Design and characterisation of a ferroelectric liquid crystal over silicon spatial light modulator

Dwayne C. Burns
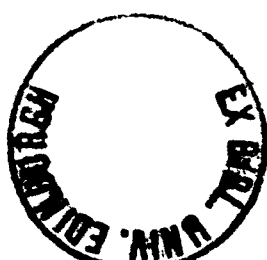
Abstract

Many optical processing systems rely critically on the availability of high performance, electrically-addressed spatial light modulators. Ferroelectric liquid crystal over silicon is an attractive spatial light modulator technology because it combines two well matched technologies. Ferroelectric liquid crystal modulating materials exhibit fast switching times with low operating voltages, while very large scale silicon integrated circuits offer high-frequency, low power operation, and versatile functionality.
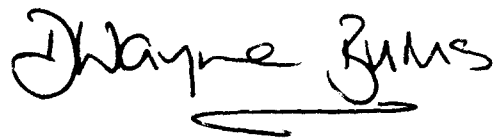
This thesis describes the design and characterisation of the SBS256 – a general purpose 256 × 256 pixel ferroelectric liquid crystal over silicon spatial light modulator that incorporates a static-RAM latch and an exclusive-OR gate at each pixel. The static-RAM latch provides robust data storage under high read-beam intensities, while the exclusive-OR gate permits the liquid crystal layer to be fully and efficiently charge balanced.

The SBS256 spatial light modulator operates in a binary mode. However, many applications, including helmet-mounted displays and optoelectronic implementations of artificial neural networks, require devices with some level of grey-scale capability. The 2 kHz frame rate of the device, permits temporal multiplexing to be used as a means of generating discrete grey-scale in real-time.

A second integrated circuit design is also presented. This prototype neural-detector backplane consists of a 4 × 4 array of optical-in, electronic-out processing units. These can sample the temporally multiplexed grey-scale generated by the SBS256. The neurons implement the post-synaptic summing and thresholding function, and can respond to both positive and negative activations – a requirement of many artificial neural network models.

Declaration

I declare that this thesis has been completed by myself and that, except where indicated to the contrary, the research documented is entirely my own.

Dwayne C. Burns

Acknowledgements

I would like to thank the following people for their invaluable assistance during the course of my PhD project:

Table of Contents

List of Figures

List of Tables

to Gail

Chapter 1

Introduction

1.1 Background

Many optical processing architectures and systems require high performance, electrically addressed spatial light modulators (SLMs). The hybrid SLM technology of ferroelectric liquid crystal over silicon (FLCOS) is attractive because it combines two well matched technologies. Ferroelectric liquid crystals exhibit fast switching times coupled with low switching voltages, while very large-scale silicon integration permits the implementation of large active pixel arrays that can be addressed at high frame rates. These electrically addressed devices function as electronically-written, optically-read memory devices that are suitable for use in display applications, and as input or filter planes in optical systems.

Many of the medium-resolution FLCOS SLMs designed so far have been based on the one-transistor dynamic random access memory (DRAM) circuit. The DRAM-type circuit offers high pixel density but has some performance limitations, including photo-induced charge leakage and limited FLC drive capability. In this thesis, I present a general-purpose 256×256 FLCOS device that incorporates a static-RAM latch and an XOR-gate at each pixel. The pixel circuit overcomes the limitations of the DRAM-type pixel at the expense of area and fill-factor. However, the high pixel count of this SRAM-type device raises many backplane design issues that are not as relevant to DRAM-type devices or smaller SRAM-type arrays. The device is particularly suitable for applications where the

1

read-beam is intense, say at the input stage of a cascaded SLM system, or in system where a pulsed light source cannot be used.

In recent years there has been a dramatic increase in research on artificial neural networks (ANNs). ANNs are useful tools for solving many classification, modelling, and prediction problems. They normally consist of parallel arrangements of many simple processing units (or neurons) connected together by variable strength connections (weights). Most work has centred on understanding the behaviour and limitations of many different network architectures and algorithms, through theoretical analysis and software simulation on digital processors. However, to realise the full potential of ANNs, efficient hardware implementations must be developed. Among other approaches, including VLSI neural chips, the optoelectronic approach is expected to play an important role in the implementation of medium-scale ANNs (100's–1000's of neurons). Optoelectronic implementations are especially attractive for classification applications where the input exists in the optical domain. FLCOS SLM technology can be used to implement the input, interconnection weights, and neuron functional blocks required by many optoelectronic ANN systems. Furthermore, if the image is converted to the coherent optical domain, powerful preprocessing functions can be performed on it 'for free' by using the techniques of Fourier optics.

