since the cover glass is not thus warped, the LC-filled cavity has a varying thickness.

The effects of wavefront distortion are well documented in the literature of **adaptive optics**[36], although the Fourier transform operation involved here is not normally considered.

Thickness Variation

As a uniaxial birefringent element, the key quantity of the modulation layer is the difference between its two refractive indices, n_e and n_o. For a birefringent material of thickness d, the difference in phase between plane polarised light propagating with its plane of polarisation parallel to the extraordinary (optic) axis and parallel to the ordinary axis is $\frac{2\pi}{\lambda}d(n_e - n_o)$. For a half-wave plate made of a material with some given n_e and n_o, d is chosen so that this phase difference is π, i.e. $d = \frac{\lambda}{2(n_e - n_o)}$. For light polarised at some angle to the optic axis, the net effect of the half-wave plate is to reflect the plane of polarisation in the optic axis, as illustrated in figure 2.5.

In order to determine the output states from a birefringent plate whose thickness

d varies from the half-wave plate thickness, the well-known Jones matrix method can be used to trace the passage of polarised light through birefringent materials, as described in Goodman[37] — see also [38, 39]. I use a Jones matrix treatment to derive a description of elliptically polarised output states intermediate between the two plane-polarised extremes of the half-wave and full-wave plates.

In this scheme, a polarisation state is represented by a vector containing its complex x- and y-components, and all operations are described by 2×2 matrices. The Jones matrix for a uniaxial retarder, L_{uni}, is given by Goodman as:

$$\begin{aligned} L_{uni} &= L_{rot}(-\theta)\, L_{retard}(\beta d)\, L_{rot}(\theta) \\ &= \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & e^{-\imath\beta d} \end{pmatrix} \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \end{aligned} \tag{3.1}$$

where θ is the angle between the optic axis and the y-axis, and

$$\beta = \frac{2\pi(n_e - n_o)}{\lambda} \tag{3.2}$$

I define η to be *the variation, of the o-wave and e-wave phase difference, from the half-wave plate case*:

$$\eta = \beta d - \pi \tag{3.3}$$

Assume an input wave of unit amplitude polarised parallel to the y-axis:

$$A_{in} = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \tag{3.4}$$

If we set $\theta = \pi/8$, the cone angle of a typical FLC, the output vector becomes:

$$\begin{aligned} A_{out} &= L.A_{in} \qquad (3.5) \\ &= \begin{pmatrix} \frac{1}{2}\left[\frac{1}{\sqrt{2}} + \frac{1}{\sqrt{2}}\cos\eta - \imath\frac{1}{\sqrt{2}}\sin\eta\right] \\ \frac{1}{2}\left[1 - \frac{1}{\sqrt{2}} - \left(1 + \frac{1}{\sqrt{2}}\right)\cos\eta + \imath\left(1 + \frac{1}{\sqrt{2}}\right)\sin\eta\right] \end{pmatrix} \qquad (3.6) \end{aligned}$$

Thus, in general the x- and y-components are both complex i.e. have some phase offset depending on the material thickness. In figure 3.3 I plot the amplitudes and phase offsets of these components as η varies over a 2π range. In an optimally-oriented wave plate, with its optic axis angled at 45° to the input polarisation, there is a phase difference between the resolved x- and y-components at each thickness. Figure 3.3(c), showing the phase lead of the y-component over the x-component, indicates that the output elliptical polarisation states are further complicated by a changing phase relationship. A series of intermediate elliptical polarisation states for various thicknesses of plate may be deduced.

The established way to make use of phase modulation in an optical system is to analyse the output light from the FLC layer with a polariser crossed with respect to the input polarisation. Given A_{in} above, this is equivalent to polarisation along the x-axis, which leaves the x-component of A_{out} unchanged and sets the y-component to zero. Then, from equation 3.6 and figure 3.3, there are two effects detrimental to the operation of FLC as a phase modulator as the thickness varies: the x-component decreases in amplitude, resulting in an efficiency drop as more light is lost at the analyser, and the x-component's phase offset — relative to that for an ideal FLC thickness — increases, causing a phase difference between output from areas of the cell. For phase holography it is reasonable to assume that this will be a problem. Although the on and off x-component output states from any given pixel remain π out of phase, the output states from other pixels are not phase-synchronised.

This offset in phase between pixels can be illustrated by calculating example output polarisation states for some different FLC thicknesses. From equation

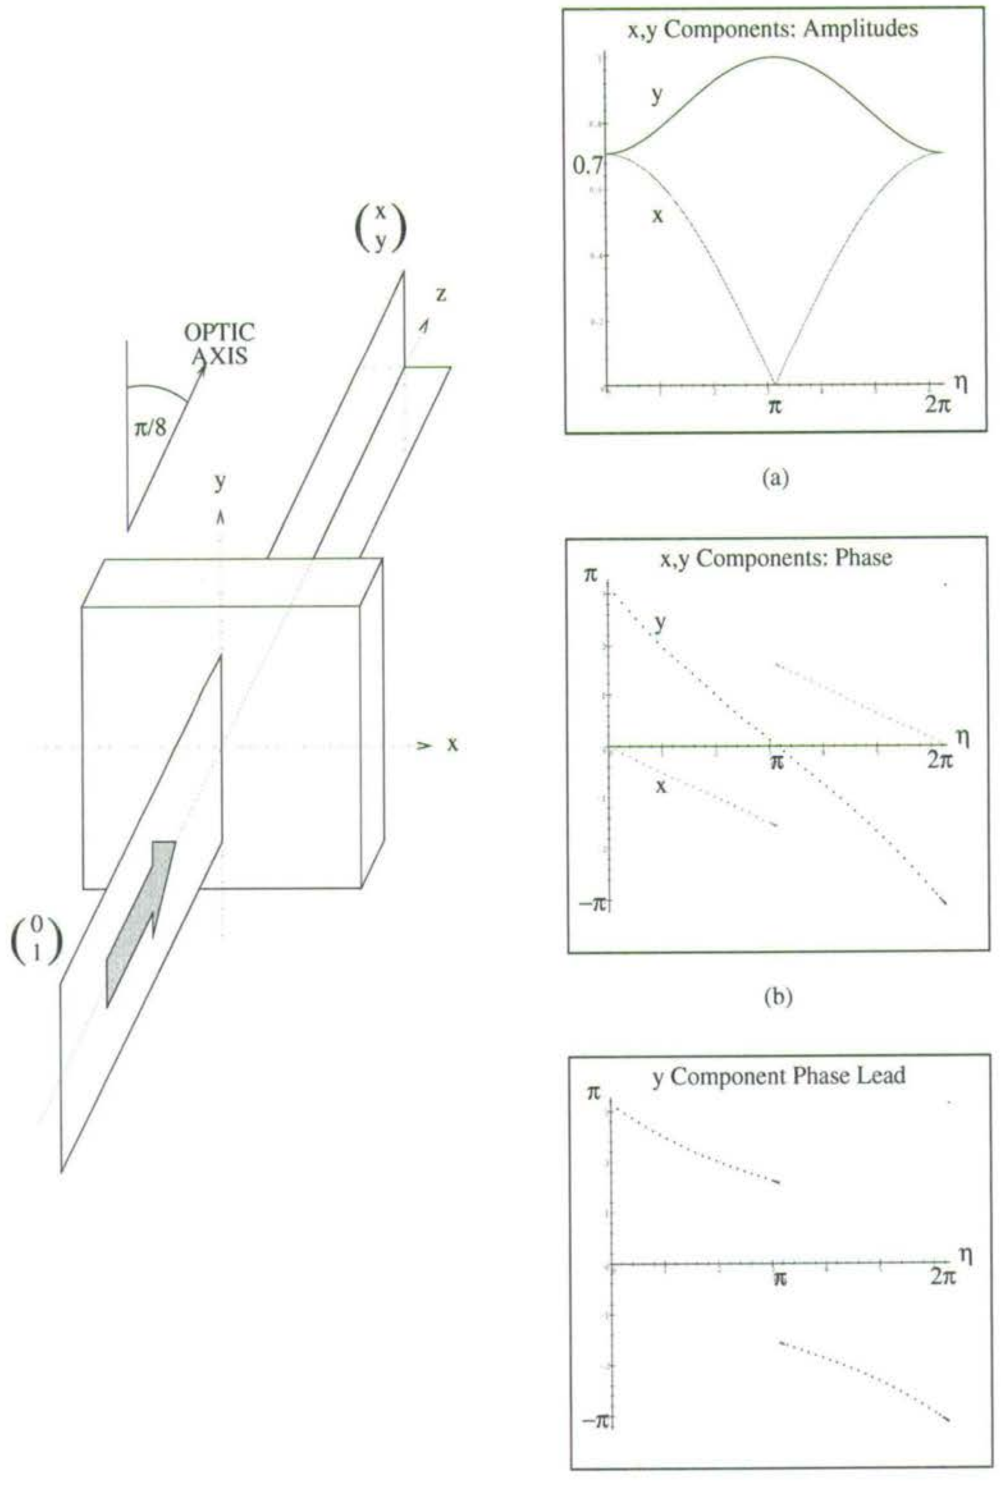

(a)

(b)

(c)

Figure 3.3: Graphs of (a) amplitude and (b) phase of the components of A_{out} against birefringence depth η, with (c) phase lead of the y-component over the x-component.

3.6, for a perfect half-wave plate ($\eta = 0$) we are not surprised to find

$$A_{out} = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} \end{pmatrix} \tag{3.7}$$

that is, a plane-polarised state at 45° to the input. The x-component is what we hope to use in an optical system.

Now suppose there is an area of the FLC cell where $\eta = \frac{\pi}{4}$, that is, birefringence depth is wrong by $\frac{\lambda}{8}$. Equation 3.6 yields

$$A_{out} = \begin{pmatrix} \frac{1}{2}\left[\frac{1}{\sqrt{2}} + \frac{1}{2} - \imath\frac{1}{2}\right] \\ \frac{1}{2}\left[\frac{1}{2} - \sqrt{2} + \imath\left(\frac{1}{\sqrt{2}} + \frac{1}{2}\right)\right] \end{pmatrix} \tag{3.8}$$

and two effects are obvious. Firstly, the x-component has decreased in amplitude (0.71 down to 0.65). This is expected. Secondly, the x-component is complex, and has a phase offset of $-\frac{\pi}{8}$.

Thus the *birefringence depth variation* can imprint a spatial phase variation on the outgoing SLM wavefront.

3.1.3 FLC Defects

FLC is a highly-ordered phase of liquid crystal, and complex crystal defects arise during the phase transitions from the isotropic state through to FLC. I will consider one effect which affects the SLM's performance as a phase modulator: local variations in the optic axis direction caused by misalignment of the FLC molecules. The FLC molecules are intended to align with molecular director oriented parallel to a single axis — which is parallel to the cell faces — across the entire area of the mirror array. However,

- the FLC forms domains wherein molecular alignment is local to the domain,
- alignment is largely set by direction of FLC flow during cell filling[28],
- there is a pre-tilt to the molecules which tilts them out of a plane parallel to the cell faces and reduces the useful switching angle, and

- the molecules are generally not homogeneously aligned through the thickness of the cell; it is an approximation to treat them so.

Neglecting, initially, the effects of molecular tilt out of a plane parallel to the cover glass, we can consider areas of the cell to be FLC domains with varying molecular orientation. Effectively, the cell's optic axis direction is varying with respect to the incoming polarisation direction, and, again, we can use the Jones matrix technique to investigate the effect this may have on the output state.

The simplest treatment of the effect of optic axis misorientation on phase modulation is to assume that the modulator functions as a half-wave plate between crossed polarisers and use Malus' Law:

$$I = A^2.\cos^2\phi \tag{3.9}$$

which gives the transmitted light intensity where ϕ is the angle between two linear polarisers. For the half-wave plate case, if θ is the angle between the first polariser and the optic axis of the plate, and the second polariser is crossed, then:

$$I = A^2.\sin^2 2\theta \tag{3.10}$$

Note that in the signed 'amplitude' version of the above expression:

$$A_{out} = A_{in}.\sin 2\theta \tag{3.11}$$

the amplitude A_{out} indicates, by its sign, the sense in which the polarisation has been rotated, and thus a relative 0 or π phase difference in the output. The two optic axis directions of FLC can be represented by $(\theta \pm \frac{\alpha}{2})$, where α is the FLC switching angle and θ is now the angle of the molecular director, to plot a sinusoidal variation in the amplitude of the output states as the director angle θ is varied. This is shown in figure 3.4. Two curves are plotted, one for each of the modulation states of an FLC half-wave plate with switching angle $\alpha = 45°$. The two states are exactly π out of phase when the input polarisation angle lies anywhere between the two optic axis positions (i.e. an angular range, for binary phase modulation, of 45°). The two output states have equal amplitude when the input polarisation bisects the switching angle, and this amplitude is $\sin 45° = \frac{1}{\sqrt{2}}$ times the input amplitude. This loss of light, even at the optimum phase modulation angle, is unavoidable for a FLC with less than 90° switching angle.

Pre-tilt is likely to reduce the switching angle of the FLC, reducing the angular range for binary phase modulation and also the amplitudes of the output states. In figure 3.4, this would shift the two sin curves closer together. For a 45° switching angle (a 22.5° cone angle), $\frac{1}{4}$ of the possible range of the molecular director orientation results in a binary phase modulation of output states, albeit not symmetrical in amplitude. Smaller cone angles reduce this phase modulation range, and hence also the likelihood that a misaligned FLC domain will modulate phase when between crossed polarisers. An amplitude modulation will arise, however, from all orientations except when the molecular director lies parallel to one of the polarisation axes.

Variations in *intensity* modulation due to both FLC thickness and molecular orientation are combined in figure 3.5, which shows the intensity of the two phase-modulated components after having been analysed at right angles to the input polarisation plane.

3.2 Experimental Observation of SLMs

Experimental work using FLC SLMs was undertaken in a progression of systems of increasing sophistication. However, each has the T-shaped core optical system shown in figure 3.6

This arrangement is necessary for a reflective modulation element; a polarising

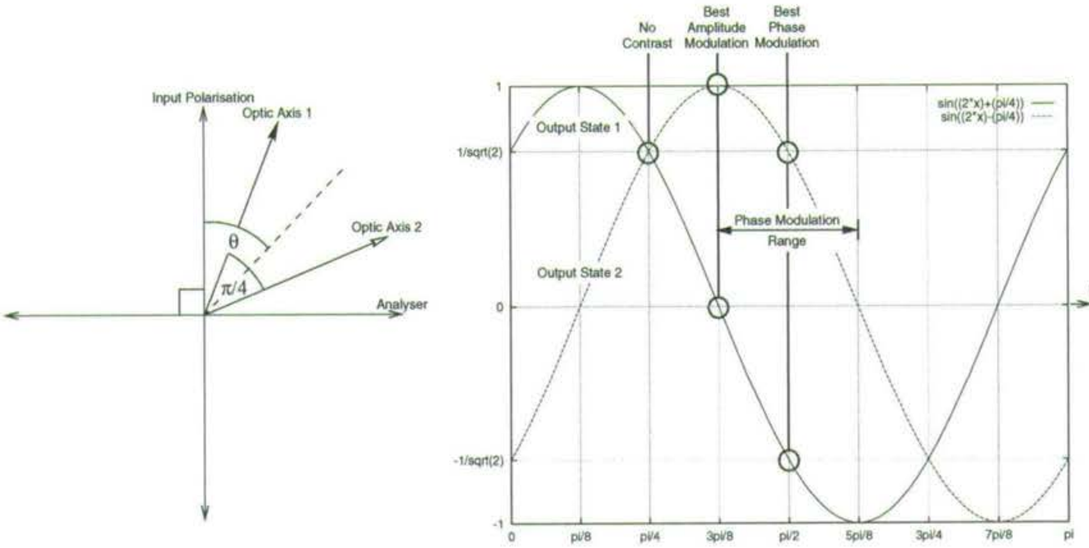

Figure 3.4: Sinusoidal (Malus' Law) variation in the two possible output states from an FLC half-wave plate against optic axis direction θ. Negative values indicate a π phase difference.

beamsplitter is used to separate the incident and reflected SLM beams, and also serves as crossed polarisers. A half-wave plate is used to rotate the plane of polarisation of the incident beam in order to align it at the correct angle with the FLC optic axes. The SLM is carried on a header PCB which can be mounted on the optical bench, as in figure 2.6. Illumination is generally a collimated coherent beam, for phase modulation systems. Incoherent illumination has a limited use, such as in a mapping of an array of input light channels to a binary *amplitude* modulator.

Both the 256^2 and the 512^2 SLM were used as dynamic phase modulators in experimental systems. The 256^2 SLM system was less complex in two main respects, both due to the problems of the 512^2's DRAM addressing: electronically, less interfacing is required, and optically, the 256^2 can be viewed under constant illumination without contrast reversal or blanking frames. The move to the higher resolution device was made in anticipation of the availability of 512^2 SLMs, and in order to evaluate this new DRAM device as a coherent phase modulator. The failure of a planarised wafer, as mentioned earlier, was an unfortunate setback to

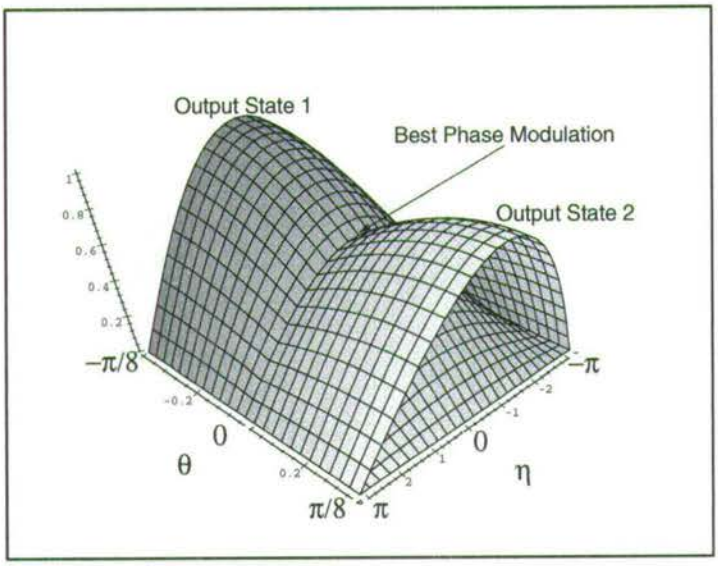

Figure 3.5: Variation in intensity of the two output states from an FLC cell against optic axis direction θ over the $\pi/4$ range of the switching angle, and birefringence depth η for the half-wave to full-wave plate range.

further work. It is anticipated that the 512^2 SLM will be the device available for future coherent systems work, in common with the experience of researchers who take their modulator devices from the miniature display industry.

3.2.1 System Configuration

The final system is shown schematically in figure 3.7, and the SLM mount in the photographs in figure 2.6. One of the salient features of this system, when compared to earlier versions, is the addition of an acousto-optic switch to shutter the laser in synchronisation with the display of frames on the DRAM 512^2 SLM. The AO switch uses acoustic waves in a crystal to diffract the incoming beam out of its zero order, when a signal is applied to the transducer. The device is specified by the manufacturer to diffract 80% of incoming light into the 1st diffraction orders. However, rather than use a diffracted order I take the zero order through the system. This gives a 4:1 shuttering contrast without the alignment problems

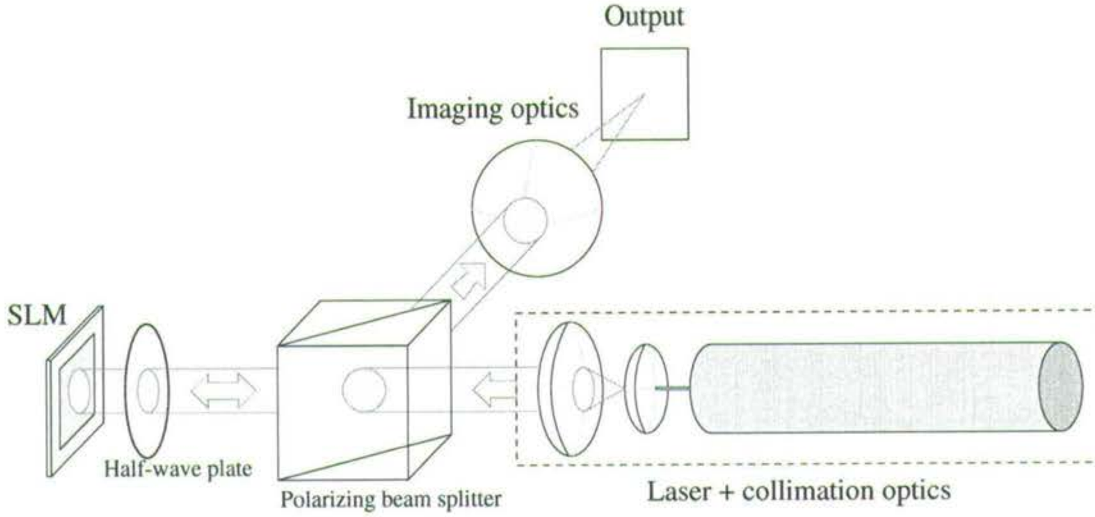

Figure 3.6: The basic system for illuminating and capturing an image from an SLM

of an off-axis beam. For phase holography, Babinet's theorem means that the inverse SLM frame is equally as useful, but the AO switch is still used in order to damp out the blanking and switching portions of the SLM addressing cycle. The AO switch receives one of the LED driving signals from the SLM interface.

3.3 Capture of the SLM Thickness Profile

It is possible to capture an image of an SLM when mounted in the system, in particular to assess the type and severity of defects. The thickness profile of the cell can be estimated from readily observed fringes due to reflection interference in the cell and the birefringence of the FLC. The defect modes of FLCoS SLMs are illustrated by the 256^2 SLMs in figure 3.8, and by the two faulty 512^2 SLM examples whose active areas are shown in figure 3.9. These two devices are referred to as SLM 1 and SLM 2 from here on, as indicated in the figure.