1.2   Objectives

The objectives of my project were to:

1. Develop a general-purpose 256 × 256 FLCOS SLM based on the static-RAM type pixel.

2. Demonstrate real-time temporally multiplexed grey-scale on such a device.

3. Develop a prototype neural-detector backplane capable of sampling the temporally multiplexed grey-scale generated by the SLM.

4. Optimise both devices for incorporation into a future optoelectronic ANN system.

1.3   Thesis outline

I begin Chapter 2 with an overview of some spatial light modulator technologies, and then introduce the technology of FLCOS. I then give a description of some pixel designs that have been incorporated into FLCOS SLM devices. These range from the single-transistor DRAM-type pixel, to so-called 'smart pixels' that can perform primitive processing functions.

In Chapter 3, I review some applications for which FLCOS SLMs are particularly suitable, including optical crossbar switching, and optical correlators. Then I consider some methods for generating grey-scale on FLCOS binary mode devices. The chapter closes with a review of some optoelectronic ANN implementations.

In Chapters 4 and 5, I present the SBS256 SLM – a 256×256 pixel device that comprises a SRAM latch and an XOR-gate at each pixel. Chapter 4 concentrates on backplane design issues such as the pixel circuit, architecture, transient current spikes, and power dissipation. In Chapter 5, I present electrical and optical characterisation results. The electrical results include current spike measurements, operating frequency, backplane electrical yield, and the effects of post-processing planarisation. The optical characterisation results include contrast ratio, frame rate, temporally-multiplexed grey-scale, and uniformity.

Chapter 6 covers the development of OANN - a 4×4 prototype neural-detector backplane. The backplane is capable of sampling a 2's complement variant of the temporally multiplexed grey-scale generated by the SBS256. Photo-induced charge-leakage is identified as a potential problem for the neuron activity storage circuits. A simple technique is presented to significantly reduce this

problem. Electrical and optical characterisation from the neuron circuits are also presented.

Finally in Chapter 7, I evaluate the performance of the two integrated circuit designs presented in earlier chapters and discuss possible improvements and future work using the devices.

Chapter 2

Spatial light modulators - an overview

The advantages gained from the inherent parallelism of many proposed optical processing and computing architectures rely critically on the availability of high performance two-dimensional spatial light modulators (SLMs). SLMs are used as data displays, spatial filters, incoherent to coherent image converters, two-dimensional logic elements, and as routing elements. In this chapter I give an overview of some SLM technologies that have shown promise before focusing on the technology of ferroelectric liquid crystal over silicon (FLCOS).

2.1 Introduction

SLMs can modulate properties of an optical wavefront such as intensity, phase, or polarisation. Spatial light modulation can be accomplished via electrooptic, acoustooptic, magnetooptic, mechanical, photorefractive, or electroabsorptive effects in a variety of materials. There are two major classes of SLMs: optically and electrically addressed. There is such a wide variety of SLM technologies, that it is impractical to discuss all of them in detail (see SLM review papers [18,23, 74], or Applied Optics special issues on Information Processing for more recent

advances). In this chapter I will concentrate on electrically addressed devices, but will also consider optically addressed devices where appropriate.

The main advantage of electrically addressed devices is their ability to interface with both electrical and optical systems. The characteristics of electrically addressed SLMs are likely to define the performance of the systems they are used in. These characteristics include pixel count, frame rate, contrast ratio, physical size, cost, and power dissipation. I begin with an overview of the major SLM technologies.

2.2 Major SLM technologies

In this section, I consider three SLM technologies that have shown promise: digital micromirror devices, Si-PLZT devices, and self electrooptic effect devices.