Only SLM 1, of these two 512^2 devices, is fully working electrically; both have poor optical properties due to typical (cell profile and FLC quality) faults. SLM 1 is a severely wedged cell, with two superimposed fringe systems visible. SLM 2 has a

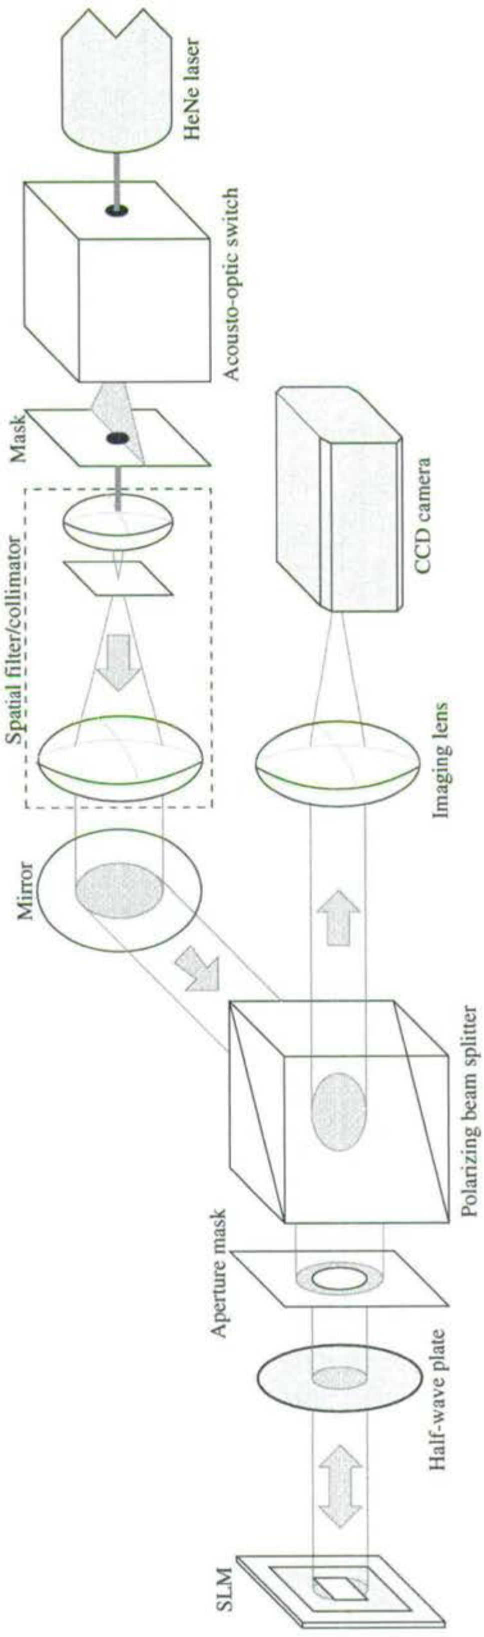

Figure 3.7: A schematic of the phase modulation system as used

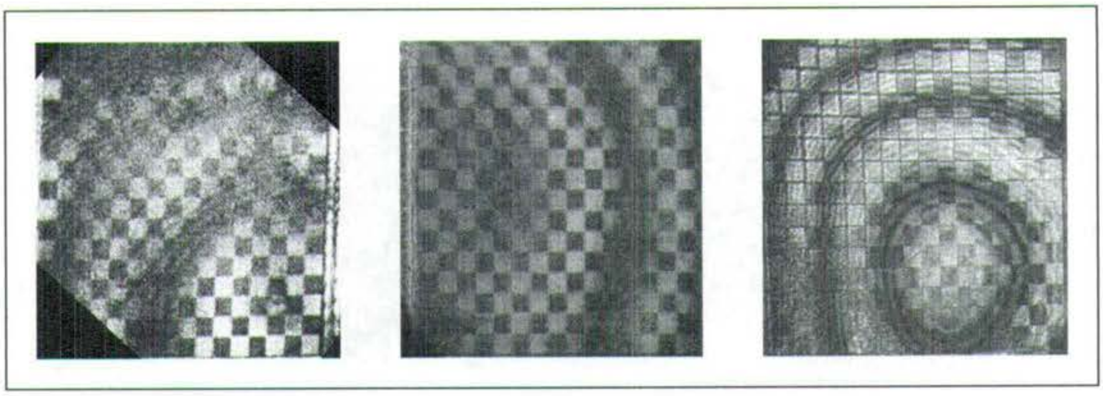

Figure 3.8: 256^2 SLMs displaying a 16^2 checker pattern; note the fringing and contrast reversal caused by varying cell gap.

warped cell profile and low-contrast FLC, perhaps due to insufficient outgassing during filling. Between them they serve to highlight the problems described previously. SLM 1 was filled using the technique described in chapter 2, and shows improved LC modulation quality. Both SLMs have $\approx 64^2$ areas usable for monochrome display.

It is possible to deduce the approximate thickness profiles of these cells by analysing the fringe patterns in figure 3.9. However, in order to gain a clear single set of interference fringes two further methods were tried with some degree of success.

3.3.1 Fizeau Interferometry

A transmissive optical flat was introduced into the system immediately before the SLM, with the intention to produce fringes due to the profile of the SLM backplane, as in a Fizeau interferometer. Obviously, any fringes so produced are superimposed on (multiplied with) the existing image. Hence the new image was divided, pixel by pixel, by the image taken before the introduction of the optical flat. This should leave only the Fizeau fringes; the result for SLM 1 is shown in figure 3.10. This method suffers due to the low contrast of the fringes, and moreover the many sources of spurious fringing in the system. Other

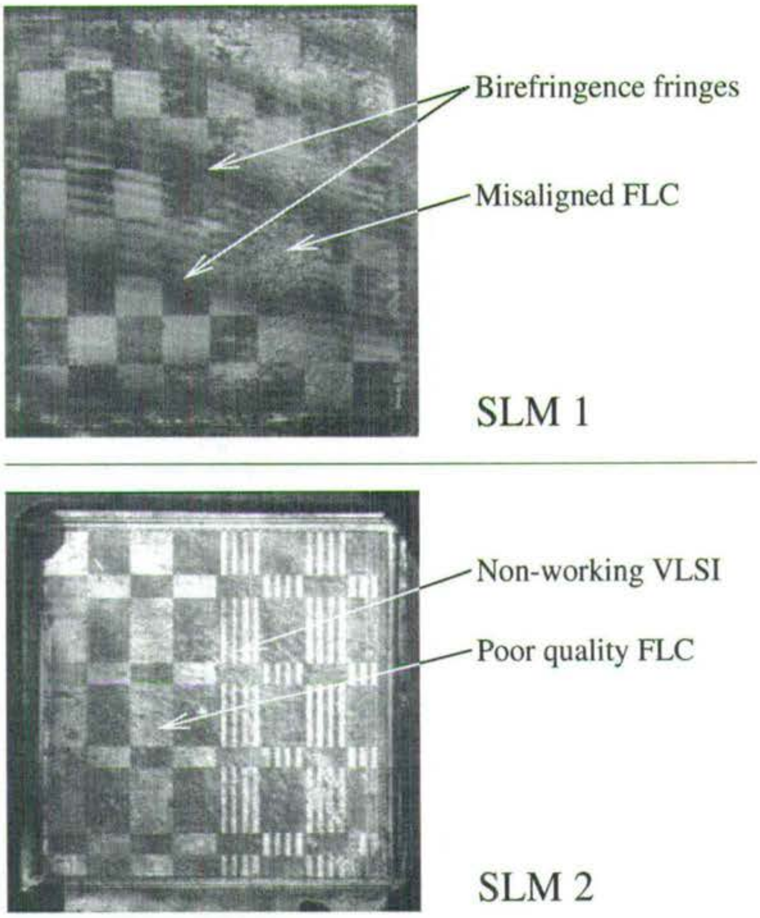

Figure 3.9: Two severely defective 512^2 SLMs, each displaying a 64^2 checker pattern.

interferometer configurations may have produced clearer results with a similar technique; in search of a simpler approach, the SLM was removed from the system and illuminated directly.

3.3.2 Capture and Use of Zernike Coefficients

There is no need for an external reference flat to gain an approximate thickness profile of a cell which is sufficiently wedged — the internal reflection fringes will serve. Further, a standalone interferometer with the reference flat removed provides good illumination and image capture for this purpose. Interference arises


(a) Original (b) Interferogram (c) Difference

Figure 3.10: An attempt to create an interferogram by adding an optical flat to the SLM system.

between the backplane and the front glass surface, since the glass/LC interface is more nearly index-matched. Internal fringe patterns were thus obtained for the two SLMs using an Interfire benchtop interferometer. This method relies, ironically, on having rather poor cells: we would hope not to have enough fringes for it to work well (~10 fringes). It also give no absolute measure of backplane profile, although the cover glass can be taken as a reference flat as a reasonable approximation[1]

Order	Name	SLM 1	SLM 2
1st	Focus	0.94	-0.24
	x Tilt	12.53	0.8
	y Tilt	-3.19	-1.01
3rd	Spherical	0.07	-0.02
	x Coma	0.99	0.35
	y Coma	0.07	0.2
	Astigmatism	0.02	-1.07
	45° Astigmatism	-1.82	0.53

Table 3.1: Zernike coefficients for the cell thickness profiles of two 512^2 SLMs, taken from internal interference patterns.

---

[1] The cover glass on the 256^2 SLMs yields thickness interference fringes in its own right; that on the 512^2 SLMs does not.

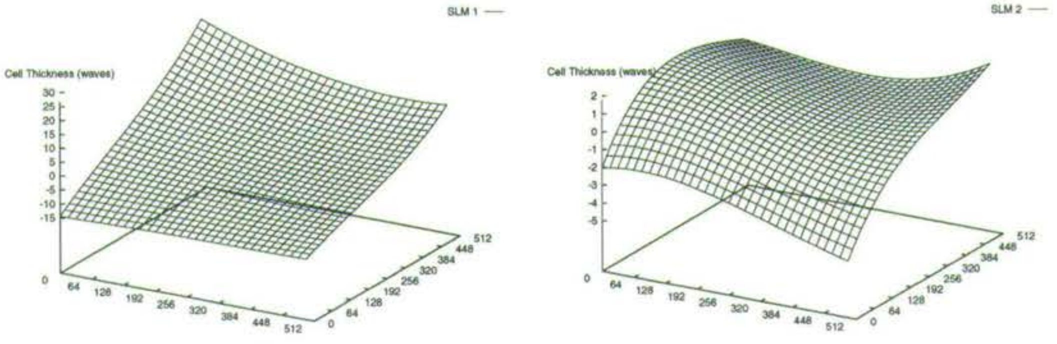

(a) (b)

Figure 3.11: Surfaces generated according to captured Zernike coefficients from (a) SLM 1 and (b) SLM 2. These correspond to thickness profiles of the cells.

To analyse the fringes I use the software package Apex[40]. The software relies on picking out fringe centres, and runs a best-fit algorithm to obtain Zernike coefficients[41]. These are shown in table 3.1 for the two SLMs, and surfaces reconstructed from these coefficients are shown in figure 3.11. Apex returns the surface parameters in terms of waves, where the interval between two adjacent fringes is a height variation of one wave, that is, λ of the (HeNe) laser source used by the interferometer. In chapter 5 these results will be used to adapt phase modulation masks to suit individual SLM properties.

Neither of these cells complies with the $\frac{\lambda}{4}$ limit as described above. However, in persevering with these captured coefficients in chapter 5, we will provide a good test of the ideas in the following chapters.

3.4 Summary

The aim was to classify defect phenomena of FLCoS SLMs and, if possible, quantify their effects; the concept was of an SLM as a complete device. While FLC defects — and the failure modes of VLSI — have been described elsewhere, the practical implications on SLM operation have been explored here.

To that end a Jones Matrix model of a birefringent cell was used to investigate phase modulation in a cell of varying thickness. This produced the interesting result that not only does birefringence depth affect the polarisation state of transmitted light, it introduces a phase offset which depends on the cell thickness. For a cell between crossed polarisers, spatially varying thickness produces spatially varying amplitude and phase in the output beam. Obviously, this is a complication for a phase modulator.

Further problems arise from a spatial variation in the direction of the effective optic axis, due to misalignment of FLC. The sinusoidal variations of the two polarised output states of FLC were shown superposed, for a varying optic axis direction. The phase- and amplitude-modulation regimes lie on this range, separated by an angle equal to the FLC cone angle. Over half the range there is no phase separation between the states, and thus the SLM does not function as a phase modulator.

An optical system to illustrate these effects in live SLMs was set up, and captured images show the three classes of SLM defect: electrical, crystal, and cell thickness. Quantifying the severity of defects for a given SLM is more difficult. Cell thickness information was sought by interferometry, initially with an external reference flat. Due to the difficulty in producing useable interference fringes this way, the cover glass itself was later used as a reference. The fringes captured in this way yield Zernike coefficients which describe — not a surface — but the thickness profile of the cell. Once this is known, the effect on specific SLM operations can be estimated, and this will form the subject of chapter 5.

Chapter 4

Binary Phase Hologram Generation

Here was another task ideally suited to the computer's particular form of high-speed idiocy
James Gleick, *Chaos*

4.1 One-to-One Data Connection

The most straightforward use of a FLC SLM as a data switch is to map a spatially separated group of optical data channels directly onto the pixel array, and then to address a given pixel (or SLM area) in order to turn the corresponding channel on or off, shown in figure 4.1. Thus a N^2-pixel SLM can modulate a maximum of N^2 channels. This approach is, clearly, highly vulnerable to SLM pixel drop-out, alignment errors and cross-talk.

Consider a hypothetical system in which each spatially separated channel in a 2D array represents a different transmitted *value*; the effects of alignment and cross-talk errors can then be reduced by using a coding strategy which groups data channels such that neighbouring pixels represent similar values. One such scheme is the *Karnaugh Map* employed in digital system calculations.

The Karnaugh Map is a means of encoding binary values on a 2D grid such that neighbouring values differ by the value of only one bit. The most common

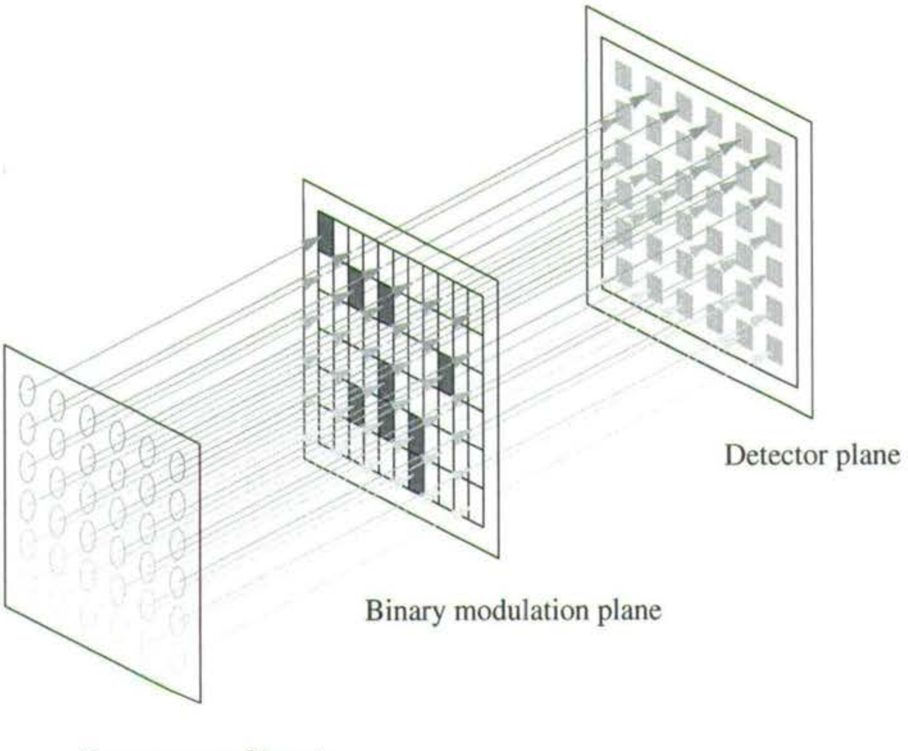

Figure 4.1: An SLM can act as the mask in a one-to-one mapping of emitters and detectors

example is a 4x4 grid which encodes all possible 4-bit numbers. As each grid position has 4 neighbours (treating boundaries as periodic), the neighbours of a 4-bit value are those formed by flipping each of the bits. The scheme for placing values in the grids is as shown in figure 4.2.

In the Karnaugh Map, value is coded by 2D position; this can be implemented optically as a data encoding scheme using 2D components. First the Karnaugh Map scheme is expanded to cover a wider range of values.

For 8-bit numbers (i.e. bytes) we will need a 16x16 grid, and each position will not have enough direct neighbours to hold all eight 1-bit different values. Since we (Matthew Hart & PJS; see [42]) had a 16×16 input device (an array of LEDs) we generated a pseudo-Karnaugh Map scheme, analogous to the 4-bit version, to encode 8-bit values. This is shown in figure 4.3. In our scheme the 4-bit map is

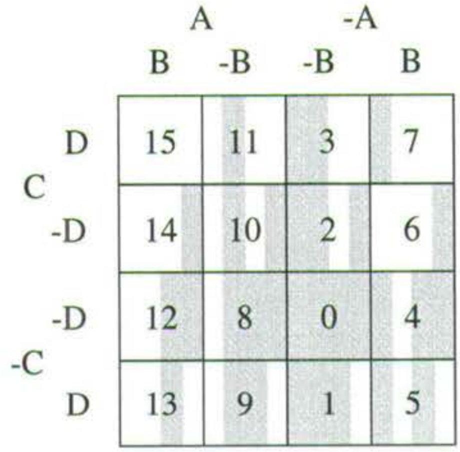

		A		-A	
		B	-B	-B	B
C	D	15	11	3	7
	-D	14	10	2	6
-C	-D	12	8	0	4
	D	13	9	1	5

Figure 4.2: The Karnaugh Map for a 4-bit number $ABCD$.

copied and reflected for each of the two extra bits encoded on either axis.

By encoding values on a 2D array in a Karnaugh Map scheme, a degree of *fuzziness* is introduced in the data values which pass through the system, in that areas of the array represent similar values. This positional encoding degrades gracefully with small alignment errors, and a neighbouring channel may be detected downstream due to crosstalk with similarly-valued modulator areas. This might be a useful property of a 'fuzzy-logic' switch, for example. A prototype system featuring this property, using a FLCoS SLM, was built.

The SLM-based incoherent illumination system used is shown in figure 4.4. A miniature 16×16 LED array was used as a binary 2D input device, mounted on a dedicated header board and interfaced from a PC. When fed a stream of bytes, the Karnaugh Map driving program on the PC illuminates the corresponding LED for each in turn, according to the pseudo-Karnaugh Map scheme. The LED array was imaged onto a 256^2 SLM so that when turned on, each LED illuminated a SLM pixel. We set one SLM pixel 'on' to allow light at that position to propagate through the system, and then 'looked' for it with byte values fed to the software, using a detector as shown. The output from the two detectors in figure 4.5 shows how a value is 'recognised' when an LED and its corresponding area of the SLM are activated simultaneously. In this example a matching and a non-matching

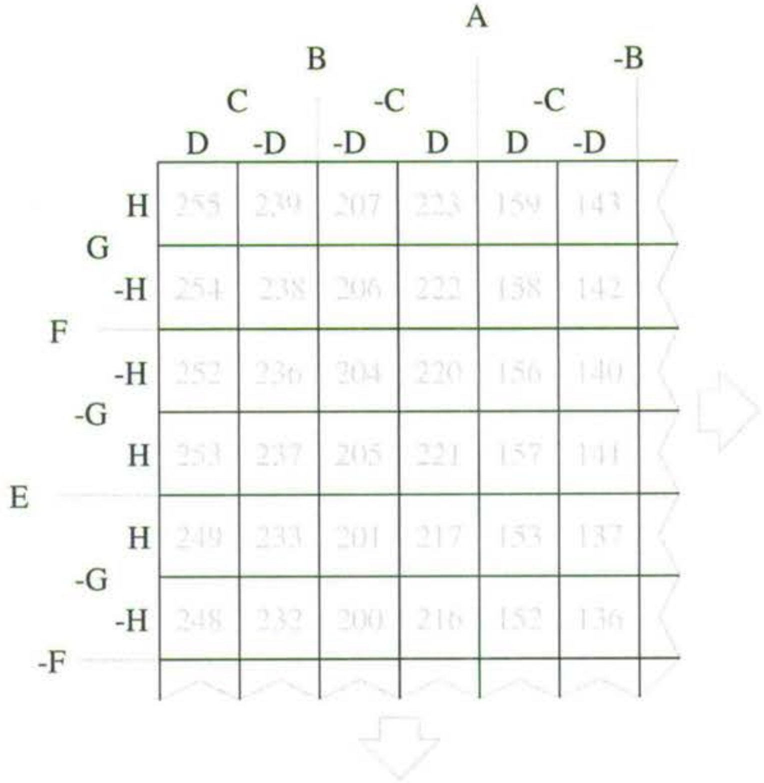

	D	-D	-D	D	D	-D
H	255	239	207	223	159	143
-H	254	238	206	222	158	142
-H	252	236	204	220	156	140
H	253	237	205	221	157	141
H	249	233	201	217	153	137
-H	248	232	200	216	152	136

Figure 4.3: Part of the pseudo-Karnaugh Map for a 8-bit number $ABCDEFGH$.

pixel are activated alternately. It is clear from the trace that this system suffers from a low contrast; this is due in part to the low light level of the illumination device, which results in a low signal-to-noise ratio for the system.

Thus, we have an individually switchable set of 256 incoherent free-space optical channels, or, more broadly, a 1-to-1 mapping of a binary source array to a binary modulator array. This would in itself be sufficient for a binary 'Hebbian' neural system, where a binary 'weight plane' correlates input vectors with output vectors — also known as a vector-matrix multiplier. Moreover, consideration of this incoherent system points up some general concerns of practical SLM data switching. It is clear that the *contrast ratio* of the device is crucial if we use the SLM for amplitude modulation, in order to discriminate on and off states; also, that the linear mapping approach is unavoidably vulnerable to modulator and system defects. For these reasons the phase-holography approach is greatly

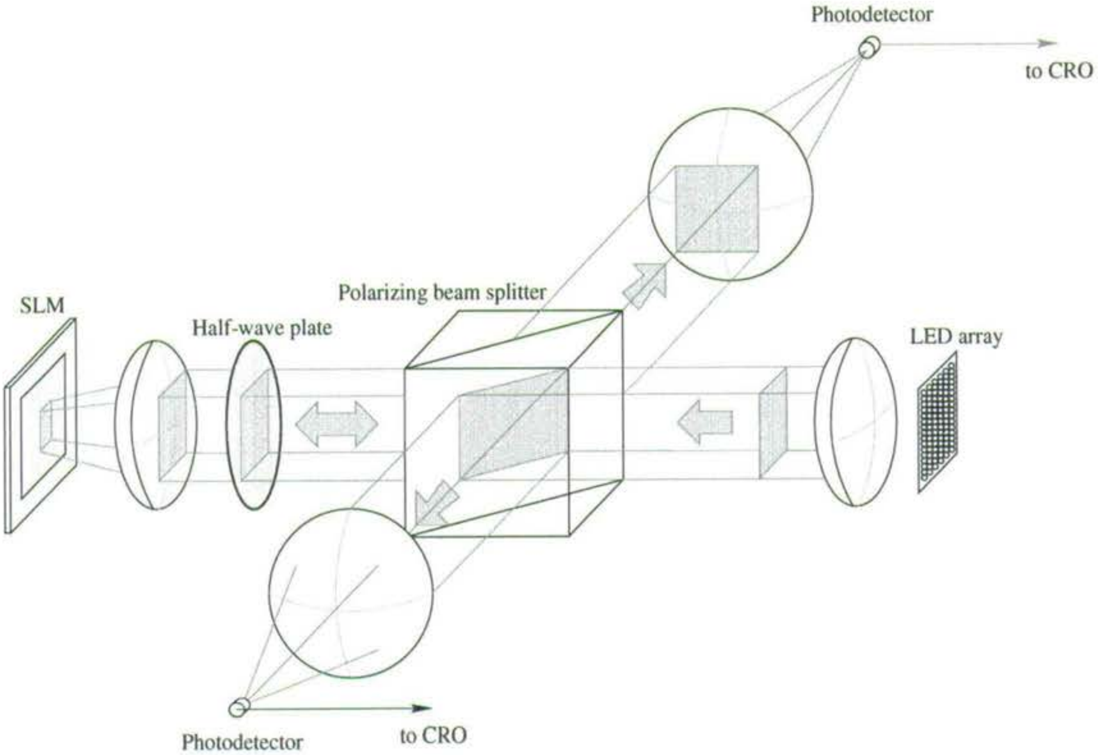

Figure 4.4: Optical system for 8-bit Karnaugh map demonstration

preferable, even though it necessitates the move to coherent light sources.

4.2 Holographic Fanout System

With a pixellated 2D phase modulator, we can impose phase control over an incident wavefront. We calculate the required phase adjustment for each of a sampled set of points in the modulator plane. These are the phase modulation depths to be used as pixel values by the modulator.

In practice, spatial phase modulation can be achieved by either optical thickness, as in say, etched fused silica[43] or silicon nitride[22], or by restricting the optical path, as in detour-phase holography[44]. Note that the latter — as in a zone plate — is an amplitude modulation technique in the sense that it blocks certain portions of the incoming light.

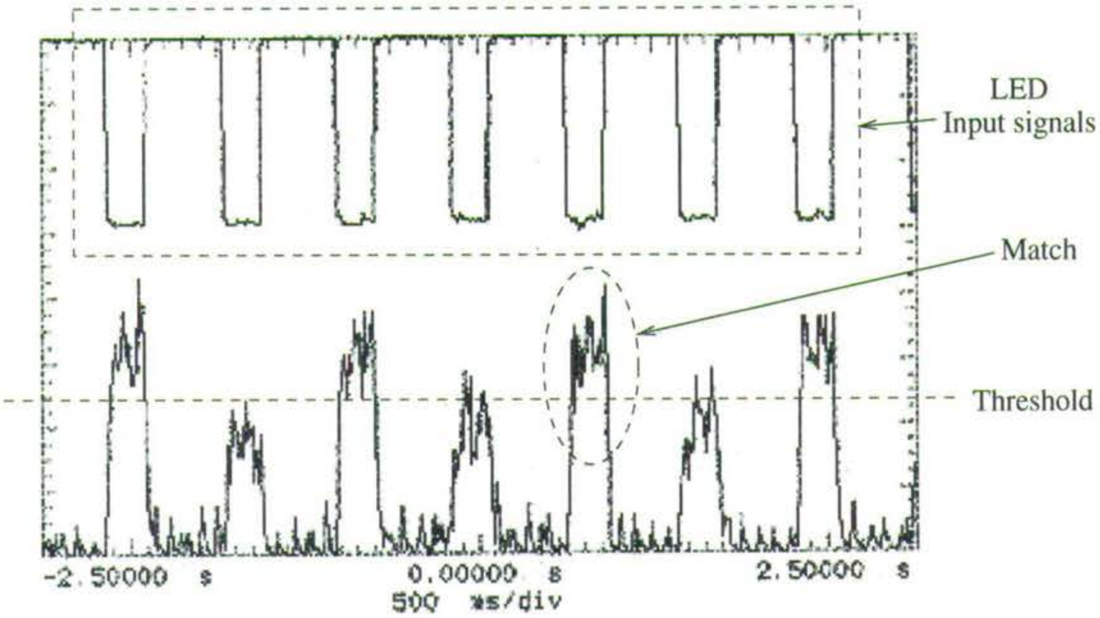

Figure 4.5: Coincidence between a lit LED and an SLM 'on' pixel, seen from detector traces

Calculation of spatial phase masks is an example of computer-generated holography; we seek to apply a function to the wavefront which creates some given optical distribution downstream. In general such 2D functions, or the optical elements which implement them, are often termed *kinoforms*. Furthermore, in the case of the FLCoS SLM the restrictions placed on the final mask (i.e. binary phase-only pixels) limit the degrees of freedom available to a computational algorithm. As explained below, these constraints suggest an algorithm which proceeds in a simple iterative fashion. I will use the term *binary phase hologram (BPH)* to mean a 2-level phase-modulation function such as may be displayed on a FLCoS SLM.

The computed BPH is effected in practice by an optical diffraction element, which uses the phase-modulating 2D pattern to generate some given optical intensity distribution upon diffractive evolution of the phase-controlled light — especially by then using a converging lens to perform an effective Fourier transform. This so-called *Fourier holography* can be treated as a straightforward application of Fourier (scalar wave) optics. Moreover, a phase — as opposed to amplitude —

2D hologram can be optimised for high optical efficiency, and thus is an obvious application for a switchable phase modulator such as a liquid crystal SLM. Widely-used nematic LC SLMs, such as cannibalised LCTV display elements[45] can modulate in an *analogue* 'phase-mostly' mode, which also entails some amplitude modulation. FLC, however, can (ideally) be used as a *binary* phase-only modulator. Although the binary nature of the phase modulation imposes some restrictions on the possible kinoform output, FLC has the advantages of fast switching speed and simplicity of modulation over a nematic device.

4.3 Fourier Phase Holography

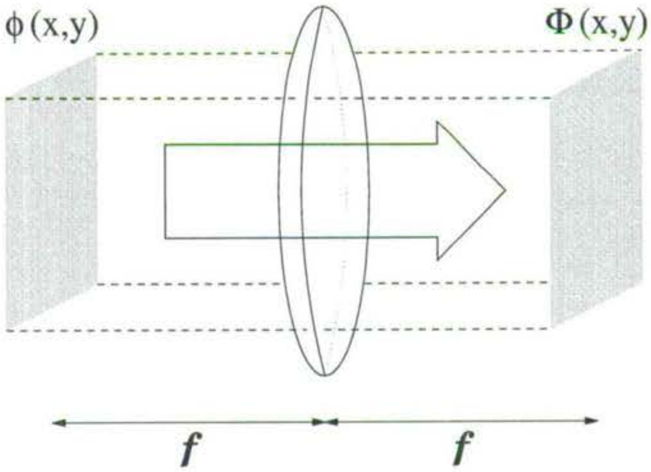

Figure 4.6: A scaled Fourier transform, Φ, of a spatial distribution $\phi(x,y)$, at the front focal plane of a converging lens, is seen at the rear focal plane.

Fourier holograms take advantage of the Fourier-transform (FT, $\mathcal{F}$) properties of a lens; a scalar diffraction treatment of plane wave propagation through a converging lens shows that for an input (complex) distribution $\phi(x,y)$ at the front focal plane, the scaled Fourier components of the input pattern appear at the 'Fourier plane', one focal length f behind the lens. As indicated in figure 4.6, it can be shown[37] that the complex distribution $\Phi(x,y)$ at the Fourier plane, of an original function $\phi(x,y)$ at the front focal plane, is given by:

$$\Phi(x,y) = \frac{A_0}{\imath\lambda f}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\phi(u,v)e^{-\imath\frac{2\pi}{\lambda f}(ux+vy)}\,du\,dv \tag{4.1}$$

Thus, by focussing the phase-modulated reflection from an SLM using a lens placed one focal length downstream, we can generate the FT of the binary phase pattern displayed on the SLM. From this it is clear that the main effort in usefully applying this result will be in efficiently finding binary phase patterns which transform to some required output distribution. Optical interconnection schemes which need a 'broadcast' one-to-many connection are feasible with this approach because of the convolution inherent in the FT: any modulation (spatial or temporal) on the input is convolved with the FT of the phase mask at the output plane. If the phase mask transforms to a comb of isolated points the effect is to copy the input to many spatially separated outputs. This might form the basis of a reconfigurable optically-connected parallel architecture, for example. With this in mind, the examples of output targets used below are 2D combs or 'spot arrays'. A generic version of this strategy is shown in figure 4.7.

Two measurements are taken for the quality of output from a kinoform: the proportion of output light that is diffracted into the desired output pattern ('efficiency'), and the deviation from the desired intensity distribution in the output pattern. Since the 'fan-out' kinoforms concerned here are designed to diffract equal amounts of input light into each output spot, we may term the latter measure 'uniformity'. Specifically, uniformity is calculated as $\frac{max-min}{mean} \times 100$, i.e. the maximum variation as a percentage of the mean.[1]

Note that the output target is defined only as intensity; there is no attempt at full complex control of the output. Obviously with a finite set of binary phase masks available, there is a corresponding finite set of complex Fourier transforms. However, the feasibility of generating a quantised solution for an arbitrary complex target is doubtful. Wyrowski[47] has considered the question of imposing amplitude and/or phase requirements on the target output, using

---

[1] This is a different definition from Samus[46], for example.

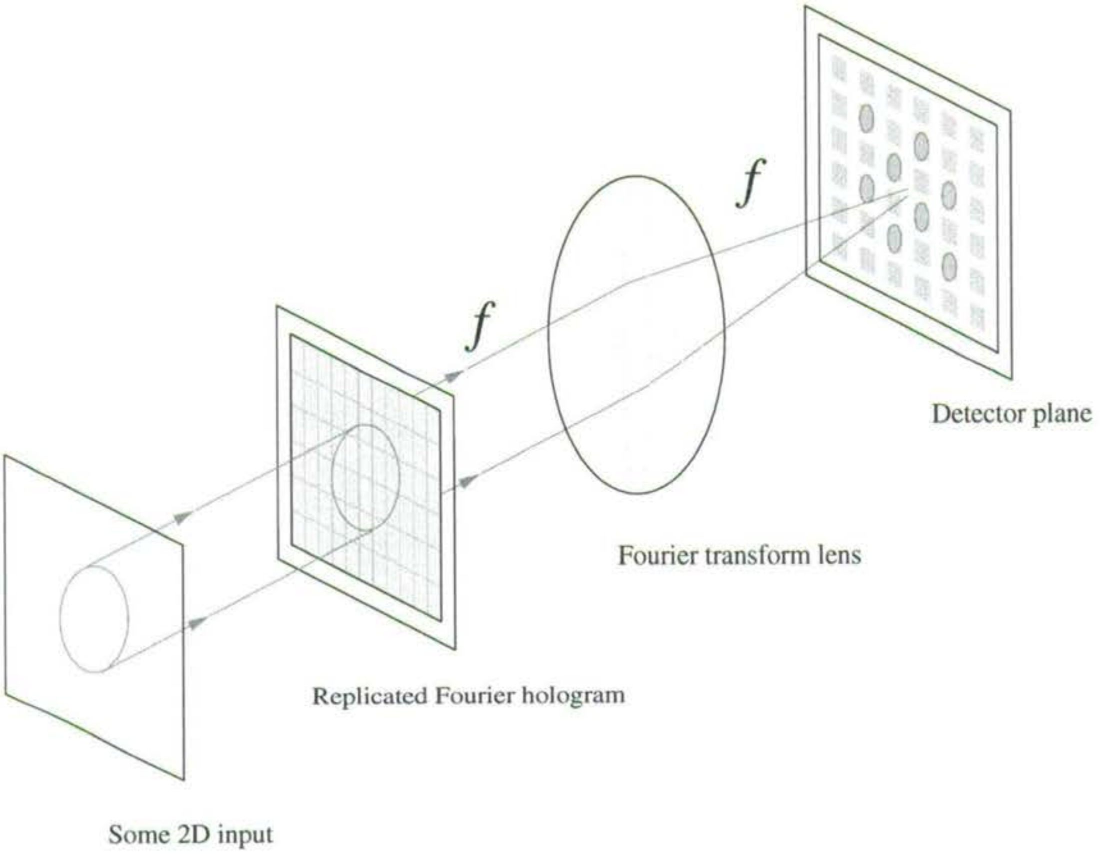

Figure 4.7: Holographic interconnection

analogue phase levels during calculation. In the following, the output target will be treated solely as an *intensity* distribution, such as would be detected by the receivers in a hypothetical interconnection system, and the phase mask will be treated as a binary ($\pm\pi$) modulator such as a FLCoS SLM.

In a typical system based on this approach, the SLM displays some phase modulation pattern at the front focal plane of a converging lens. The bottom line is that to achieve a given intensity distribution $I(x, y)$ in the lens' Fourier plane, we need to design a binary phase pattern $\phi(x, y)$ whose Fourier transform squared is as near as possible to a scaled version of $I(x, y)$. That is, we require:

$$A\left|\mathcal{F}(\phi(x,y))\right|^2 \rightarrow I(x,y) \qquad (4.2)$$

with an arbitrary scaling constant A.

Hence, the question arises: *is* there a binary phase-only solution for the inverse FT of an arbitrary intensity pattern, and if so, how can it be found? This is an example of a so-called *phase retrieval* problem, such as is encountered in adaptive optics. Hayes[48] notes that not only is the phase part of the Fourier transform sufficient in itself to reconstruct the essential features of an image, the phase of the FT can even be binarised before reconstructing a recognisable version of the original. He concludes that the image information is encoded almost entirely in the phase of the FT. While this is so, a binary phase image necessarily transforms to an intensity distribution which is symmetric about the origin. Hence:

1. Any asymmetric intensity distribution to be generated by BPH needs to be off-axis to avoid confusion with its reflection.
2. An off-axis asymmetric target can have at most half the input light directed into its reconstruction, since it will be replicated by the BPH.

For intensity targets which are symmetric about the origin, we can immediately double the possible light efficiency of a BPH, since the replicated order contributes to the desired output. Moreover, we need only calculate the BPH for half of the target, which speeds up the process. For true asymmetric output, however, more phase modulation levels are needed.

Among the techniques which have been applied to the BPH design problem are two families of iterated algorithm. The **Gerchberg-Saxton** method[49] repeatedly Fourier transforms the plane to and from the Fourier domain. It applies the phase-only constraint to the BPH and keeps desired portions of the output. **Simulated annealing**[50], by contrast, applies *single cell* changes to the BPH and keeps them if the change reduces some 'cost' function applied to the Fourier transform. Additionally, some random changes may be allowed with a decreasing probability. The stumbling block in each is the computationally-intensive operation of performing a 2D Fourier transform. Gerchberg-Saxton avoids this by requiring relatively few iterations. However, Gerchberg-Saxton does not provide a solution for quantised phase levels. Simulated annealing must, in principle, perform a full Fourier transform upon each pixel change, and this is where BPH design diverges from the well-known Ising model for a 2D binary array of states. Dramatic speed-ups have been achieved in the Ising model by switching groups of cells etc.[51], but in the BPH case it is the *evaluation* of the pixel change which slows the process. There is, as will be described below, a short cut in this evaluation; even so, only one pixel is switched on each iteration.

4.4 Superposition of Gratings

It would be much easier if we could find a set of BPH basis functions from which to build BPHs, and thus optimise or even circumvent the iterative design process. A naive assault, based on those phase masks which *can* be calculated analytically, goes as follows. The SLM has a 2D pixellated modulator array. We can treat the BPH and its FT as pixellated 2D arrays of side n, and use the discrete form of the 2D Fourier transform:

$$\begin{aligned} F(u,v) &= \mathcal{F}\left[f(x,y)\right] \\ F(u,v) &= \frac{2\pi}{\lambda n^2} \sum_{x,y=-\frac{n}{2}}^{\frac{n}{2}} f(x,y) e^{-i2\pi\left(\frac{ux}{n}+\frac{vy}{n}\right)} \end{aligned} \tag{4.3}$$

Any given output intensity distribution is a pattern of discrete pixel values, $I(u_i, v_j)$. If each non-zero output pixel is treated as a separable output, it is straightforward to *calculate* a set of phase gratings which generate each pixel individually. The equation of a phase diffraction pattern to generate a diffraction order at (a, b) in the Fourier plane is:

$$\phi(x,y) = \mathcal{F}^{-1}\left[A.\delta(a,b)\right]$$

$$= A.e^{\imath 2\pi(ax+by)} \tag{4.4}$$

where δ is the Kronecker delta function.

That is, the phase modulation is a sinusoidal wave with wave vector $\bar{k} = (a, b)^T$, which generates non-zero output, of intensity A^2, at (a, b) and $(-a, -b)$. This suggests a possible strategy for generating arrays of output spots: we might try to approximate a reasonable BPH by summing the (possibly binarised) phase gratings for the constituent non-zero pixels of the desired output, then normalising the sum between 0 and 2π — see [46]. There is a further complication of converting the result to a binary phase mask.

In practice this gives acceptable results for low numbers ($< 4^2$) of pixels in the output. To see how it breaks down, consider the square array of spots shown in figure 4.8. There are 8×4 non-zero pixels on a $1 : 1$ mark-space ratio in a 128^2 array. Each spot has a binary phase grating associated with it, given by binarising equation 4.4. The 32 patterns sum as in figure 4.9 to give a composite containing 33 grey levels. However, binarising this composite itself in order to use it as a BPH is a nonlinear operation that does not preserve phase information: because phase is cyclic the operation of adding several masks and thresholding the result is — in general — an invalid way of representing the composite operation of successively applying each mask. Neither the sum or its binarisation transform to give an acceptable output; the 8×8 pattern is buried in aliasing noise and most of the power goes to the DC spot. Even so, some success has been reported with small numbers of output spots[46]. Pixel-by-pixel operations such as error diffusion as an attempt to improve the mask would seem to defeat the point of operating at the large-scale, and bring us back to annealing algorithms anyway.

4.5 Iterative Generation of Phase Masks

In general there is no analytical way to calculate an optimal binary phase mask for a given output. It is necessary instead to adopt a controlled trial-and-error

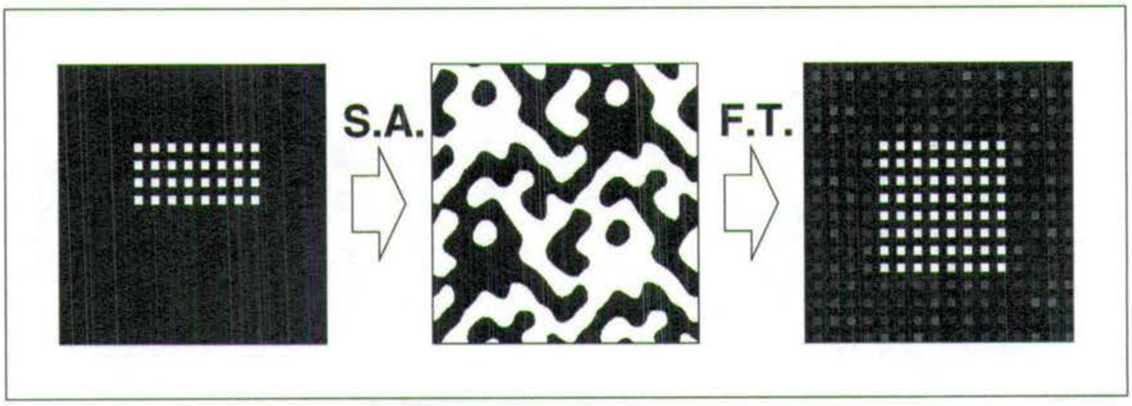

Figure 4.8: 8 × 8 spot array: target, binary phase hologram and generated output.

approach, by an iterative method which seeks to minimise some cost function by randomly varying the BPH pixels. An algorithm developed by Mark Dames[52] and used by Sergei Samus[46] goes as follows:

Start: A BPH $\phi(x, y)$ generates an intensity output $\Phi(u, v)$. Initially this can be entirely random.