2.2.1 Digital micromirror devices

The digital micromirror device (DMD[1]), developed by Texas Instruments [84], is a monolithically fabricated SLM that employs electrostatically deflected mirrors as the light modulating elements. Each mirror is suspended by thin hinges over one or more address electrodes which are in contact with the underlying address circuitry. When an electrode is addressed, the overlying mirror, which is held at a fixed potential, is attracted to it. The mirror element design can be tailored to a particular application. For example, torsion beam DMDs are built to modulate amplitude, while flexure beam DMDs modulate mostly phase, and cantilever DMDs modulate amplitude and phase in a coupled fashion. The torsion-beam mirrors can be deflected in either of two directions about the axis defined by the support hinges (see Fig. 2-1). When a pixel mirror is tipped in one direction,

---

[1]DMD can also refer to deformable mirror devices, the forerunners to the digital micromirror technology

incident light is thrown out of the optical imaging system, so that the pixel appears dark. When it is tipped in the other direction, light passes through the imaging system so that the pixel appears bright.
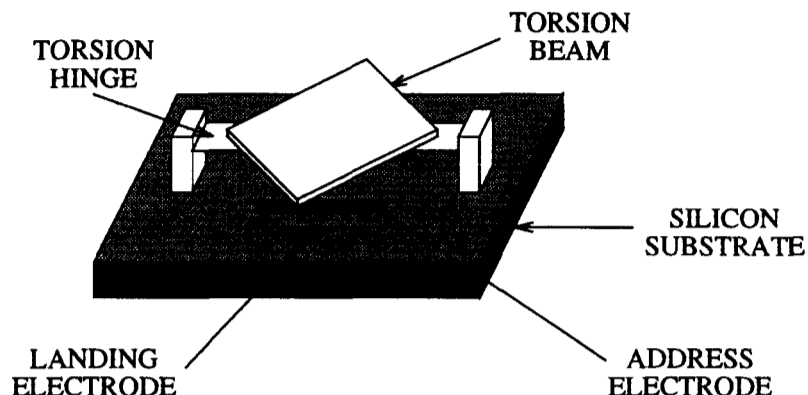
Figure 2–1: Cross-section of a torsion-beam deformable mirror pixel. The mirror can be deflected in either of two directions about the support hinge axis by the underlying address electrodes.

DMDs have been addressed using charge coupled device (CCD) and static random access memory (SRAM) technologies. Boysel [9] developed a 128 × 128 pixel virtual phase CCD-addressed device to establish an analogue charge image, which is then transferred to the overlying electrodes to deflect clover-leaf deformable mirrors. It accepts 16 analogue input signals at a clocking rate of up to 4 MHz. The underlying CCD array permits frame addressing so, unlike row addressed architectures, there is no skew, and a high 96% duty cycle. The contrast ratio of this device is approximately 2:1, as it was designed as a phase modulator. A 768 × 576 pixel SRAM addressed DMD SLM with torsion-beam mirrors, has recently been developed by Texas Instruments capable of displaying PAL broadcast quality pictures for use in a projector display system [112]. Each frame of video is divided into three colour fields and grey scale is achieved by pulse width modulation within each colour field. Pixels can be switched in 10 μs and exhibit a contrast ratio of 50:1. Texas Instruments are also fabricating a 2048 × 1152 pixel device for a prototype high definition television display system [116]. The die will be 37 by 22 mm in size, and so is likely to suffer yield and reliability problems.

2.2.2 Si-PLZT devices

Si-PLZT devices combine mature silicon technology for processing and control, with the modulator material, lead lathanum zirconate titanate (PLZT). PLZT is a transparent, electrooptic, ferroelectric, ceramic material. It exhibits low switching times ($\leq$ 100 ns), high contrast ratios ($\geq$ 500:1), and analogue modulation capability. However, it also requires high drive voltages (20–100 V) to give reasonable transmission through crossed polarisers.

The main technique used for combining Si and PLZT substrates involves flip-chip bonding an Si chip (or chips) to a PLZT wafer – flip-chip bonding is widely used in the emerging technology of multi-chip-modules. Control and driver circuits are placed on the Si chip and connected to the PLZT modulator using metal solder bumps. The basic configuration of a unit cell is shown in Fig. 2–2. The drive transistor gate can be varied by electrically or optically addressed [20] control circuitry. The resulting electric field between the switching and ground electrodes, spaced 15–50 μm apart, can vary the refractive index of the overlying PLZT material to phase modulate a read-beam. Amplitude modulation can be achieved by passing the read-beam through suitably oriented crossed polarisers. The output signal rise/fall times are limited by the drive circuitry. To give an idea of the transmission characteristics inherent with these devices, consider one of the prototype modulators fabricated by Jin *et al* [48] where an electrode spacing of 15 μm was used. For an applied voltage of 0 V the transmission was $\sim$ 0%, while for 40 V and 120 V it was $\sim$ 4.5% and $\sim$ 60% respectively.