1. Perturb ϕ and hence Φ by flipping the state of pixel (i, j) in $\phi(x, y)$ (i.e. change its phase modulation by π). This represents a change in $\Phi(u, v)$ of:

$$\Delta\Phi(u, v) = \left|e^{\imath 2\pi(iu+jv)}\right|^2 \tag{4.5}$$

2. To calculate how well $\Phi(u, v)$ corresponds to a desired output function $I(u, v)$, we consider the non-zero pixels in $I(u, v)$. These form a set of coordinates (u_i, v_i), and the differences between $\Phi(u_i, v_i)$ and $I(u_i, v_i)$ form the core of a cost function which determines the evolution of $\phi(x, y)$. There are several possible cost functions; in particular we can minimise the sum of the differences between $I(u_i, v_i)$ and $\Phi(u_i, v_i)$:

$$C_{sum} = \sum_i [I(u_i, v_i) - \Phi(u_i, v_i)]^2 \tag{4.6}$$

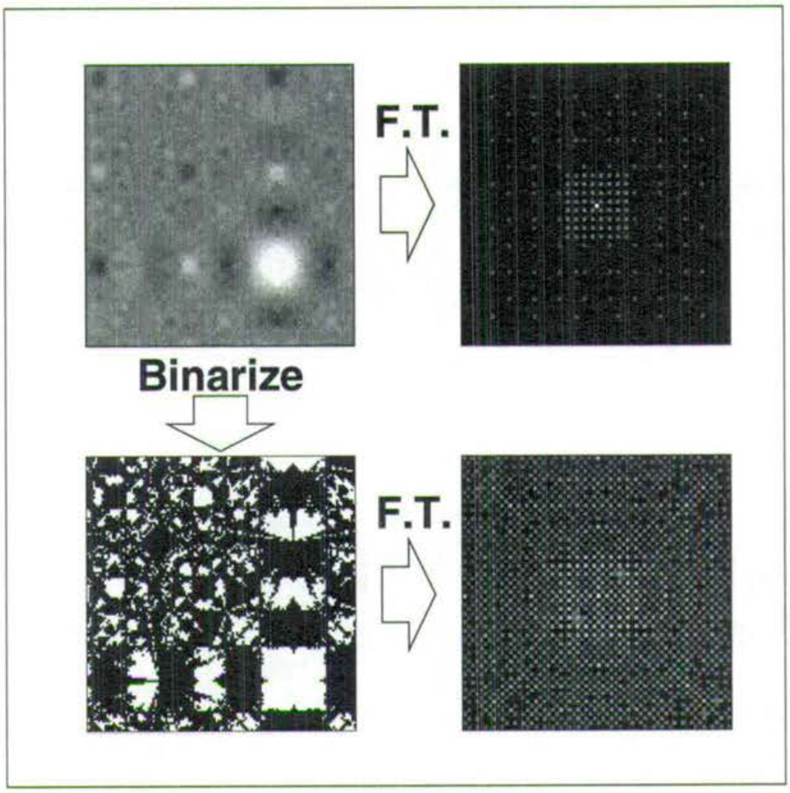

Figure 4.9: Superimposing the individual gratings for the spots in an 8×8 array: sum and binarised sum shown with Fourier transforms.

Alternatively, we can minimise the largest single difference between $I(u_i, v_i)$ and $\Phi(u_i, v_i)$:

$$C_{max} = \max_i [I(u_i, v_i) - \Phi(u_i, v_i)]^2 \tag{4.7}$$

3. Preserve the change in $\phi(x, y)$ if and only if C has been lowered by it.

The algorithm proceeds by flipping random pixels in $\phi(x, y)$, testing for a new value of C, and preserving the change in $\phi(x, y)$ if and only if C has been lowered. Over many iterations ϕ converges to a solution for the phase mask which generates a good approximation to $I(u, v)$.

'Simulated annealing' also implies some analogue of 'temperature' — the possibility of a random energy-increasing change in $\phi(x, y)$. This can knock the current

ϕ out of a local minimum and hence into a better solution. For instance, there are at least two solutions for the 4×4 fanout. The usual method is to 'cool' the BPH as the solution condenses. This was not found to be necessary in order to converge to a good solution from a randomised starting BPH.

This method does not conserve the total power in the running version of $\Phi(u,v)$. Parseval's theorem says that $\sum |f(x,y)|^2 = \sum |\mathcal{F}[f(x,y)]|^2$, but the algorithm only considers the change in a small subset of $\Phi(u,v)$ with each change in $\phi(x,y)$, and *keeps changes in* $\Phi(u,v)$ when C is reduced. The effect is a discrepancy between the running version of $\Phi(u,v)$ and the true FT of $\phi(x,y)$. In particular, although the algorithm may report an efficiency to within 1% of the 80% theoretical maximum[53], the efficiency seen from the FT of $\phi(x,y)$ may be $\sim$2% lower. Hence it is desirable to let the algorithm run but to perform a full FT at intervals, and especially when a solution is found. This also has the effect of removing the cumulative rounding errors in the running values of efficiency and cost, and keeps the algorithm 'on course' to a solution.

'Finding' a solution in this context means that the algorithm has not managed to lower the cost function in some given number of iterations. The code used to implement this algorithm is given in Appendix A. 512^2 holograms were generated on an SGI Indigo, working at 32-bit precision. The 4×4 hologram takes 2M iterations to converge to a stable solution, which takes 600 seconds of machine time. The tolerance (number of iterations without improvement) is arbitrarily set at 10000 iterations. A range of 256^2 BPHs is shown in figure 4.10, with generation times. Clearly, the annealing process stops when further improvements to the mask drop below the minimum change detectable at this precision.

Generation time, moreover, scales with N, the number of non-zero pixels in the output, and n^2, the array size. There are a number of speed-ups possible in the process, discussed in [46]. These tend to exploit symmetries in the BPH design and were not employed due to the asymmetric distortions described in the following chapter.

4.5.1 Pattern Replication

A BPH is periodic, and most of the time we need only calculate a single period of the pattern, as shown in figure 4.10. This is the resulting BPH for a spot array target with a mark-space ratio of 1: the irreducible pattern. However, when used in fanout applications this single period is replicated across the modulator area to produce a diffractive plane. Due to the Fourier transform relationship between BPH and output, smaller periods produce more widely-spaced outputs, which is usually desirable in optical systems. This 'replication' of the pattern is no more than a recognition of the frequency space inhabited by a Fourier hologram; the periodic nature of the hologram is implicit in the computational Fourier transform. In an optical system it is necessary to explicitly replicate the BPH over the light beam.

4.5.2 Height Coefficient

One of the factors determining the final efficiency and uniformity of an algorithmically generated BPH is the *height coefficient* h, representing the desired intensity of the output spots given that the BPH is necessarily $< 100\%$ efficient. Suppose we have a normalised output plane, i.e. $\sum \Phi(u,v) = 1$. Then, for a 100% efficient BPH generating N *equally-weighted* spots, $\Phi(u_i, v_i) = (\frac{1}{N})^2$. However, the maximum efficiency we can expect from a BPH is $\approx$80%, so set e.g. $h = 0.8$ and C becomes:

$$C_{sum} = \sum_i [(\frac{h}{N})^2 - \Phi(u_i, v_i)]^2 \tag{4.8}$$

In effect, we are asking the algorithm to generate a BPH which supplies spots with height h. This allows an investigation of the efficiency-uniformity tradeoff as h is varied through achievable heights and upwards. Values of $h > 0.8$ cause the algorithm to sacrifice uniformity to achieve maximal efficiency i.e. to channel intensity into a subset of the output spots, even though efficiency cannot climb more than a few % beyond 80% for any h. Values of efficiency and uniformity against h for generated BPHs are shown in figure 4.11 for the 4×4 and 8×8 targets. It can be seen that there is an optimum h at about 0.85 where a high efficiency coincides with a good (i.e. *low*) value for uniformity.

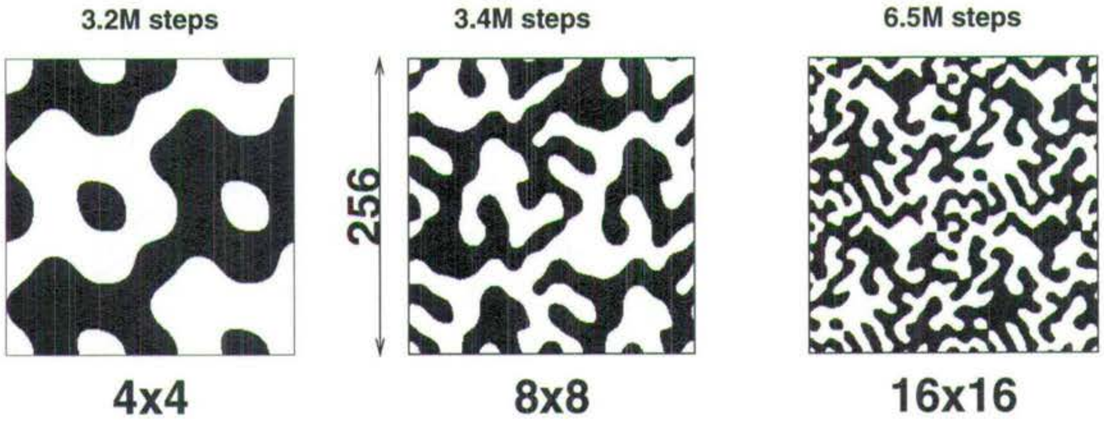

Figure 4.10: Binary phase holograms for (a) 4×4, (b) 8×8 and (c) 16×16 arrays on a 256^2 grid.

Therefore, the best BPHs are obtained from the use of a cost function which is attempting to generate a phase mask which has an efficiency slightly higher than that which is, in fact, achievable. This suggests a procedure for generating progressively better holograms by incrementing the height coefficient. For an achievable efficiency, the algorithm effectively stops when the computer 'runs out of decimal places' — the evaluated cost function drops to a lower value than can be represented. If the height coefficient is increased slightly the cost function value becomes significant and is again reduced. At some value of h the cost function will not vanish, and the unchanging cost function value stops the iteration after some (user-defined) number of steps without a change. We can expect further increases in h after that to worsen uniformity without marked efficiency improvement.

Comparison of these algorithmically-generated BPHs with those from the 'sum-and-threshold' method above shows a similarity between the shapes of the phase patterns. It is clear that the complexity of the pattern must be equivalent, and the general form is too. One notable difference is the asymmetry of the repeated

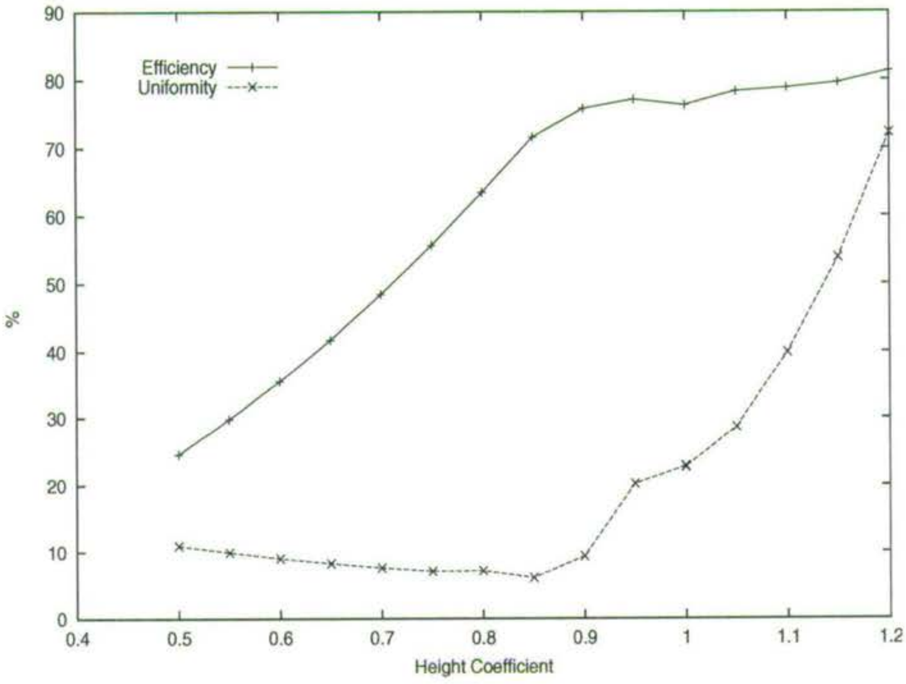

(a)

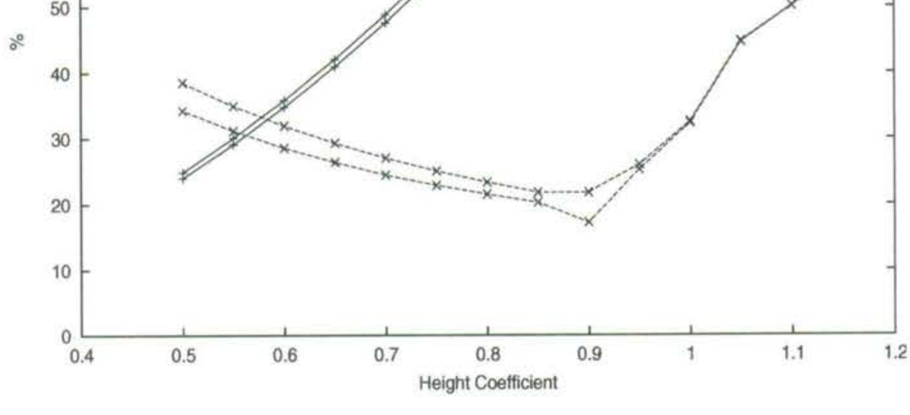

(b)

Figure 4.11: Efficiency and uniformity of generated BPHs depend critically upon the height coefficient h; (a) shows results for the 4×4 target, and (b) two runs for the 8×8 target.

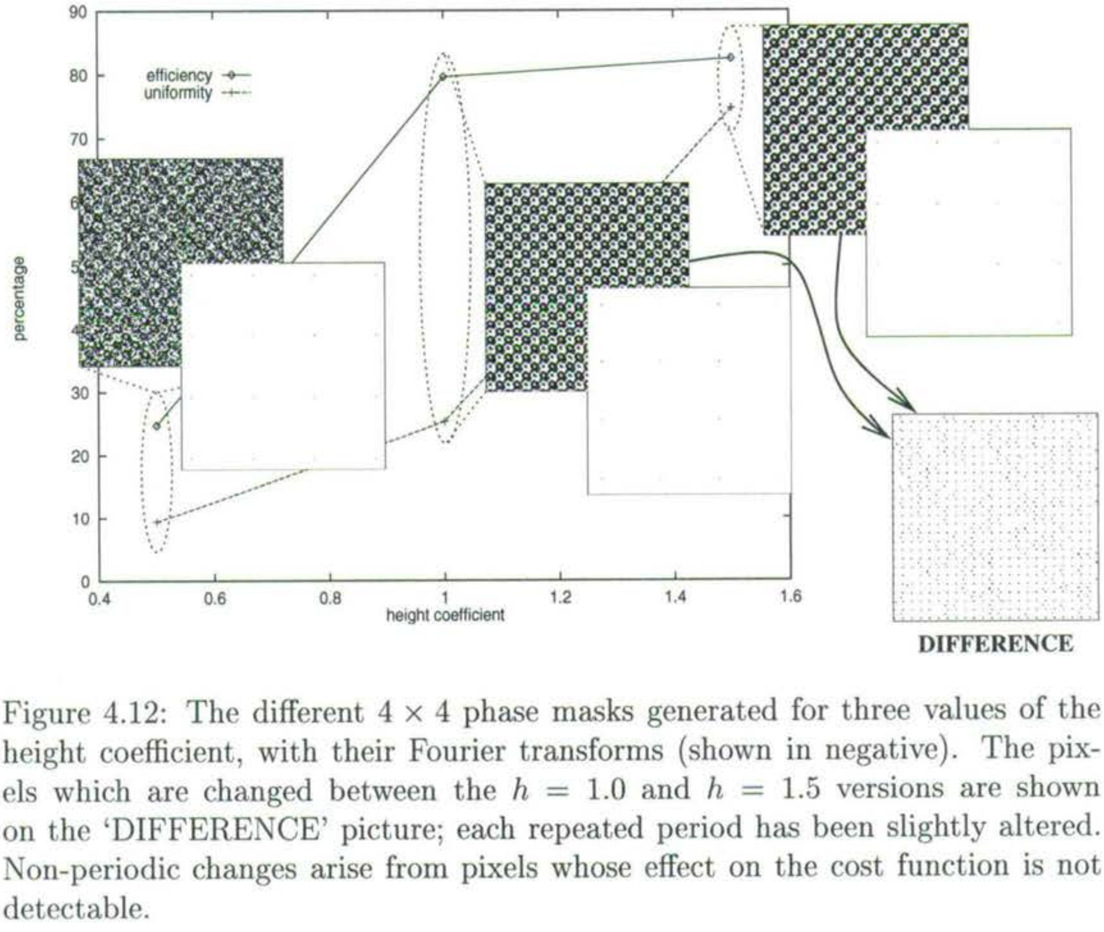

Figure 4.12: The different 4 × 4 phase masks generated for three values of the height coefficient, with their Fourier transforms (shown in negative). The pixels which are changed between the $h = 1.0$ and $h = 1.5$ versions are shown on the 'DIFFERENCE' picture; each repeated period has been slightly altered. Non-periodic changes arise from pixels whose effect on the cost function is not detectable.

unit in the algorithmic version, which defies analytical attempts at reproduction.

4.6 Summary

Although it is possible to build a dynamic data switch with incoherent light channels using a FLCoS SLM, we must contend with the vulnerability of such a system to low modulator contrast and pixel drop-out. The Karnaugh Map, where value is encoded by spatial position, is one potential scheme for limiting the error due to misalignment and crosstalk between channels. This may even be desirable in certain applications.

However, the strength of the FLCoS SLM as a phase modulator can be truly exploited in coherent light. For data-switching in particular, the SLM is used display a Fourier phase hologram to replicate its input to an array of positions in the Fourier plane of a lens. Hence generation of suitable binary phase holograms becomes necessary. The simulated annealing algorithm for the generation of binary phase holograms is a well-known and widely-used method, albeit inelegant and compute-intensive.

That there is no analytical alternative to an iterative process such as simulated annealing is due to the lack of a basis set for Fourier holograms. For instance, the grating superposition method tried here produces inferior results, and then only for very simple output targets. That it generates a recognisable version of the target at all is perhaps surprising, and of course tantalising. The imposition of the requirement that the hologram be binary is a particular problem for this method, whereas it is this condition which makes simulated annealing attractive.

Fine-tuning the iterative hologram generation program, to produce the optimum phase mask, largely involves the form and parameters of the cost function. The program tries to generate a mask which reproduces a target output with a felicity determined by the height parameter, h. This is a ratio of output intensities; $h = 1$ implies an attempt to reproduce the target exactly. Values greater than

$h = 0.8$, while not theoretically achievable, more effectively goad the algorithm into producing an optimum result than achievable h values. However, above $h = 0.9$ the divergence of a subset of the output points from the target starts to increase, as intensity is channelled preferentially into the remaining desired output points. This is recorded as a loss of uniformity among equally-weighted output spots. The optimum height coefficient for an efficiency/uniformity tradeoff is found to be $h = 0.9$ for the cost function given in equation 4.6. With this setting, holograms are generated at the theoretical maximum efficiency of 80% with a total range of output intensities of <10% of the mean intensity.

Chapter 5

Compensating for BPH Defects

> *If we knew what we were doing, it wouldn't be research.*
> Steve Jackson Games' *Illuminati*

5.1 Degradation of Holographic Output

In effect, figure 3.5 amounts to a description of the intensity response of an FLC phase modulator subject to defects of FLC layer thickness and molecular alignment. We can use this result and those above in figure 3.3 to introduce phase and amplitude errors into the iterative BPH generation process. First we ask: how does the BPH output degrade due to such effects? Then: how will the generated BPH differ if we can factor in FLC defects at design-time?

The SLM defects have three implications for a holographic application:

1. **Light directed into DC spot,** due to inactive areas of device.
2. **Lowering of contrast,** as less light is successfully diffracted into the intended output.
3. **Broadening of spots,** because the FT of the uneven SLM phase profile is convolved with the output spots.

5.1.1 Phase Holographic Output from SLM System

The performance of the 512^2 SLMs with a standard 4×4 fanout is shown in figure 5.2. All three above degradations are apparent. There is also an unwanted diffractive effect which results in a replication of the pattern in one direction. This is an artefact of the 512^2 SLM's D-RAM addressing cycle, as sampled by the acousto-optic modulator. Thus the AO-switch in this 'straight-through' configuration is not a sufficient method for pulsing illumination of this device for phase modulation; this underlines the difficulty in adapting this device for coherent processing applications. Figure 5.3 shows similar results for an off-axis 4×4 fanout pattern.

In comparison, figure 5.1 shows a 8×8 fanout as generated by a 256^2 SLM. The result is more distinct than that from the the 512^2, without unwanted diffractive orders apart from the large DC spot. However, the three degradation modes are apparent here too, even with non-working parts of the SLM array masked out.

5.1.2 Resolution and Crosstalk

From the above, it is clear that SLM uniformity defects contribute to a broadening and weakening of phase holography output. Although this is optically a graceful degradation, for an interconnection application it will eventually cause the system to develop crosstalk and potentially fail in a far-from-graceful manner. To combat this, it is desirable to increase the resolution of the hologram plane in order to spatially separate the output spots, and to use a small repeated hologram. The 512^2 FLCoS SLM is a high space-bandwidth device, more than capable in theory of generating resolvable arrays of output spots of the sizes considered here. More complex holographic output (such as a larger interconnection) requires a larger repeated hologram unit to maintain efficiency and uniformity. Output spots are correspondingly harder to resolve. The reduction of the width of output spots, by an SLM pattern which acts to counter phase error, is potentially a way of lowering the risk of crosstalk and allowing more performance to be extracted

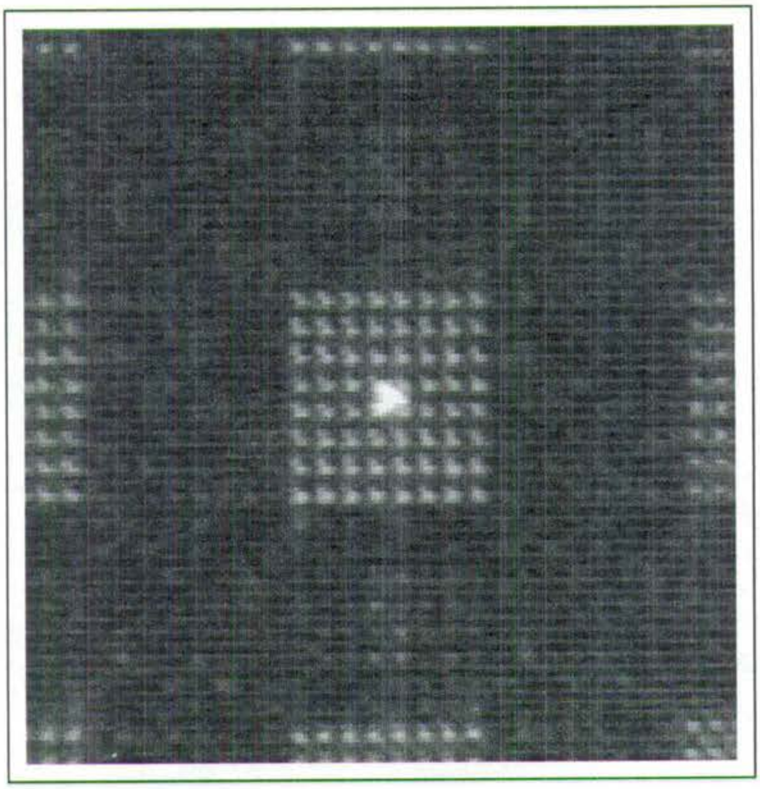

Figure 5.1: 8 × 8 fanout pattern from a 256^2 SLM, using only a portion of the active area. There is a significant DC spot.

from a system built of imperfect components.

5.1.3 Modelling BPH Defects

In order to model and quantify the effects of various SLM errors for the BPH application I present further computational results based on an iteratively-designed 'fan-out' phase mask, which generates a square array of intensity spots as before. As a benchmark I will use 4x4 and 8x8 fan-out phase masks, as shown in figure 5.4. These are 512^2 binary arrays with a small repeating unit, designed so that when displayed on a 512^2 resolution modulator they will produce usable arrays of spots: spatially separated and tolerant to SLM defects. The use of many periods of a small repeated pattern is an obvious way in which phase holography circumvents defects in the modulator plane, through replication and delocalisation

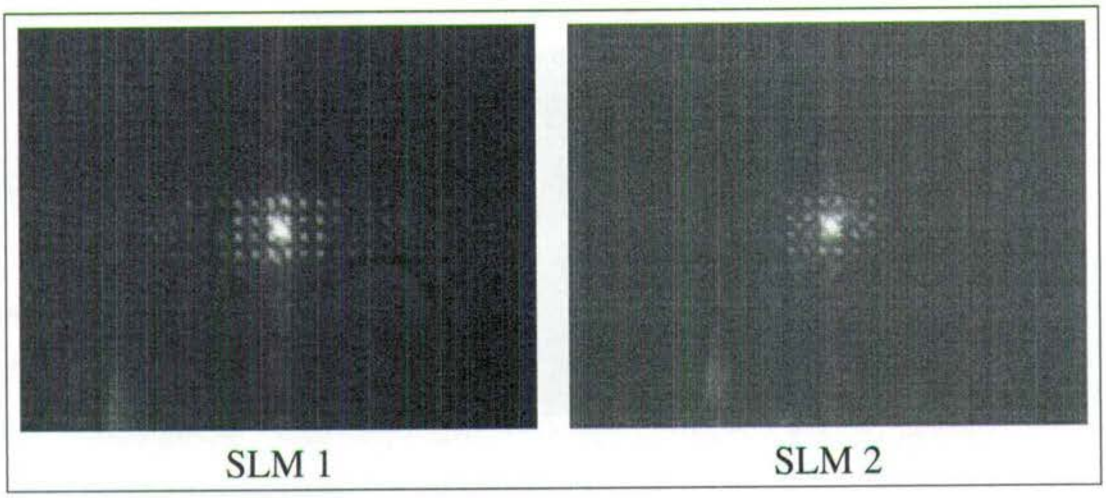

Figure 5.2: 4 × 4 fanout patterns from the 512^2 SLMs in figure 3.9.

of the information in the hologram. Naturally, it suffices to calculate the small repeated period of the mask and replicate it across the large array, *except* that when phase and amplitude errors are introduced which vary slowly across the array it becomes necessary to calculate the mask as a whole. Therefore certain shortcuts which rely on the symmetry of the generated mask are not available.

5.1.4 Random Pixel Noise

Figure 5.5 shows an example of the degradation of BPH performance under the effect of *random pixel noise*. This was generated by successively setting pixels of the BPH at a permanent random value: the graph shows how the efficiency and uniformity of the 8 × 8 fanout BPH degrade as an increasing percentage of its pixels are randomised. Although the output spots remain of uniform amplitude as more noise is introduced, efficiency drops steadily as the bright spots sink into the noise generated by pixel phase errors. Further application of the annealing algorithm is not able to significantly improve efficiency, as the noise pixels are no longer available to the algorithm for adjustment. The graph shows that such efficiency improvement as is possible is achieved at the cost of uniformity, as the algorithm starts to direct power into a subset of the output pixels. Given that

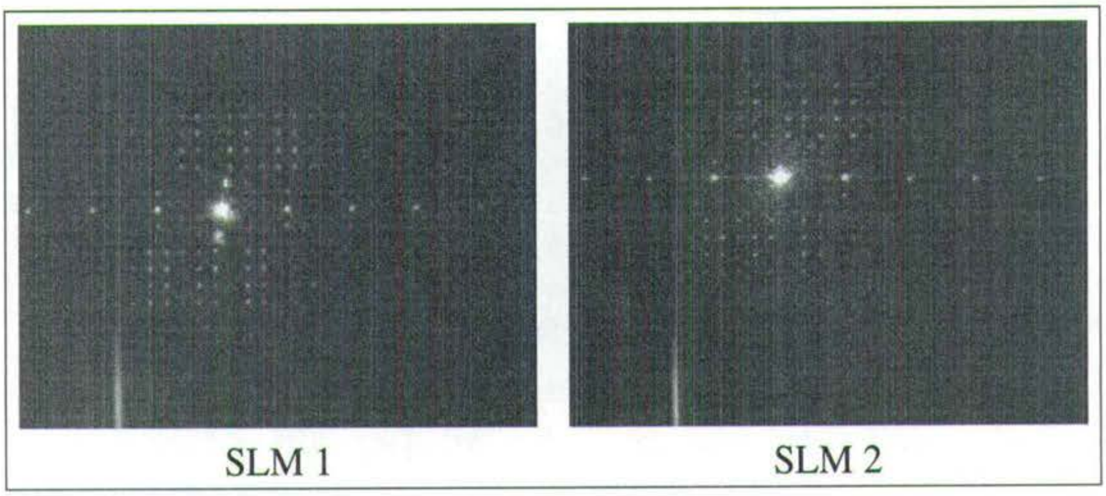
SLM 1 SLM 2

Figure 5.3: Off-axis fanout patterns, as figure 5.2.

such compensated annealing is of little efficacy, how problematic is pixel dropout? A 10% expected *electrical* pixel loss is, in practice, unreasonably pessimistic. The conclusion is that the phase holography application is not severely degraded by realistic levels of electrical failure.

5.1.5 Bowed Surfaces

The more prevalent and serious effects which degrade SLM performance, the variations in FLC thickness and alignment over significant areas of an SLM, are harder to treat as a quantified amount of error. For instance, the variation in thickness of a cell is due a warp in the silicon which has an effectively arbitrary shape. In the field of adaptive optics the solution has been to consider Zernike polynomials as a test of wavefront-correction systems[54].

One simple warp which we can consider is to assume that the silicon is a shallow paraboloid, centred at the centre of the SLM. This corresponds to a typically-seen effect during SLM fabrication: a bow in the chip is induced during processing and cutting which causes a thickness variation in the cell when spaced at the edges against a flat glass cover. This is manifest in a feature of the filling behaviour of

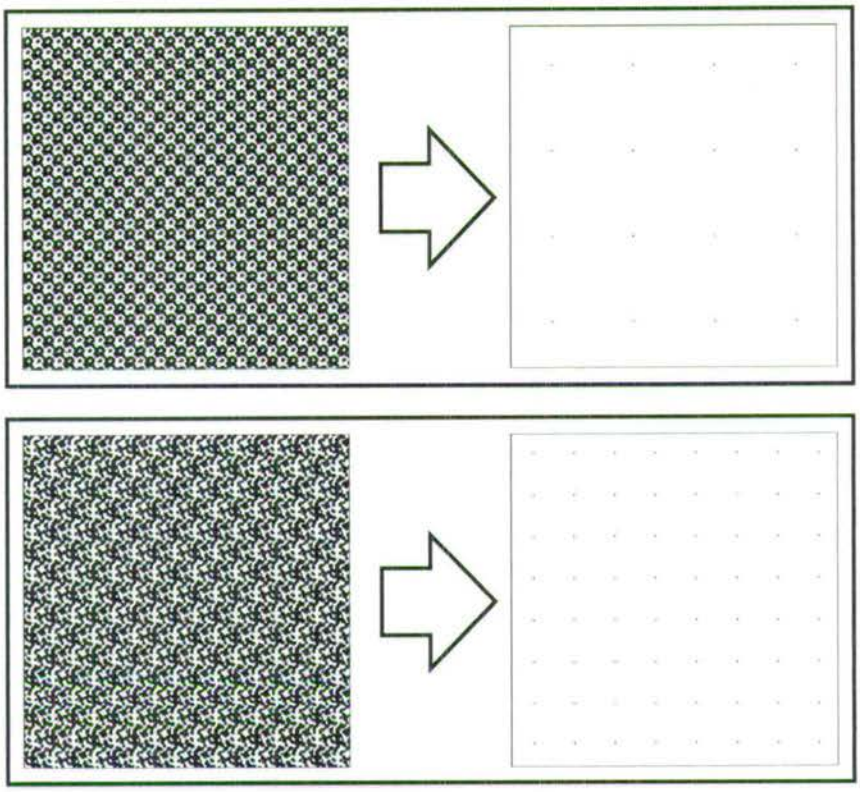

Figure 5.4: BPHs for the 4x4 and 8x8 fanout arrays. Note the small repeated unit in each phase mask: this gives a widely-spaced output. BPHs are calculated on a 512^2 array, and their Fourier transforms are shown in negative.

FLC, whereby there is a discrepancy in filling speed between FLC at the edges and centre of the cell; this is caused by a variation in the strength of capillary action with a varying cell gap.

As stated earlier, the FT is simply a convolution of the output spots with the FT of the combination height/cell thickness phase and amplitude function. This is illustrated in figure 5.6. As the cell errors broaden the output spots, an interconnection application runs the risk of crosstalk. To simulate a (paraboloidally or otherwise) warped backplane, we need to change the phase and amplitude components contributed to the FT by each pixel, depending on its position in the SLM array. One way to do this is by a look-up technique whereby we construct arrays for the on- and off-state pixel modulation values and refer to them during

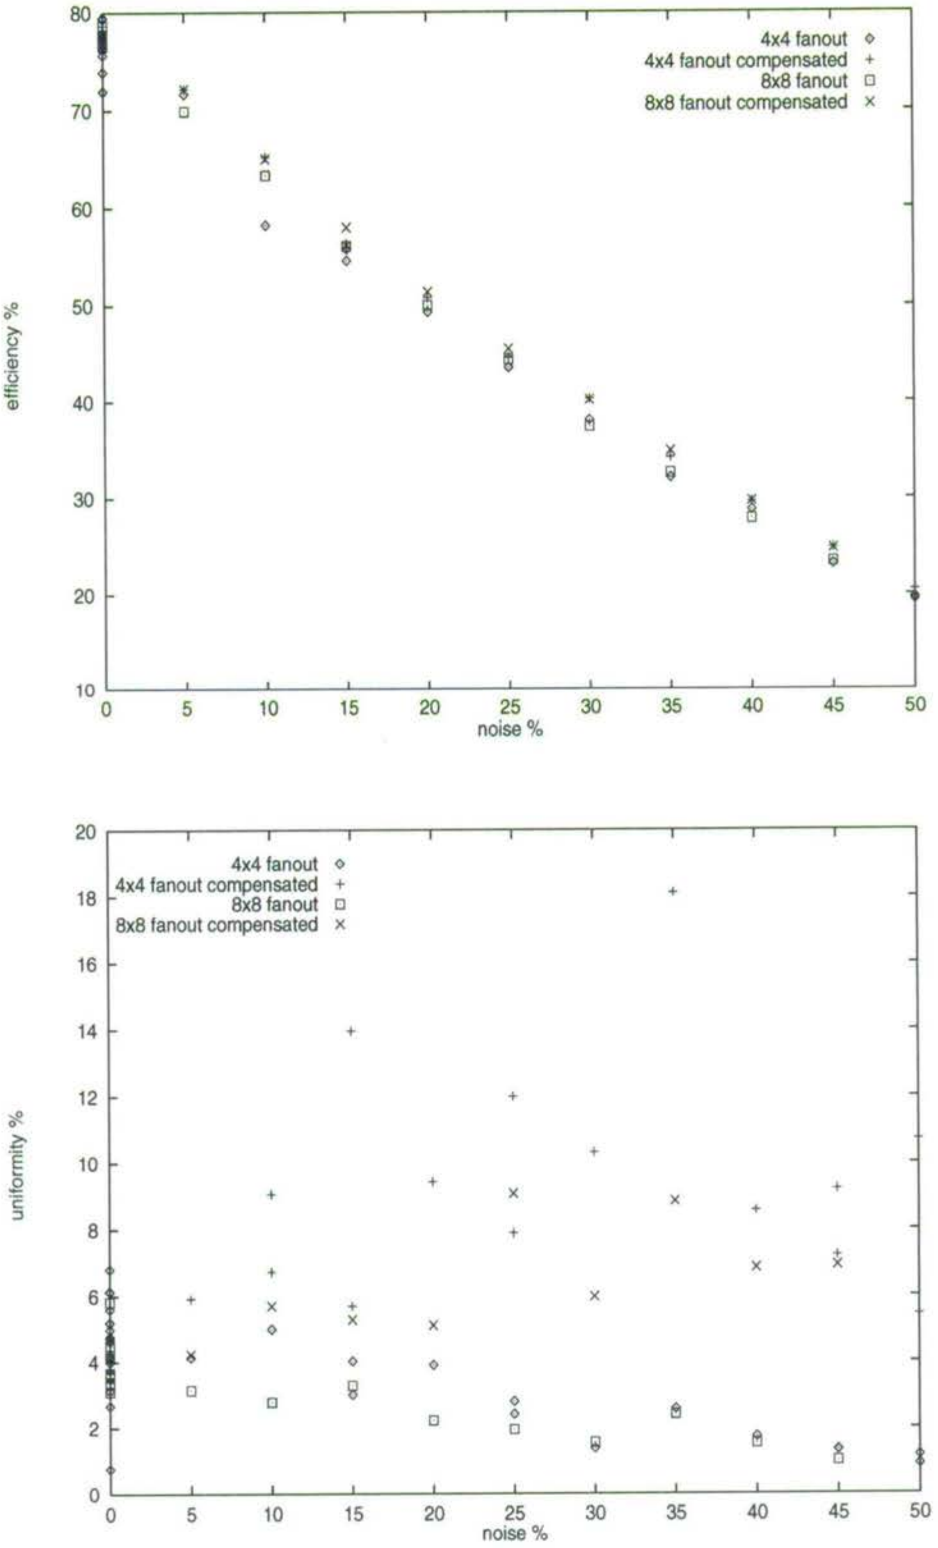

Figure 5.5: The (a) fall in efficiency of a BPH with rising percentage of randomly-stuck pixels, with minimal gain from further annealing, with (b) the uniformity loss caused by further annealing.

the annealing algorithm; see Appendix A for the implementation of this.

Due to the cell thickness variation caused by backplane warp (i.e. assuming a flat upper surface) the pixel amplitude and phase values are a combination of those shown in figure 3.3 and a further phase offset due to the height variation of the backplane. The x-coordinate in the figure 3.3 graphs is now given by the distance of the pixel from the SLM array centre. Hence for a pixel at offset (i, j), let the backplane height be $H(i, j)$ and find an expression for the output state phasor (r, ϕ), as follows.

ξ is the phase difference caused by a height of $H(i, j)$, taking into account the mean refractive index of the FLC, $\bar{n}$:

$$\xi = 2\pi \frac{2\bar{n}H(i,j)}{\lambda} \tag{5.1}$$

η is the retardation caused by $H(i, j)$, dependent on the birefringence Δn:

$$\eta = \Delta n \times \xi \tag{5.2}$$

r can be calculated as the magnitude of the vector sum of the usual — wholly real — output phasor and the complex offset described by η:

$$r = \frac{\sqrt{2}}{4}\sqrt{[1 + \cos(\eta)]^2 + [\sin(\eta)]^2} \tag{5.3}$$

Whereas, ϕ_{on} depends on the phase offset ξ added to the phase of r, and ϕ_{off} is always π behind it:

$$\begin{aligned} \phi_{on} &= \xi + \arctan[-\frac{sin(\eta)}{1 + cos(\eta)}] && (5.4) \\ \phi_{off} &= \phi_{on} + \pi && (5.5) \end{aligned}$$

When an existing BPH is used to reference these on- and off- look-up planes, and then transformed, the degradation in the output is as shown in figure 5.6, due to the warp function applied to the input plane. The BPH can now be re-generated using phase values from the look-up planes to see if an improvement can be gained once this effect is taken into account. As above, I refer to this as *compensated annealing*.

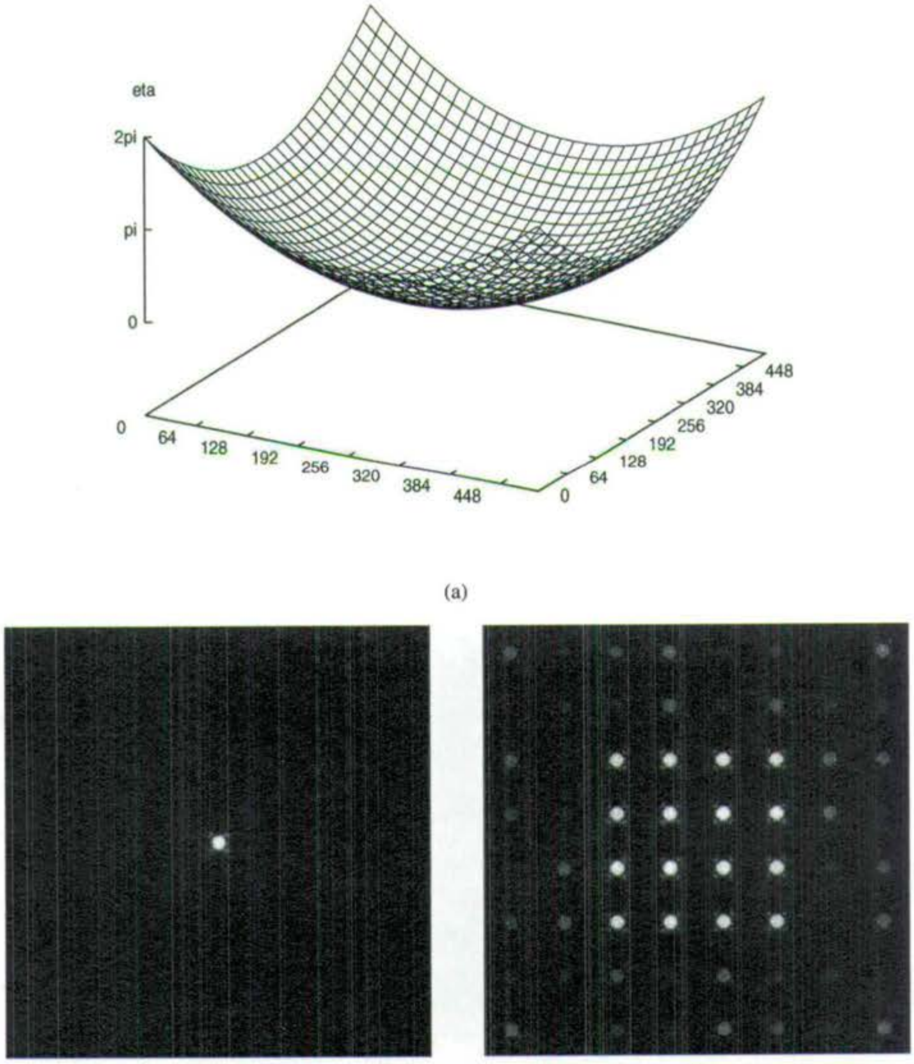

(a)

(b)

(c)

Figure 5.6: The output of a BPH on a warped cell is simply the convolution of the FT of the warp and the clean BPH output; (a) is the warped backplane, (b) its Fourier transform, and (c) the predicted output of a 4 × 4 BPH with this phase aberration.