With the flip-chip bonding approach, there is a voltage compatibility problem between the 20–100 V required for the modulator drive circuits (which may contain more than the single transistor illustrated in Fig. 2–2) and the control logic circuits operating at 5 V. Jin *et al* have also studied methods of bonding thin Si films with driver circuits directly onto bulk PLZT and then flip-chip bonding an Si chip containing the control circuitry. The low-voltage control signal is passed across the solder bump to control the high-voltage drive circuit. This will allow them to use high-density, foundry-processed VLSI chips with sophisticated

LIGHT | ↑ LIGHT
IN ↓ | OUT

PLZT
WAFER

SWITCHING
ELECTRODE
&
SOLDER
BUMP

n+   n+

Si CHIP

GROUND
ELECTRODE

DIELECTRIC
MIRROR

DRIVE TRANSISTOR
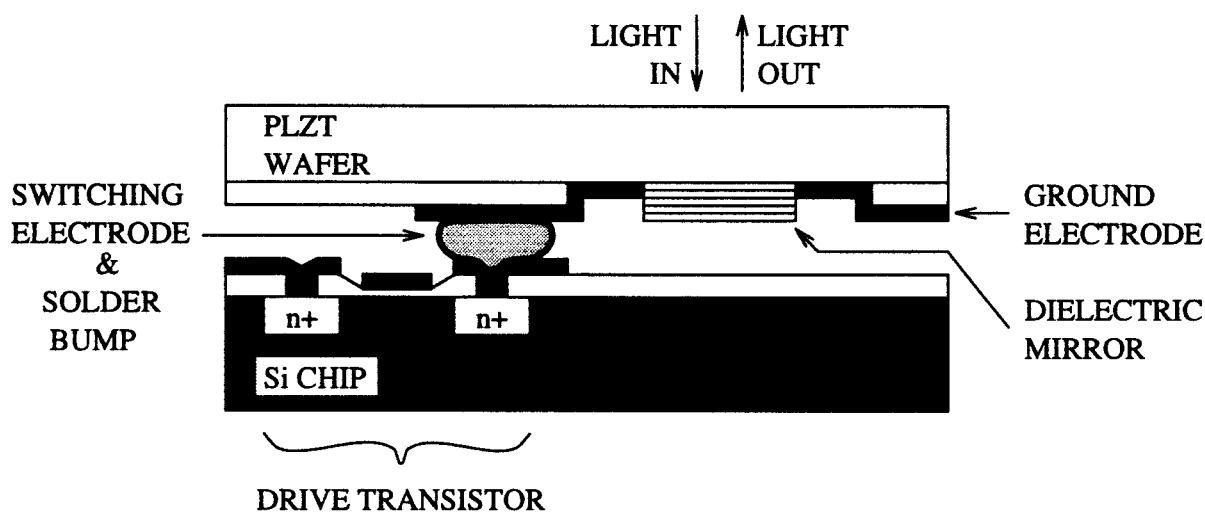
Figure 2–2: Cross section of a unit cell of a Si-PLZT SLM implemented using flip-chip bonding of the Si chip to the PLZT wafer. The electric field between the switching and ground electrodes varies the refractive index of the overlying PLZT material.

control circuits. Esener [19] reviews several other approaches used for combining the Si and PLZT substrates, each with their attendant advantages and disadvantages. He also describes a technique that incorporates self-tuned Fabry-Perot cavities into the modulator structures to reduce the required drive voltage by a factor of at least five. An optically recorded volume hologram is used as one of the highly reflective coatings of the Fabry-Perot cavity. It is recorded in place using interference between read-beam illumination and the reflected output. The phase angle of the resulting hologram automatically compensates for thickness variations across the modulator layer. This technique could also be applied to a wide variety of other SLM technologies where modulator uniformity is difficult to achieve over large arrays.

The large drive transistors required for high speed operation, coupled with the considerable power dissipation at high operating voltages, will probably limit this technology to relatively small, low fill-factor arrays. However, their analogue modulation capability makes them suitable for many optical processing and neural network applications.

2.2.3   Self electrooptic effect devices

In the past few years several variants of the self electrooptic effect device (SEED) technology have been developed. SEEDs make use of the non-linear effect called the quantum confined Stark effect, which describes the changes in optical properties of quantum wells when an electric field is applied perpendicular to multiple quantum well (MQW) layers. Since the quantum confined Stark effect is an almost instantaneous process, the modulation frequency is limited only by the frequency with which the electric field can be modulated. Devices must be operated at a relatively narrow optical bandwidth, which is determined by the structure and properties of the MQW layers.