5.2 Compensated Annealing

The goal is to use a dynamic phase modulator both as a hologram and a phase correction element by iteratively generating a hologram given prior knowledge of the modulator's phase distortion. This technique is separated into attempts to correct firstly a warped mirror and then a wave plate of varying thickness.

The method throughout is to first calculate what phase and/or amplitude modulation is possible for each pixel in the artificially warped array. Then the BPH generation program can be run using these new pixel states, aiming as usual to match the Fourier transform of the array with a target intensity pattern. The key additions to the program are a method of creating and importing warped surfaces, and the calculation of individual pixel modulation states.

In effect, we are seeing what the simulated annealing algorithm can do when given varying and imperfect modulators to work with.

Hence we have three complementary results for each case: the quality of the original BPH output using the unwarped array, the (reduced) quality when using the warped array, and the quality of the new compensated BPH output using the warped array.

5.2.1 Backplane Warp

Backplane warp, with the SLM acting as a curved mirror, can be modelled by adding a phase term to each pixel representing the difference in (system) optical path at this point. The effect of such distortion on the BPH system is to convolve the FT of the SLM's surface profile with each output spot, broadening them, lowering efficiency and promoting crosstalk.

This is similar to the problem of wavefront correction in adaptive optics, with two key differences:

- Only binary phase modulation is available.
- The Fourier transform of the wavefront is taken as the output.

Binary phase modulators have been tried for wavefront correction[55], with a reported Strehl ratio (ratio of corrected to unaberrated signal) of 40%. FLCoS SLMs are fast and convenient but inflexible for this task, which is normally the preserve of analogue phase modulators. As well as astronomy, there are potential applications in laser communications and imaging in a turbulent medium. Compensated phase holograms are primarily applicable to communication applications where some optical processing is required by the modulator.

Since phase corrective plates are commonly used to remove wavefront phase errors from systems, we might expect the algorithm to at least partially compensate for this purely additive phase function. The key issue is the extent to which a binary element can correct for continuous phase error. In figure 5.7, and again in figures 5.11 and 5.8, I show how the simulated annealing algorithm adapts the phase mask to compensate for phase error.

In each case, a clearly 'corrected' BPH is seen to be created. Generating a BPH to compensate such a distortion produces the interesting result that the BPH 'warps' to compensate the added phase function, rather than simply, say, reversing contrast. The new BPH has continuous areas of $\pm\frac{\pi}{2}$ phase, in a non-repeating pattern which complements the added phase. This pattern is a BPH with built-in corrective phase plate, and as such loses the symmetry of the uncorrected BPH. The final modulated *wavefront* has continuous phase levels, even though each pixel is a binary modulator.

Experimentally the efficiency for the new compensated mask is found to be 50–60% for a maximum warp of around λ, even for asymmetric Zernike distortions — for the on-axis spot arrays. This is down from the near 80% of the unwarped version, and up from the $\approx 10\%$ of power which falls into the output spots from the raw warped mask. An analogous 'Strehl ratio' could therefore be given as 60–75%. For off-axis targets such as figure 5.8 all efficiencies are reduced, but the

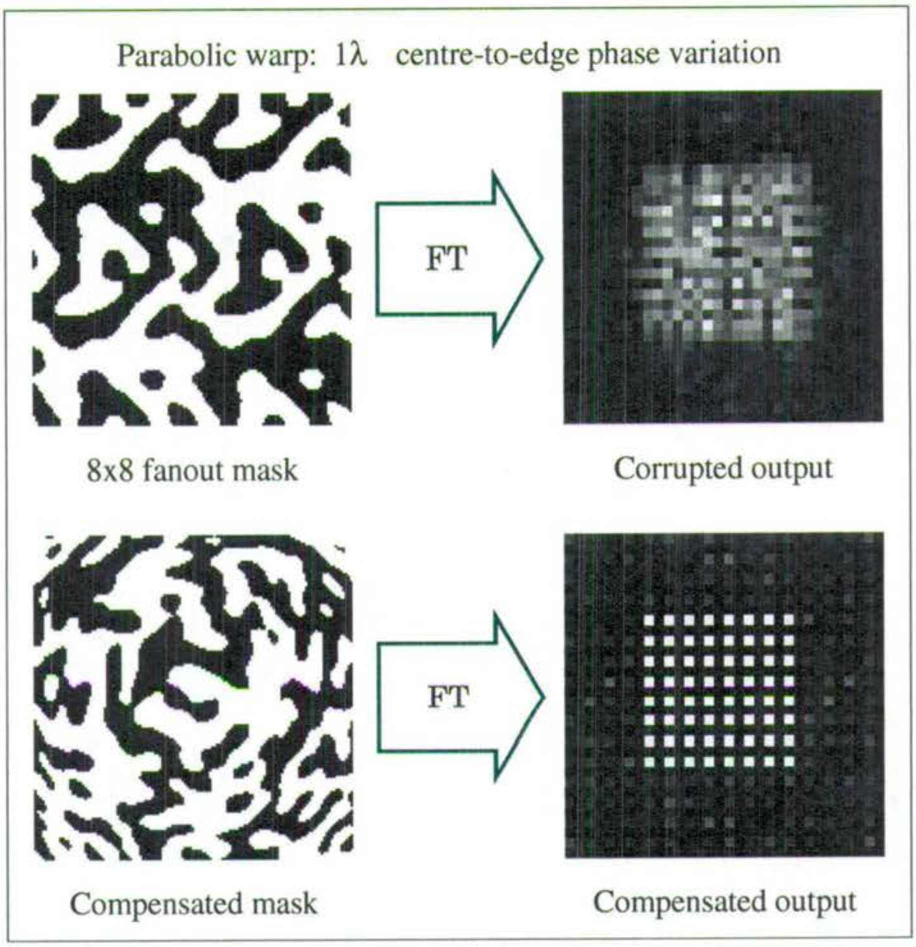

Figure 5.7: A 8×8 fanout mask with a parabolic phase distortion of $\sim 1\lambda$ produces a badly degraded output. Upon compensation the output spots are recovered.

ratio is similar.

We can conclude that it is possible to correct for wavefront distortion with a binary phase modulator in the Fourier phase holography application.

5.2.2 Varying Cell Thickness

The problem of generating a binary phase hologram is usually stated as constraining the inverse FT of some given 2D function to having only the values $+\imath$ and $-\imath$. In a real SLM, we have a binary device which has a set modulation

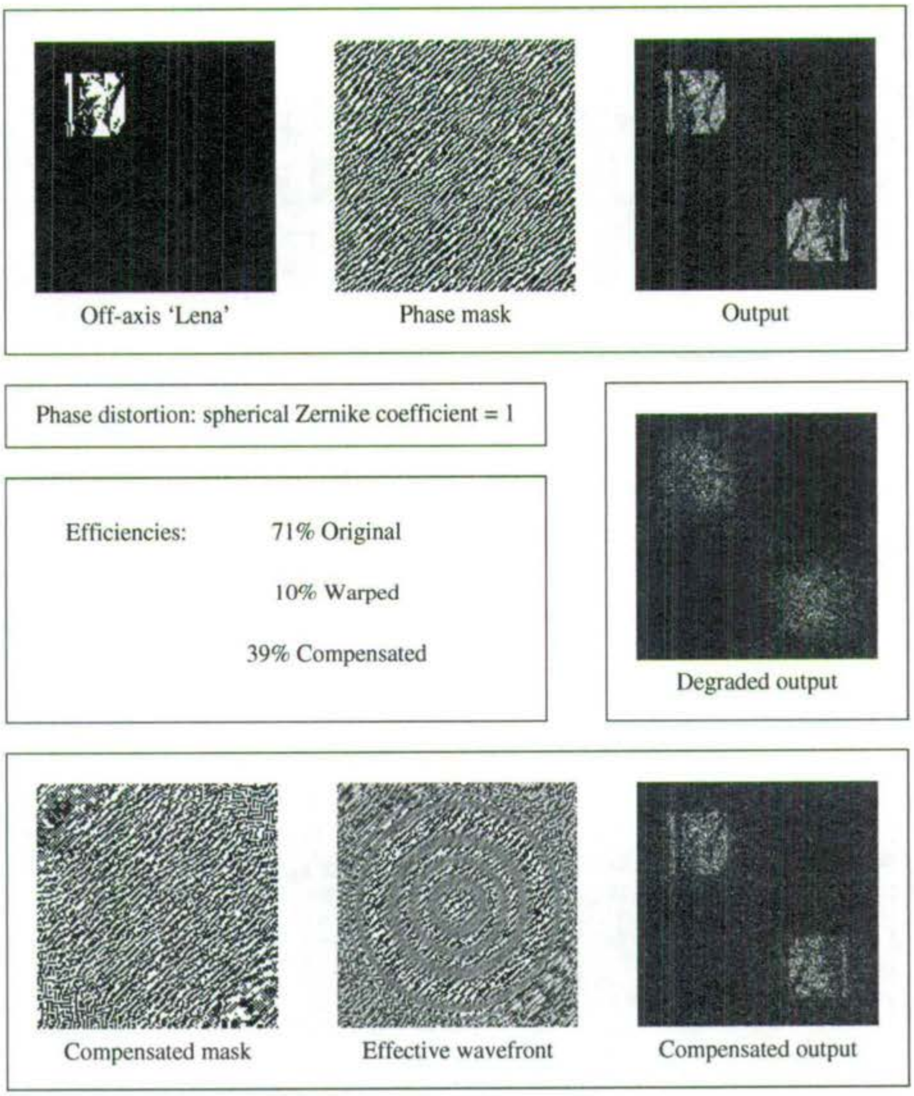

Figure 5.8: An off-axis target with a spherical phase distortion, and its compensated phase mask. On the top row are the target, its generated BPH, and the output of the BPH. Below this output picture is the effect of imposing a spherical distortion (spherical Zernike coefficient = 1) on the BPH — the output is lost. On the bottom row are the BPH generated to compensate for this distortion, the wavefront created by combining distortion and compensated BPH, and the resulting output, which has recovered somewhat. The efficiencies given are for the three output cases — clean, degraded and compensated.

map whose modulation varies over the full range of phase and amplitude. The problem changes to become that of finding the inverse FT of a function which takes the values offered by the SLM.

Variation in cell thickness means that some areas of the SLM will modulate phase with reduced output amplitude and an added phase offset. Additionally, there is a contrast reversal at the full-wave plate thickness.

What maximum efficiency can we expect from an arbitrary distribution of possible modulation states? As an aside, note that the phase hologram is no longer strictly a binary pattern. It is now an array of various amplitude/phase modulators each switchable into two fixed states. As such it is possible to generate asymmetric output patterns analogous to the 'pseudo 4-phase' holograms reported by Samus[46], having lifted the binary phase constraint on output noted by Wyrowski[47]. This is an ironic development: an FLC cell which is imperfect enough to generate a range of amplitude/phase states is able in principle to outperform a purely $\pm\pi$ modulator — the catch, once again, is knowing the modulation map of the device or being able to calculate holograms *in situ*. This also means that we cannot take advantage of the symmetrical output of a BPH by iterating over one half of a symmetrical target, as is usually done to save algorithm time.

Compensation solely for cell thickness error gives considerably less dramatic results than when the additive phase of a curved mirror is considered. The predicted reduction in output amplitude from some areas corresponds to the 'removal from consideration' seen for non-working pixels, which cannot be worked around. Also, the phase offset is more slowly varying than that due to the curved backplane, by a factor of $\frac{\bar{n}}{n_e - n_o}$.

Figure 5.9 and table 5.1 show the efficiency and uniformity from the 512^2 4×4 BPH on variously warped paraboloidal surfaces, before and after compensated annealing. Warp is now in terms of the birefringence error η, as defined earlier, which means these are more severe errors than previously. Figure 5.10 and table 5.2 show likewise results for the 512^2 8×8 BPH. It is instructive to see in which

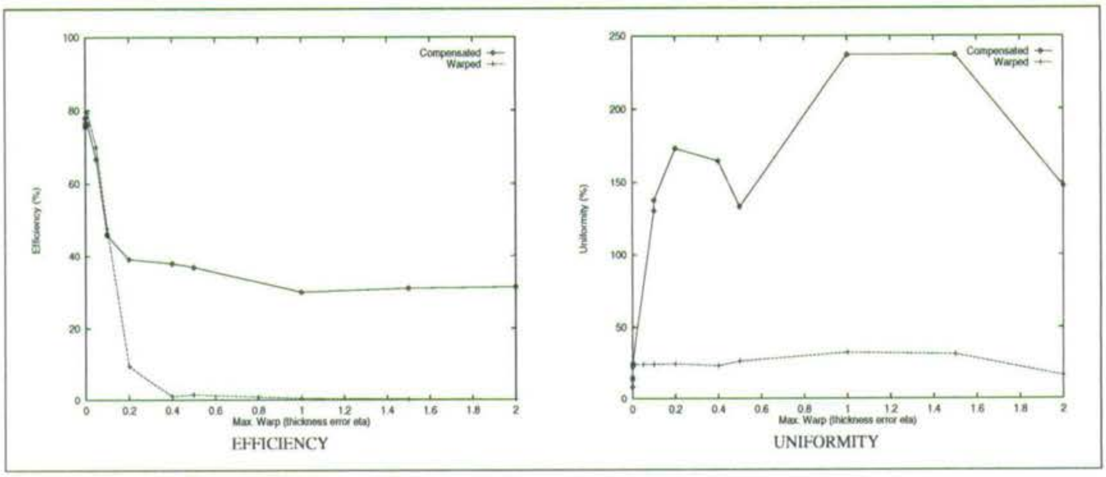

EFFICIENCY UNIFORMITY

Figure 5.9: Graphs of efficiency and uniformity for the 512^2, 4×4 fanout under various amounts of phase warp, before and after compensation.

ways the compensated annealing is or is not able to improve BPH performance, now that birefringence effects are included.

As the cell thickness changes, it moves from being a half-wave plate to a full-wave plate, and loses its ability to alter polarisation angle. The effect is to give less and less modulation effect to pixels as the cell thickness is altered from the half-wave plate ideal. At the full-wave plate thickness, $\frac{d(n_e-n_o)}{\lambda} = 2\pi$, pixels do not affect the output, and appear dark through crossed polarisers. As the thickness varies beyond this point, modulation is restored with contrast reversed, so that *on* pixels appear darker than *off* pixels. This effect is of course maximal for a $3\pi/2$ thickness, and thereafter the cycle repeats. Hence the BPH annealing algorithm is less able to ascertain a change in the cost-function as the cell thickness varies from half-wave to full-wave, since flipping a pixel has less effect. At the full-wave thickness the pixel state is arbitrary, and thereafter contrast is reversed in the generated BPH, as pixels contribute phase values of opposite sign when in a given state.

The primary effect of compensated annealing is to narrow the broadened output spots: the cost function 4.6 does not penalise directly for light falling outside the

Max. Warp (η_{max})	Uncompensated		Compensated	
	Efficiency	Uniformity	Efficiency	Uniformity
0.01	79.180	24.793	76.238	
0.05	70.059	24.845	66.699	
0.1	47.599	24.919	45.731	130.361
0.2	9.778	25.106	39.070	173.362
0.4	1.077	23.965	37.949	164.792
0.5	1.558	26.865	36.922	133.170
1	0.437	32.799	30.024	237.093
1.5	0.190	31.593	31.046	237.086
2	0.056	17.306	31.377	147.459
5	0.015	103.853	31.120	183.149

Table 5.1: Improvement upon compensation of the 512^2, 4×4 fanout under various amounts of spherical phase warp.

intended output pixels, but its action in maximising the light power in the subset of 'on' pixels suffices for narrowing the output power back into the intended spots.

Note the divergence of output pixel intensities (loss of uniformity) as mentioned previously in relation to the algorithm's height coefficient: since the cost function is now using a higher-than-possible value for the desired intensity of output spots, the cost function tends to favour some spots at the expense of others, and uniformity suffers. At first it seems that the uncompensated BPH preserves uniformity whereas the compensated version does not. However, the efficiency of the uncompensated BPH drops to such low values as to make its uniformity of no interest: there is, actually, no intensity in these 'uniform' spots. Recall that the uniformity is measured as a percentage of the mean intensity. Thus the range of intensities in the compensated output spots is roughly twice the mean in the worst cases. This is an unhappy tradeoff for the recovery in efficiency.

Moreover, since the phase distortion introduced by birefringence error is small compared to that caused by reflection from a curved surface, and the loss of pixel output amplitude cannot — as we have seen from pixel dropout simulation — be compensated for, improvement from compensated annealing is almost entirely

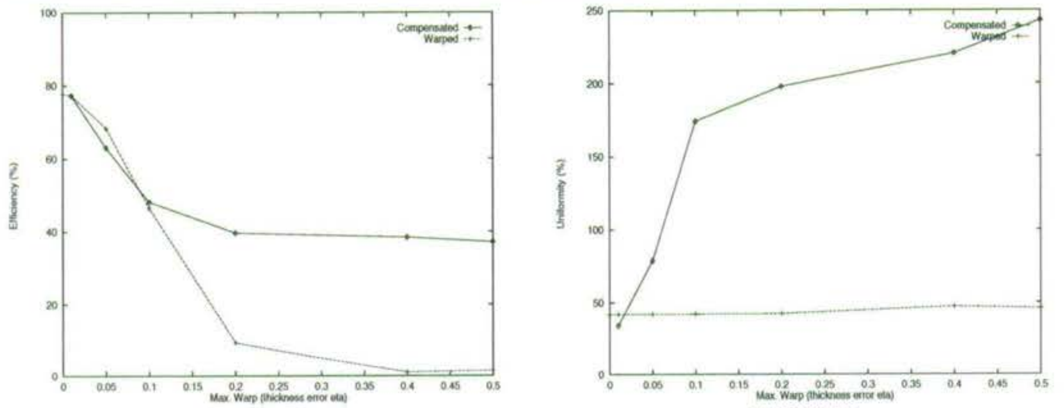

EFFICIENCY

UNIFORMITY

Figure 5.10: Graphs of efficiency and uniformity for the 512^2, 8×8 fanout under various amounts of phase warp, before and after compensation.

due to the wavefront correction effect described above. The birefringence error serves to lower total power throughput, reverse contrast and add minor phase distortion. At high birefringence variation, compensated annealing results in a reversal of pattern contrast to match the contrast reversal effect, and minor changes in pattern shape.

From the above graphs, it is clear that efficiency levels off at $\approx 30\%$ for high uniformity variations. This corresponds to a 'Strehl ratio' approaching 40%. In this result, however, the dominant compensation effect is the correction of a curved wavefront; the simulated annealing algorithm does not have such an effective way to work around a loss of light amplitude caused by changes in birefringence depth.

For SLM systems, absolute wavefront distortion may not be relevant to the quality of BPH output until typical backplane height variation can be reliably reduced to $< 2\lambda$, and systems are designed which depend critically upon the placing of components and the precision of output spots. Since the net effect of curvature and wedging in the mirror array is to move the effective Fourier plane, components can normally be adjusted to suit. The high space-bandwidth product of the

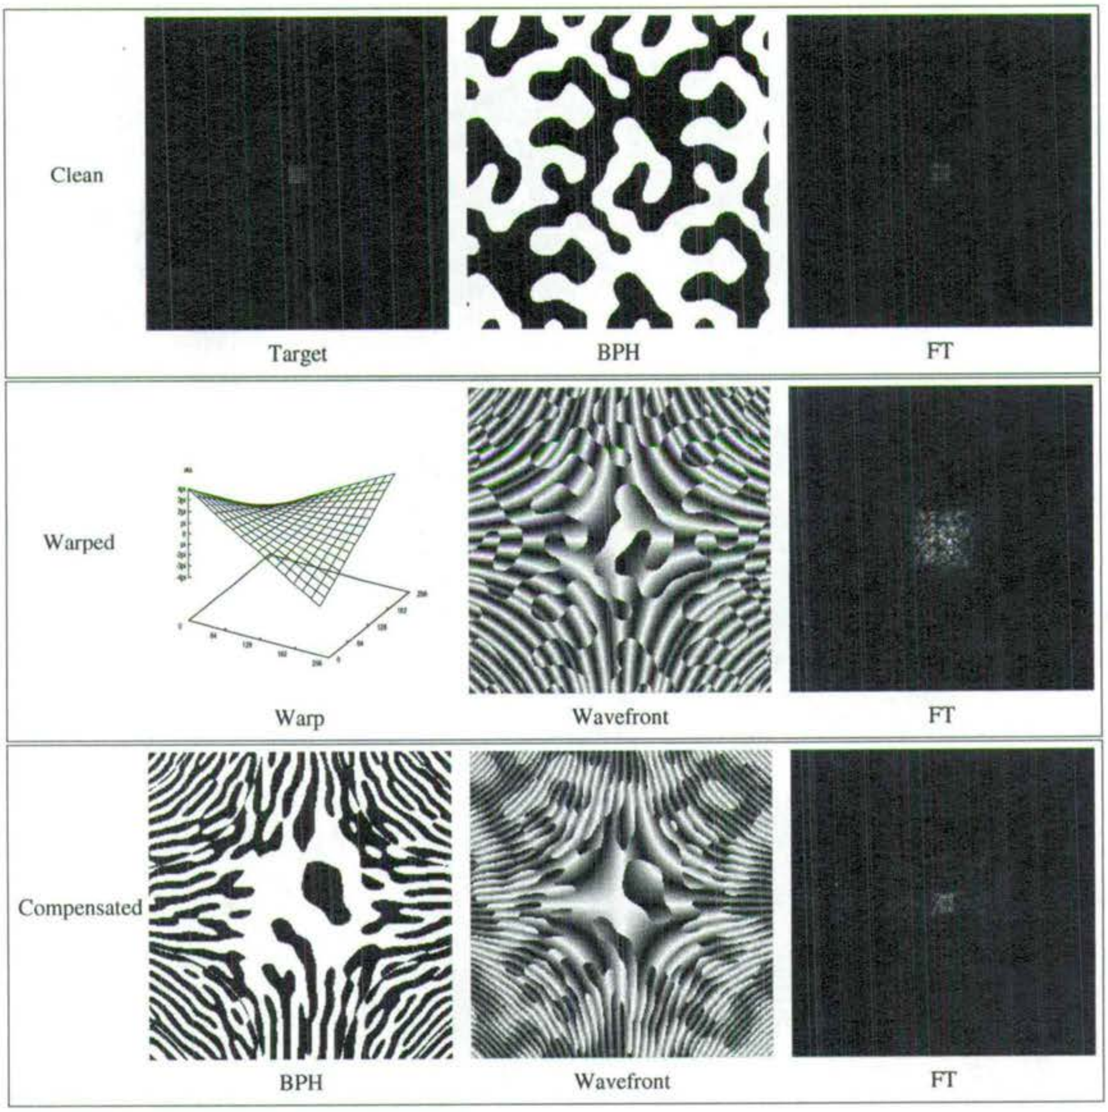

Figure 5.11: 8 × 8 in 256^2 fanout, astigmatic warp, and compensation.

Max. Warp	Uncompensated		Compensated	
(η_{max})	Efficiency	Uniformity	Efficiency	Uniformity
0.01	77.301	42.097	77.232	34.316
0.05	68.399	42.167	63.006	78.890
0.1	46.475	42.297	48.145	174.212
0.2	9.550	42.558	39.655	197.954
0.4	1.057	47.241	38.408	220.500
0.5	1.527	45.972	37.085	243.005
1	0.427	55.680		
1.5	0.187	80.428		
2	0.056	154.570		
5	0.091	686.313		

Table 5.2: Improvement upon compensation of the 512^2, 8 × 8 fanout under various amounts of spherical phase warp.

	No warp	Warped	Compensated
Efficiency	77.125	4.439	36.667
Uniformity	36.522	444.808	217.001

Table 5.3: 8 × 8 fanout in 256^2 array with astigmatic warp as shown in figure 5.11.

devices can likewise make up for broadened output spots, by spatially separating them. I suggest that when devices are of high enough quality to escape the worst effects of birefringence depth variation, wavefront compensation is available to improve a Fourier-plane output which is subject to additive phase error. This, of course, need not be due to the SLM itself. I have demonstrated that in principle a binary phase modulator used as a holographic element can compensate for *system* wavefront distortion, and this may have application in active and adaptive optics.

5.2.3 Importing Captured Cell Profiles

To import the measured cell thickness profiles from chapter 3 into the BPH generation software I build a 'warp' image from the captured Zernike coefficients, which is then recognised by the code. Translated into the BPH software, this corresponds directly to an effective thickness variation in units of $\bar{n}\lambda$, where $\bar{n}$, as before, is the effective refractive index in the cell. This can be converted to η, the birefringence depth, by recollecting that, for CS1031, $\delta n = 0.17\bar{n}$.

BPHs are then generated to compensate for the captured thickness profiles of the SLMs. In figure 5.12 I show a 4 × 4 BPH as calculated for each. As expected, these compensated masks are calculated to give ~30% optical efficiency, as shown in table 5.4. The improvement in calculated output can be seen from the Fourier transforms of each in figure 5.13. The convolution of the FT of the distorted reflector and the required output is clearly seen, together with the quenching of this effect upon compensated annealing.

Note again that as the efficiency (table 5.4) of the warped masks is so low, the superficially good values for uniformity are barely meaningful. Moreover, this should be borne in mind for the rescaled images in figure 5.12, which greatly increases the contrast of the warped FT relative to the compensated version. The warped FTs simply would not be visible if the same greyscale was used throughout. The *figure* illustrates the concentration of output power back into the spots upon compensation, and the *table* shows the recovery in efficiency. The apparent 'cost' in uniformity upon compensation is academic, since the efficiency results show that the 'uniform' uncompensated FTs are, in fact, uniformly dark.

		No warp	Warped	Compensated
SLM 1	Efficiency	79.448	0.001	30.998
	Uniformity	22.416	247.809	214.222
SLM 2	Efficiency	79.448	1.954	37.265
	Uniformity	22.416	27.576	184.890

Table 5.4: Calculated performance of the compensated masks shown in figure 5.12.

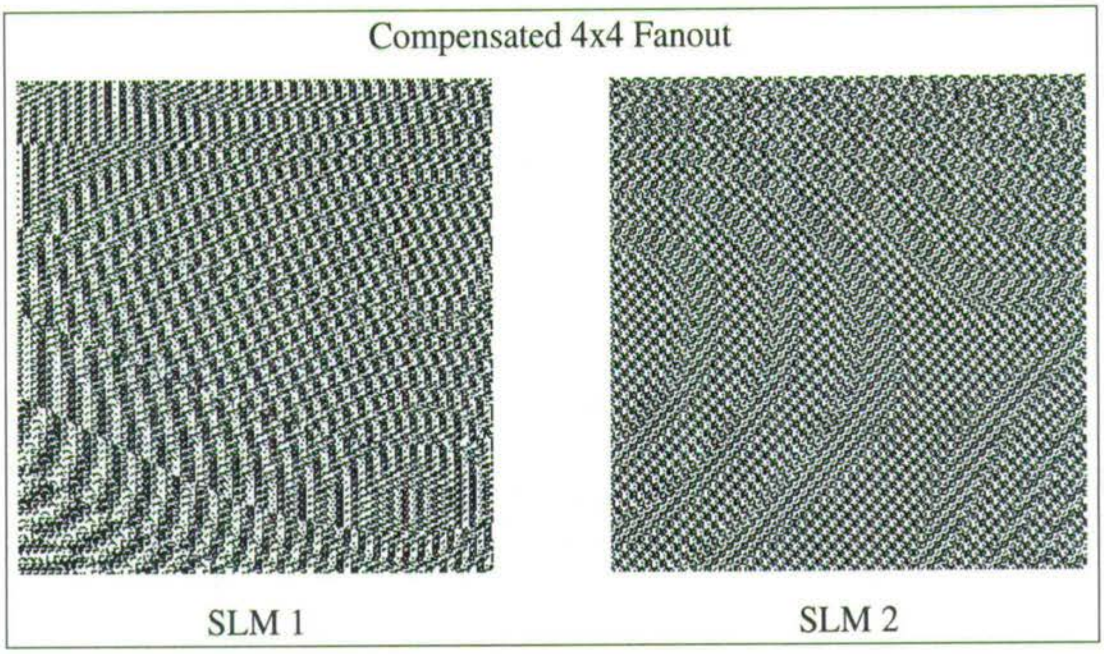

Figure 5.12: A 4 × 4 fanout BPH, as generated for SLM 1 and SLM 2 using the captured thickness profiles shown above.

5.3 Summary

5.3.1 Binary Phase Holography As Adaptive Optics

As established in chapter 3, a measurably imperfect SLM is a calculably imperfect phase modulator. Chapter 4 introduced ways of generating phase masks for given target outputs using only the knowledge of how any one mask pixel affected the output. This chapter has synthesized both, by calculating how defective pixels will modulate phase, and then using these various modulation states as the raw material for the BPH generation algorithm.

The result has been to create phase masks which are matched to specific defect profiles, and are able to compensate for the defects in order to function as phase holograms. By analogy with adaptive optics, which also seeks to compensate out phase error, this technique yields a 30% Strehl ratio. The complication over a conventional adaptive scheme is the Fourier tranform step between modulator

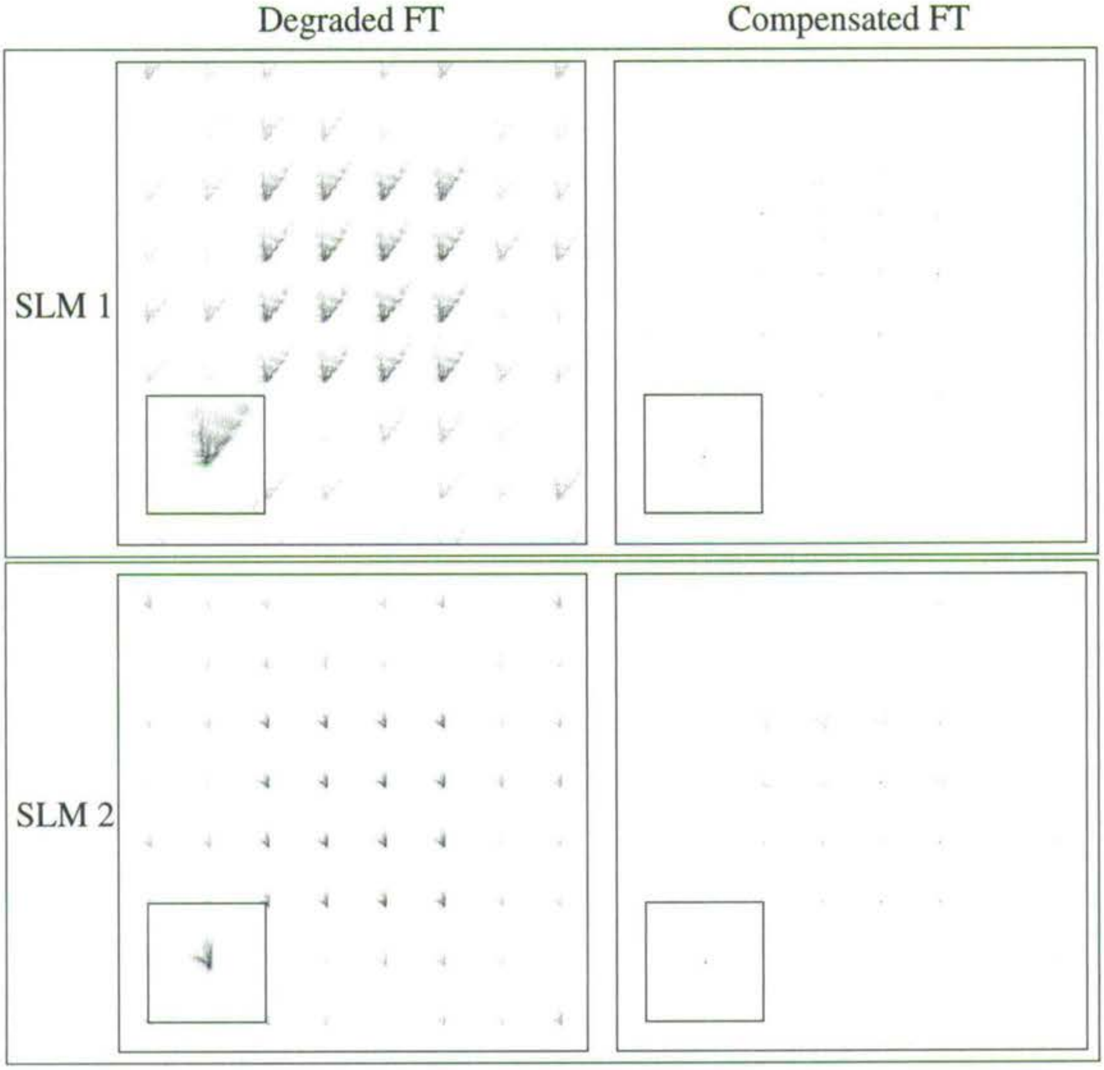

Figure 5.13: Calculated FTs (shown in negative) for the captured cell profiles for SLM 1 and SLM 2, using both the unmodified 4 × 4 BPH and the compensated versions in figure 5.12. For each FT, a closeup of a single spot is shown.

and output, which gives the situation a unique twist.

5.3.2 Further Work

The next logical step is to test how well this works in real systems. Unfortunately pressure of time made this impossible, particularly due to the time entailed in further SLM fabrication. As noted earlier, there is a minimum useable quality of device even for defect compensation tests.

Another area to be explored is compensation for misalignment, whereby the displayed BPH can be adapted to steer the output to where it should be. This would have particular application in free-space data links, where alignment is largely an unsolved problem.

5.3.3 Projected Integrated System

One implementation of this work would be a real-time feedback system for designing phase masks, similar to that reported by Paige[56]. Unfortunately we are once more up against the lack of a basis-function set for binary phase FT patterns. This means that it is difficult to make systematic changes to the phase pattern on an SLM which will produce measureable differences in the Fourier transform.

The computing power required to anneal phase masks is considerable, however as a one-off corrective step this is acceptable. The power to capture and compensate for phase variations in a holographically-connected optical communication system may yet be a useful adjunct, especially when the components will not easily be physically moved.

A system is envisaged where the composite phase distortion can be captured once upon construction, and then used in the design of the set of phase masks which are to be used by the SLM. Further updates of the phase masks would be possible, to cope with system deterioration. In such a way, a delicately-aligned free-space coherent system might be tolerant to faults in its components, construction and operation.

Conclusions

> — *I am flush with the heady glow of scientific discovery.*
> — *Take five.*
> Berke Breathed, *Bloom County*

The Edinburgh Applied Optics group concerns itself with virtually every aspect of LC SLM technology, from materials to systems; perhaps inevitably, to those involved the most salient features of SLMs have come to be their difficulties and faults. Complementary strategies have been pursued in the group: to understand the defect-producing phenomena; to improve fabrication techniques; and to quantify the effect of SLM faults and seek to compensate for them.

Classifying SLM Defects

A menagerie has been established of particular faults which affect the FLCoS SLM as a complete device, as opposed to its constituents when each is taken separately. There are three classes: electrical faults, crystal defects and loss of cell uniformity. For the last of these it has proved possible to work through a matrix treatment to show how modulated light is affected by variations in the FLC layer thickness. As cell thickness varies from the ideal, output light states degrade in both amplitude and phase.

Effect on Binary Phase Holography

The demanding task of using FLCoS SLMs as coherent phase-modulating elements is adversely affected by cell defects. Through simulation and live results it has been demonstrated how these errors impinge on the SLM when used as a holographic fan-out element. This application is not chosen arbitrarily; phase masks are a standard part of many proposed optical interconnection schemes, and the FLCoS SLM stands as one of the few electrically configurable high-resolution phase mask technologies.

It is also a uniquely fault-tolerant application. Even when an essentially one-to-one mapping is required, holographic interconnection offers the robustness of non-localized information and redundancy which means a lost pixel does not imply a lost information channel. In fact one failed pixel in ten will not catastrophically affect the operation of a $\sim 200^2$ fan-out mask.

The severest problems with FLCoS SLMs, though, are not electronic but optical — backplane height and cell thickness variation. These effects are lethal for a phase-modulation application.

Defect Compensation

However, it is possible to correct wavefront with a binary phase modulator in a Fourier transform system. Efficiency of phase holography drops from 80% to unusably low levels as a smooth phase function is superimposed on the modulator, to simulate the phase error caused by a defective cell. Then, after the iterative annealing algorithm is re-applied to compensate for this warp, the efficiency recovers to $\approx 30\%$ — a 'Strehl ratio' of 40%, and indeed this has a relation, and perhaps an application, in adaptive optics systems.

This re-application of the annealing algorithm using imperfectly-modulating pixels is referred to as 'compensated annealing'. It is shown to work using computed phase error derived from actual SLM cell thickness profiles.

Future Work

Further coherent systems using FLCoS SLMs will need better devices than the regular standard required for miniature displays. FLCoS SLM fabrication will need to be characterised and made reliable. However, there are ways and means for improving not just SLM, but system wavefront errors, and this seems the natural province for a switchable phase modulator.

A data-switching system which uses a FLCoS SLM is able to adapt. Either to respond to inherent device shortcomings, emergent ones during operation, initial and current system misalignment, wavefront aberrations and so on. While fulfilling its role as a dynamic one-to-many switch, the SLM can also be a powerful way of adjusting the system to deal with optical problems during operation.

I have shown how a variation in the phase mask displayed can improve the Fourier-transformed output. The errors corrected for need not originate with the SLM itself. I propose a system which captures its own output wavefront at intervals during operation, and then recalculates its set of phase masks to compensate for the predicted deviation from its correct output. In effect, an automation of chapters 3 and 5.

To return, then, to the theme of chapter 1: it is accepted that electronic computing technology is travelling in a cul-de-sac; whether the next generation of photolithography is the penultimate or not, at some stage there is not going to be enough room, speed, or *dimensionality* in metal-over-silicon. When a new technology *is* at last inevitable, there will be various possibilities of which an optics/electronics hybrid is a relatively mature and proven candidate. Optical interconnection is extant as fibre networks; the shape of generic *intra-machine* information optics, though, has yet to be designed. The essence of light as an information technology is to make data connections and exploit its unique properties. Dynamic phase modulators, and Fourier holography, are a powerful and purely optical solution. The extent to which this solution is workable in the real world of imperfect devices has been demonstrated here.

Appendix A

BPH Generation Code

Here are some sections of the C code written in the course of this work. All of it makes use of a library of image-manipulation functions written by William Hossack[57], and the implementation of the Dames algorithm[52] for BPH generation is built on work by Sergei Samus[46].

A.1 Simulated Annealing Algorithm

A.1.1 Code

This is the C function which implements the Dames simulated annealing algorithm for BPH generation. The calculation now works at double precision, and I have tried to keep as little in the iteration loop as possible. There is a new calculation of the phase contribution which allows for arbitrary look-up images holding on- and off- pixel values.

```
void anneal(T0, delta_t)
float T0, delta_t;
{
  float chance, T = T0, th, re, im, a;
  double err_old, err_delt, err_last;
  int fnord, i, j, k, l, m, n;
  long n_cycles, n_same_err=0;
  fcomplex *trial_point, *trans_point, *on_point, *off_point;
  unsigned char *noise_point;
```

```
  /* first generate *real* error */
/*
  im_transform = ft_full(im_current);
  im_stats(im_transform);
  im_copy(im_transform, im_trial);
*/
  if(anneal_type == ANNEAL_MAX)
    err = calc_max_err(im_transform);
  else
    err = calc_tot_err(im_transform);
  err_old = err;
  err_best = err;
  err_last = err;

  printf("\n first... error: %2.10f%%  max single error: %2.12f\n\n",
    err, calc_abs_err(im_transform));
```

The iterated part of the algorithm begins here:

```
  for(n_cycles = 0; n_same_err<ANNEAL_TOLERANCE; n_cycles++)
  {
    halfway = 0;

    /* random point, but don't use noise points */

    wh = (long int)(ran() * n_sq);

    if (noise > 0.0)
    {
      noise_point = B_DATA(im_noise) + wh;
      while(*noise_point == 1)
      {
        wh = (long int)(ran() * n_sq);
        noise_point = B_DATA(im_noise) + wh;
      }
    }

    /* binary phase factor */

    on_point = C_DATA(im_mask_on) + wh;
    off_point = C_DATA(im_mask_off) + wh;
    re = ((*off_point).r - (*on_point).r) / n_sq;
    im = ((*off_point).i - (*on_point).i) / n_sq;
    if(flip_pixel(im_current, wh) > 0) { re = -re; im = -im; }

    /* coords of flipped pixel */

    m = wh / sq_size;
```

```
n = wh % sq_size;

/* effect of flipped pixel */

trans_point = C_DATA(im_transform);
trial_point = C_DATA(im_trial);

for(i = 0; i < paddrsize; i ++)
{
  j = paddr[i];

  k = j / sq_size;
  l = j % sq_size;

  a = -2*PI * (m*k + n*l) / sq_size;

  (trial_point[j]).r = (trans_point[j]).r + cos(a)*re + sin(a)*im;
  (trial_point[j]).i = (trans_point[j]).i + sin(a)*re + cos(a)*im;
}

if(anneal_type == ANNEAL_MAX)
  err = calc_max_err(im_trial);
else
  err = calc_tot_err(im_trial);
halfway = 1;

/* chance of a random error-increasing change */

if (T == 0.0)
  chance = 0.0;
else
{
  T = T/(1.0 + delta_t);
  err_delt = fabs(err_old - err);
  chance = exp(-err_delt/T);
}

/* test for error reduction or random fluctuation */

if(err_old > err || ran() < chance)
{
  err_old = err;

  /* copy im_trial -> im_transform */

  for(i = 0; i < paddrsize; i ++)
  {
    j = paddr[i];
    (trans_point[j]).r = (trial_point[j]).r;
    (trans_point[j]).i = (trial_point[j]).i;
  }
}
```

```
    else
    {
      /* copy im_transform -> im_trial */

      for(i = 0; i < paddrsize; i ++)
      {
        j = paddr[i];
        (trial_point[j]).r = (trans_point[j]).r;
        (trial_point[j]).i = (trans_point[j]).i;
      }
      flip_pixel(im_current, wh);
    }

    if(err < err_best)
    {
      err_best = err;
      n_same_err = 0;
    }
    else
      n_same_err++;

    if((n_cycles % ANNEAL_STEP) == 0)
    {
      printf("cycle: %d\terror: %3.12f\tchange: %3.12f\n",
        n_cycles, err_best, err_last-err_best);
      err_last = err_best;
    }
  }
```

End of iterated part. Hopefully a solution has been found!

```
  im_copy(im_current, im_solution);
  im_copy(im_trial, im_transform);

  if(autowrite)
    show_props();
}
```

The error calculation functions are as follows. Firstly, for the cost function:

$$C_{sum} = \sum_i [I(u_i, v_i) - \Phi(u_i, v_i)]^2:$$

```
double calc_tot_err(im_toterr)
image *im_toterr;
{
  int i;
```

```
  float re, im;
  double err_tot = 0.;
  fcomplex *err_point;


  err_point = C_DATA(im_toterr);

  for(i = 0; i < paddrsize; i++)
  {
    re = err_point[paddr[i]].r;
    im = err_point[paddr[i]].i;
    err_tot += pow((re*re)+(im*im)-(height*height), 2.0);
  }

  return(err_tot);
}
```