The MQW structure typically consists of 50–100 alternate layers of GaAs and $Al_x Ga_{1-x} As$ [2] grown by molecular beam epitaxy. In a SEED, a MQW is placed within a p-i-n diode structure, and can be configured as optical input and output devices [70]. When a SEED is connected to an external load, and is under strong optical illumination, the photocurrent generated in the SEED induces a large voltage drop across the load. The field across the MQW is reduced which, in turn, increases its absorption and photocurrent, so that by positive feedback, the transmittance of the device becomes significantly reduced. The modulation is fast, but the contrast ratio is low (2:1), and operation is restricted to a particular wavelength (typically 850 nm for $GaAs/Al_x Ga_{1-x} As$ MQWs).

In symmetric-SEEDs (S-SEEDs), two SEEDs are connected back to back so that each acts as a light controlled load resistor for the other. The device can be switched between two stable states by momentarily changing the ratio of the two optical inputs. As switching is controlled by the ratio of two inputs rather than absolute intensities, device operation is more robust for real optical systems. The optical input and output power supply beams must be clocked, but this is not seen as a problem for most digital optical systems, where the architectures

---

[2] Where x is typically $0.25 \leq x \leq 0.35$

and algorithms tend to be synchronous anyway. A 256×128 optically addressed array and a 16×8 electrically addressed array have been fabricated [11].
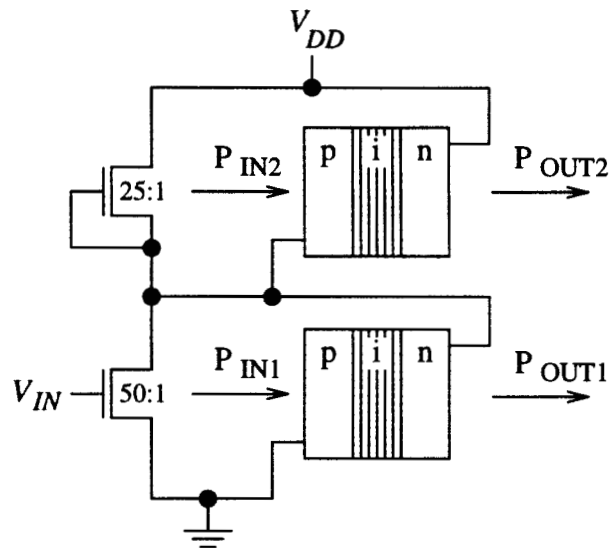
Figure 2–3: Cross section of a unit cell of an electrically addressed GaAs FET-SEED amplified differential modulator. implemented by Lentine *et al* [56]. Note that the drive transistors are wide compared to their length (W:L given in microns) to give high current sourcing capability for high frequency operation.

S-SEEDs have also been successfully integrated with various electronic technologies. When S-SEEDs are monolithically integrated with GaAs field effect transistors (FETs), very high speed, electrically addressed, differential modulator arrays can be implemented. A small 6×6 [56] electrically addressed integrated array has been fabricated with a differential amplifier at each pixel (see Fig. 2–3). The input voltage swing to the differential amplifiers can be as small as a few hundred millivolts, while the voltage across the MQW modulators can be as high as 10 V. This permits the array to be interfaced to standard Si bipolar emitter-coupled logic or GaAs FET electronic families that have typical voltage swings of less than 1 V. The device can operate at modulation frequencies of up to 2 GHz, limited only by the current sourcing capabilities of the drive transistors. Note that as with high speed Si-PLZT drive circuits, the transistors must be made wide to give high current sourcing capability for high frequency operation. They therefore take up significant area on the GaAs substrate which tends to limit pixel density, resolution and active area fill factor. GaAs circuits also use only

n-channel transistors, so circuits such as inverters sink current in one of their states, and thus may be susceptible to overheating if care is not taken in their design.

Goodwill *et al* [30] have recently flip-chip bonded a 16 × 16 array of S-SEEDs to a silicon driver backplane. Although inherently slower than GaAS FET technologies, silicon CMOS processes offer higher packing density, lower cost, higher yield and less overheating problems. The modulators are designed to work in the 1047–1064 nm region to match the high power available from diode pumped Nd:YLF and Nd:YAG lasers. They believe the device could work at up to 70 MHz with contrast ratios of 2–4:1.