The variant cost function $C_{max} = \max_i[I(u_i, v_i) - \Phi(u_i, v_i)]^2$ is coded as:

```
double calc_max_err(im_input)
image *im_input;
{
  int i;
  double err_max, next;
  fcomplex *err_point;


  err_point = C_DATA(im_input);
  err_max = pow(height * height -
                pow(mod(err_point[paddr[0]]), 2.0), 2.0);

  for(i = 1; i < paddrsize; i++)
  {
    next = pow(height * height -
               pow(mod(err_point[paddr[i]]), 2.0), 2.0);

    if(next > err_max)
      err_max = next;
  }

  return(err_max);
}
```

A.1.2 Sample Output

This is the output upon generation of a 4 × 4 fanout.

```
cycle: 1950000  error: 0.000623558961   change: 0.000000000039
cycle: 1955000  error: 0.000623558881   change: 0.000000000080
cycle: 1960000  error: 0.000623529541   change: 0.000000029340
cycle: 1965000  error: 0.000623485614   change: 0.000000043927
cycle: 1970000  error: 0.000623441671   change: 0.000000043943
cycle: 1975000  error: 0.000623425606   change: 0.000000016065
cycle: 1980000  error: 0.000623374724   change: 0.000000050882
cycle: 1985000  error: 0.000623317610   change: 0.000000057115
cycle: 1990000  error: 0.000623317610   change: 0.000000000000
cycle: 1995000  error: 0.000623299402   change: 0.000000018208
cycle: 2000000  error: 0.000623299330   change: 0.000000000072
cycle: 2005000  error: 0.000623299330   change: 0.000000000000
```

```
target: 4x4t.pgm         noise level:  0%

0.162751        0.154088        0.156974        0.169651
0.152896        0.151008        0.153212        0.159916
0.159916        0.153212        0.151008        0.152896
0.169651        0.156974        0.154088        0.162751

eff.: 79.552%           unif.: 24.047%         err. sum: 0.0007096541
max at 12: 0.028781     min at 10: 0.022803    mean power: 0.024860
height: 0.177           coeff: 1.000           heightsum: 2.828
var. from mean: 0.002   var. from mean^2: 0.000003522693
```

A.2 Defect Simulation

A.2.1 Generation of 'Warped Cell' Image

```
/* Introduces a parabolic phase variation */

image *par_phase(im_input)
image *im_input;
{
  float *f_pt, scale;
  int x, y, i;


  f_pt = F_DATA(im_input);
  scale = (max_warp*4*PI)/area;
```

```
  for(i = 0; i < area; i++)
  {
    x = (i/side)-(side/2);
    y = (i%side)-(side/2);
    *(f_pt+i) = (x*x+y*y)*scale;
  }
  im_stats(im_input);
  return(im_input);
}

/* Writes a pgm file to encode a specified phase warp */

void main(argc, argv)
int argc;
char **argv;
{
  static image IM_OUT, *im_out;
  FILE *fp;
  int i;
  float *f_pt;
  char *info;


  im_out = &IM_OUT;
  info = malloc(sizeof(char)*16);

  parse_command_line(argc, argv);

  area = side*side;
  im_alloc(im_out, side, side, 1, FLOAT);
  im_stats(im_out);

  switch(warp_type)
  {
    case PAR:
      strcpy(info, "Paraboloidal");
      par_phase(im_out);
      break;
    default:
      printf("Warning: warp type %d not known\n", warp_type);
      break;
  }

  printf("%s warp with side %d and %1.2f wavelengths maximum\n",
    info, side, max_warp);
  im_pstats(im_out);

  if((fp = fopen(outfile, "wb"))== 0)
    printf("Error: output file %s not openable\n", outfile);
  else
  {
    printf("Output file %s used\n", outfile);
```

```
    im_dump (im_out, fp);
  }
  fclose (fp);
}
```

A.2.2 Incorporation into the BPH Generator

This is how the BPH generation code uses a 'warp' image to generate look-up tables of the on- and off- contributions of each pixel.

```
void build_on_off_planes(im_input)
image *im_input;
{
  int i;
  float *warp_point, ran_warp, r, c, d, ph;
  fcomplex *on_point, *off_point;


  warp_point = F_DATA(im_input);
  on_point = C_DATA(im_mask_on);
  off_point = C_DATA(im_mask_off);

c = (sqrt(2.0))/4.0;
d = N_BAR/DELTA_N;

for(i = 0; i < n_sq; i++, warp_point++, on_point++, off_point++)
{
```

This is the final version: a weighted combination of birefringence and phase warp contributions.

```
/*
 * 4th try: combination of 1 & 3
 */

  r = c*hypot(1+cos(*warp_point), sin(*warp_point));
  ph = atan2(-sin(*warp_point), 1+cos(*warp_point))
        + d*(*warp_point);

  (*on_point).r  = r*cos(ph);
  (*on_point).i  = r*sin(ph);
  (*off_point).r = r*cos(ph + PI);
  (*off_point).i = r*sin(ph + PI);
```

This version uses the amplitude values given by the Jones matrix treatment in Chapter 4.

```
/*
 * 3rd try: Maple value of x-coord
 */

/*     (*on_point).r  =  c + (c * cos(*warp_point));  */
/*     (*on_point).i  = -c * sin(*warp_point);     */
/*     (*off_point).r = -c - (c * cos(*warp_point));  */
/*     (*off_point).i =  c * sin(*warp_point);     */
```

This version treats the variation in amplitude caused by a varying thickness of birefringent material.

```
/*
 * 2nd try: variation of amplitude
 */

/*     r = sin((*warp_point)/2 + PI/2);  */

/*     (*on_point).r = 0;        */
/*     (*on_point).i = r;        */
/*     (*off_point).r = 0;        */
/*     (*off_point).i = -r;        */
```

This version simply encodes height as a phase offset — in effect a curved mirror.

```
/*
 * 1st try: 'warped backplane' phase
 */

/*  (*on_point).r = cos(*warp_point);     */
/*  (*on_point).i = sin(*warp_point);     */
/*  (*off_point).r = cos((*warp_point)+PI);  */
/*  (*off_point).i = sin((*warp_point)+PI);  */
  }
  im_stats(im_mask_on);
  im_stats(im_mask_off);

}
```

Appendix B

Acronyms

AC	Alternating Current
AM	Active Matrix
AO	Acousto-Optic
BDH	British Drug House
BPH	Binary Phase Hologram
CCD	Charge-Coupled Device
CMOS	Complementary Metal-Over-Silicon
DC	Direct Current
DRAM	Dynamic RAM
FLC	Ferroelectric Liquid Crystal
FLCoS	FLC over Silicon
FPGA	Field-Programmable Gate Array
FT	Fourier Transform
IO	Input/Output
ITO	Indium Tin Oxide
LC	Liquid Crystal
	.
	.
	.

.
.
.

LED	Light Emitting Diode
MEMS	Micro-Electro-Mechanical Systems
PC	Personal Computer
PCB	Printed Circuit Board
PGA	Pin Grid Array
RAM	Random Access Memory
RGB	Red/Green/Blue
SBP	Space-Bandwidth Product
SLM	Spatial Light Modulator
SRAM	Static RAM
SSFLC	Surface Stabilised FLC
TFT	Thin Film Transistor
UV	Ultra-Violet
VGA	Video Graphics Adaptor
VLSI	Very Large Scale Integration
WDM	Wavelength Division Multiplexing

Bibliography

[1] Glenn Zorpette (ed.). The Solid-State Century. *Scientific American*, Special Issue, 1997.

[2] Sing H. Lee (ed.). special section on Optical Interconnects and Packaging. *Optical Eng.*, 33(5):1511–1603, 1994.

[3] A. Miller. Optical computing — elements of an optical engine. *Nature*, 323(6083):13–14, 1986.

[4] Uzi Efron, Arthur D. Fisher, and Cardinal Warde (eds.). special issue on Spatial Light Modulators for Optical Information Processing. *Applied Optics*, 28(22), 1989.

[5] 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. C. Latham. A high performance spatial light modulator. In *Devices for Optical Processing*, volume 1562 of *SPIE*, pages 107–115, San Diego, 1991.

[6] D. C. Burns, J. Gourlay, A. O'Hara, I. Underwood, and D. G. Vass. A 256 × 256 SRAM-XOR pixel ferroelectric liquid crystal over silicon spatial light modulator. *Optics Comms.*, 119(5-6):623–632, 1995.

[7] Dwayne C. Burns, Mark L. Begbie, Ian Underwood, and David G. Vass. Ferroelectric liquid crystal over silicon spatial light modulator viewable under continuous illumination. *Ferroelectrics*, 181:93–97, 1996.

[8] M. C. Wu, L. S. Lee, and C. R. King. Integrated devices make an optical bench on a chip. *Laser Focus World*, pages 64–68, Feb 1996.

[9] Berkeley Silicon Applications Center. WWW pages. www-bsac.eecs.berkeley.edu/research/microphotonics.

[10] David May (Bristol University). Optical interconnections for information processing, June 1997. talk given at one-day LEOS seminar.

[11] Jurgen Jahns. Planar packaging of free-space optical interconnections. *Proc. IEEE*, 82(11):1623–1631, 1994.

[12] H. Scott Hinton, Tom J. Cloonan, Frederick B. McCormick Jr., Anthony L. Lentine, and Frank A. P. Tooley. Free-space digital optical systems. *Proc. IEEE*, 82(11):1632–1649, 1994.

[13] M. P. Y. Desmulliez, P. W. Foulk, and M. R. Taghizadeh. Optical clock distribution for multichip module. *Optical Review*, 3(6A):379–380, 1996.

[14] J. H. Yeh, R. K. Kostuk, and K. Y. Tu. Hybrid free-space optical bus system for board-to-board interconnections. *Applied Optics*, 35(32):6354–6364, 1996.

[15] G. J. Grimes et al. Photonic packaging using laser/receiver arrays and flexible optical circuits. *IEEE trans. on Components Packaging and Manufacturing Technology B*, 20(4):409–415, 1997.

[16] Optical Networks Inc. WWW pages. www.opticworks.com.

[17] Charles K. Gary. The dependence of diffraction efficiency on storage density in digital holographic memory using bacteriorhodopsin. In *Second International Conference on Optical Information Processing*, volume 2969 of *SPIE*, pages 30–39, St. Petersburg, 1996.

[18] D. Psaltis. Parallel optical memories. *Byte*, 17(9):179, 1992.

[19] S. M. Zhou, S. Campbell, P. C. Yeh, and H. K. Liu. 2-stage modified signed-digit optical computing by spatial data encoding and polarisation multiplexing. *Applied Optics*, 34(5):793–802, 1995.

[20] M. P. Y. Desmulliez, F. A. P. Tooley, J. A. B. Dines, N. L. Grant, D. J. Goodwill, D. Baillie, B. S. Wherrett, P. W. Foulk, S. Ashcroft, and P. Black. Perfect-shuffle interconnected bitonic sorter: Optoelectronic design. *Applied Optics*, 34(23):5077–5090, 1995.

[21] D. C. O'Brien and D. J. McKnight. A compact holographically routed optical crossbar using a ferroelectric liquid-crystal over silicon spatial light-modulator. *Institute of Physics Conference Series*, 139:187–190, 1995.

[22] Andrew J. Stevens. *Experimental Investigation of Free-Space Optical Routing Systems Using Static and Dynamic Binary Holographic Elements.* PhD thesis, University of Edinburgh, 1995.

[23] C. P. Barrett, P. Blair, G. S. Buller, D. T. Neilson, B. Robertson, E. C. Smith, M. R. Taghizadeh, and A. C. Walker. Components for the implementation of free-space optical crossbars. *Applied Optics*, 35(35):6934–6944, 1996.

[24] Iain Rankin. PhD thesis, University of Edinburgh, 1997.

[25] Steve J. Elston. The optics of ferroelectric liquid crystals. *J. Modern Optics*, 42(1):19–56, 1995.

[26] N. A. Clark and S. T. Lagerwall. Submicrosecond bistable electro-optic switching in liquid crystals. *Appl. Phys. Lett.*, 36(11):899–901, 1980.

[27] D. J. McKnight. Continuous view of DC-balanced images on a ferroelectric liquid-crystal over silicon spatial light-modulator. *Optics Letters*, 19(18):1471–1473, 1994.

[28] Georg Bodammer. *Device Oriented Experimental Investigation of the Alignment of Liquid Crystals.* PhD thesis, University of Edinburgh, 1997.

[29] D. C. Burns, I. Rankin, I. Underwood, and D. G. Vass. 512 × 512 DRAM FLCOS SLM preliminary information, November 1995. internal group document.

[30] 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 by the use of dielectric coatings and chemical-mechanical polishing. *Applied Optics*, 32(28):5549–5555, 1993.

[31] G. Bodammer, J. Gourlay, A. O'Hara, and D. G. Vass. Study of indium-tin-oxide surface morphology and its influence on obliquely evaporated SiOx alignment layers for ferroelectric liquid crystals. *Ferroelectrics*, 178:135–144, 1996.

[32] Iain Rankin. Generic SLM interface, 1996. notes for PCB design.

[33] Mark Begbie. PhD thesis, University of Edinburgh, 1998.

[34] Robert A. J. Matthews. Tumbling toast, Murphy's law and the fundamental constants. *European J. Physics*, 16:172–176, June 1995.

[35] Timothy J. Drabik. Optoelectronic integrated systems based on free-space interconnects with an arbitrary degree of space variance. *Proc. IEEE*, 82(11):1595–1622, 1994.

[36] Robert K. Tyson. *Principles of Adaptive Optics.* Academic Press, 1991.

[37] J. W. Goodman. *Introduction to Fourier Optics.* McGraw-Hill, second edition, 1996.

[38] Pochi Yeh. Extended Jones matrix method. *J. Opt. Soc. Am.*, 72(4):507–513, 1982.

[39] Ralph J. Gagnon. Liquid-crystal twist-cell optics. *J. Opt. Soc. Am.*, 71(3):348–353, 1981.

[40] Interfire. Apex, 1996. fringe analysis software.

[41] M. Born and E. Wolf. *Principles of Optics.* Pergamon Press, 1990.

[42] P. Smith, M.R. Hart, G. Bodammer, and D.G. Vass. Optical byte correlator for serial data streams employing Karnaugh mapping using FLC SLMs. *Ferroelectrics*, 181:29–36, 1996.

[43] J. M. Miller, M. R. Taghizadeh, Turunen J., N. Ross, E. Noponen, and A. Vasara. Kinoform array illuminators in fused-silica. *J. Modern Optics*, 40(4):723–732, 1993.

[44] S. Sinzinger and V. Arrizon. High-efficiency detour-phase holograms. *Optics Letters*, 22(12):928–930, 1997.

[45] Uwe Dahms, Gunther Wernicke, Nazif Demoli, and H. Gruber. Modulation characteristics of the Sanyo liquid-crystal television. In *Second International Conference on Optical Information Processing*, volume 2969 of *SPIE*, pages 170–174, St. Petersburg, 1996.

[46] Sergei S. Samus. *Computer Design and Optimisation of Holographic Phase Elements*. PhD thesis, University of Edinburgh, 1995.

[47] Frank Wyrowski. Diffractive optical elements: Iterative calculation of quantized, blazed phase structures. *JOSA A*, 7(6):961–969, 1990.

[48] Henry Stark (Ed.). *Image Recovery: Theory and Application*. Academic Press, 1987.

[49] R. W. Gerchberg and W. O. Saxton. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35:237–246, 1972.

[50] A. Kirk and T. Hall. Design of binary computer generated holograms by simulated annealing: coding density and reconstruction error. *Optics Comms.*, 94:491–496, 1992.

[51] Michael Creutz. Microcanonical cluster Monte Carlo simulation. *Phys. Rev. Lett.*, 69(7):1002–1005, 1992.

[52] M. Dames, R. Dowling, P. McKee, and D. Wood. Efficient optical elements to generate intensity weighted spot arrays: Design and fabrication. *Applied Optics*, 30(19):2685–2691, 1991.

[53] U. Krackhardt, J. Mait, and N. Streibl. Upper bound on the diffraction efficiency of phase-only fan-out elements. *Applied Optics*, 31(1):27–37, 1992.

[54] Gordon D. Love. Wave-front correction and production of Zernike modes with a liquid-crystal spatial light modulator. *Applied Optics*, 36(7):1517–1524, 1997.

[55] Gordon D. Love, Nigel Andrews, Philip Birch, David Buscher, Peter Doel, Colin Dunlop, John Major, Richard Myers, Alan Purvis, Ray Sharples, Andrew Vick, Andrew Zadrozny, Sergio R. Restaino, and Andreas Glindemann. Binary adaptive optics: atmospheric wave-front correction with a half-wave phase shifter. *Applied Optics*, 34(27):6058–6066, 1995.

[56] E.G.S. Paige, R. H. Scarbrough, and G. G. Yang. Feedback generated holograms. *Electronics Letters*, 30(14):1174–1175, 1994.

[57] Will Hossack. C image library, 1996. image manipulation software.

[58] Alastair R. MacGregor. Method for computing the optical properties of a smectic C* liquid-crystal cell. *JOSA A*, 6(9):1493–1503, 1989.

[59] B. S. Wherrett and F. A. P. Tooley (Eds.). *Optical Computing*. SUSSP, 1989.

[60] H. Y. Chen and E. G. S. Paige. Creation of 3-D radiation fields to specification and demonstration using an optical SLM. *Electronics Letters*, 30(9):735–736, 1994.

[61] Jorgen Bengtsson. Kinoform design with an optimal-rotation-angle method. *Applied Optics*, 33(29):6879–6884, 1994.

[62] Jorgen Bengtsson. Kinoform-only Gaussian-to-rectangle beam shaper for a semiconductor laser. *Applied Optics*, 35(20):3807–3814, 1996.

[63] Ho Hyung Suh, Chong Hoon Kwak, and El-Hang Lee. Combined binary-phase holograms for free-space optical interconnection. *Optics Letters*, 20(20):2131–2133, 1995.

[64] Bjorn Lofving. Measurement of the spatial phase modulation of a ferroelectric liquid-crystal modulator. *Applied Optics*, 35(17):3097–3103, 1996.

[65] James E. Morris and Michael R. Feldman. Reconfigurable optical interconnects by a combined computer-generated hologram and spatial light modulator method. *Applied Optics*, 33(17):3683–3694, 1994.

[66] Nobukazu Yoshikawa and Toyohiko Yatagai. Phase optimization of a kinoform by simulated annealing. *Applied Optics*, 33(5):863–868, 1994.

[67] Ignacio Moreno, Juan Campos, Christophe Gorecki, and Maria J. Yzuel. Effects of amplitude and phase mismatching errors in the generation of a kinoform for pattern recognition. *Jpn. J. Appl. Phys.*, 34:6423–6432, 1995.

[68] Michael R. Feldman, Clark C. Guest, Timothy J. Drabik, and Sadik C. Esener. Comparison between electrical and free space optical interconnects for fine grain processor arrays based on interconnect density capabilities. *Applied Optics*, 28(18):3820–3829, 1989.

[69] Demetri Psaltis and Yong Qiao. Adaptive multilayer optical networks. In E. Wolf, editor, *Progress in Optics*, volume 31, chapter 4, pages 227–261. Elsevier Sci. Pub. B.V., 1993.

[70] Iain Cramb and Colin Upstill. An optoelectronic binary image algebra architecture. *Optical Computing and Processing*, 1(2):115–126, 1991.

[71] Alain Bergeron, Jonny Gauvin, Francois Gagnon, Denis Gingras, Henri H. Arsenault, and Michel Doucet. Phase calibration and applications of a liquid-crystal spatial light modulator. *Applied Optics*, 34(23):5133–5139, 1995.

[72] Colin Soutar and Kanghua Lu. Determination of the physical properties of an arbitrary twisted-nematic liquid crystal cell. *Optical Eng.*, 33(8):2704–2712, 1994.

[73] H. M. Kim, J. W. Jeong, M. H. Kang, and S. I. Jeong. Phase correction of a spatial light modulator displaying a binary phase-only filter. *Applied Optics*, 27(20):4167–4168, 1988.

[74] T. G. Slack and C. R. Chatwin. Interaction of multiple distortions in spatial light modulators. *Applied Optics*, 34(8):1341–1351, 1995.

[75] Roderick P. Webb. Performance of an optoelectronic neural network in the presence of noise. *Applied Optics*, 34(23):5230–5240, 1995.

[76] Christopher Tocci and H. John Caulfield (Eds.). *Optical Interconnection*. Artech House, 1994.

[77] H. Inbar and M. R. Taghizadeh. Effects of dead zones in multiple-quantum-well binary-phase modulators on optical interconnections. *Applied Optics*, 37(5):912–920, 1998.

UNIVERSITY OF EDINBURGH THESIS

Author (surname, initials): Smith P Degree: Ph.D Year: 1999

This thesis is an unpublished typescript and the copyright is held by the author. *All persons consulting this thesis, or having copies made, must sign the Copyright Declaration below.*

Copying regulations:

(1) This thesis may be copied in whole or in part for the use of individuals and for libraries wishing to add this thesis to their stock. *Copying must be done by Edinburgh University Library.*

(2) ~~This thesis must not be copied, either in whole or in part, without the applicant obtaining the author's written permission. *If permission is granted, then copying must be done by Edinburgh University Library.* (The bibliography/list of works consulted may be copied without the author's permission, provided the copying is done by Edinburgh University Library.)~~

(3) ~~This thesis may be copied only in so far as the copying does not contravene the 1956 Copyright Act. *Copying must be done by Edinburgh University Library.*~~

Copyright Declaration: I undertake fully to observe the author's copyright in this thesis, not to publish the whole or any part of it without the author's written permission, and not to allow any other person to use any copy made for me.

Date	Name and Address (BLOCK CAPITALS)	Signature	Pages copied (if any)

Smith P
PhD 1999