S-SEED technology shows great potential for smart pixel and digital optical computing applications where very high frequency operation and low switching energies are more important than high contrast ratio and pixel resolution.

2.3   Liquid crystal devices

The design of electrically addressed FLCOS SLMs has been influenced by the liquid crystal display (LCD) industry. Many of the liquid crystal cell configurations and addressing schemes are similar, so before focusing on FLCOS SLMs, some configurations and schemes will be reviewed.

Liquid crystals (LCs) are attractive modulator materials for SLMs and displays as they have high birefringence, low switching energies, and low drive voltage requirements. LCs research and development has been driven by the displays industry as it attempts to challenge the cathode ray tube's dominance of the television monitor market with flat, light-weight, low-power devices. There can be major differences between the requirements of good SLMs and those for displays [15]. Displays typically have to operate with white light illumination, have a wide viewing angle, be shock resistant, and may have to withstand extremes of temperature. They usually only require 25–50 Hz frame rates, and

contrast ratios of 5–10:1. SLMs, on the other hand, usually require high frame rates, high contrast ratios, and for coherently illuminated systems, phase flatness is also important.

LCs are a class of materials that possess properties intermediate of those of typical crystalline and liquid phases. LC molecules are mainly organic in nature with long, thin, rod-like structures formed by linking aromatic rings to hydrocarbon chains. Their anisotropic[3] structure can lead to anisotropy of certain bulk physical properties such as dielectric permittivity and refractive index. The combination of such anisotropy, characteristic of crystalline materials, with the flow properties of liquids make liquid crystals uniquely useful in a variety of applications.

Three classes of liquid crystal structures exist: nematic, cholesteric, and smectic phases, each with a characteristic crystalline structure. Commercially available nematic LCs appeared in the early 1970's and have been used extensively in the relatively mature LCD technology. However, nematic LCs exhibit relatively low contrast (5:1), and 10–100 ms switching times, restricting potential device frame rates to 10–50 Hz. Since the early 1980's, smectic phase ferroelectric liquid crystals (FLCs) have been investigated for use in high frame rate applications requiring much faster switching times (10–100 μs), and higher contrast ratios.

2.3.1 Smectic phase - the surface stabilised effect

In the smectic phase, the long axes of the molecules align themselves parallel to each other. Furthermore, the centres of gravity of these molecules exist in one plane and the planes, or layers, pile up on each other. However, there is no general regularity of spacing between the molecules in one layer. There is therefore a two-dimensional ordering in the position of the liquid crystal molecules.

---

[3]Anisotropic materials have different properties for different directions in the material.

Certain chiral[4] smectic liquid crystals exhibit ferroelectric properties, which can be utilised in electrooptic devices. Ferroelectric liquid crystals (FLCs) possess a permanent or spontaneous polarisation, and most of the commonly used ferro-electric phase is the chiral smectic C (S_C^*) phase. In a bulk sample there is no net polarisation because the director n, spirals through the material from layer to layer. However, by use of suitable boundary conditions, the helical structure can be suppressed and thin layers of S_C^* can show a net polarisation.

The most commonly used device geometry is the surface stabilised FLC (SS-FLC), first described by Clark and Lagerwell [13]. A transmissive SSFLC cell typically comprises a 2–3 μm of FLC sandwiched between two glass substrates, each coated with a transparent electrode and a SiO_x or rubbed polyamide alignment layer. The device is similar to a twisted nematic device [89] except that the cells are thinner and the alignment layers are parallel. In the idealised SSFLC configuration, the substrates are close enough to unwind the FLC's intrinsic helical structure permanently so that the molecular axes (or directors) can only exist in two energetically degenerate states oriented at $\pm\phi$ to the layer normal, where ϕ is the molecular tilt angle. ϕ, is typically close to 22.5° over large temperature ranges, permitting the optic axis to be electronically rotated through approximately 45°, by an applied electric field of the appropriate sign. Both amplitude and phase modulation are possible.

Amplitude modulation

Amplitude modulation can be achieved by placing the cell between crossed polarisers (see Fig. 2–4). If the polarisation of normally incident light is chosen either parallel or perpendicular to one of the voltage selected optical axis states, it will be transmitted through the FLC unaffected. When the other optical state is selected by applying a field of the opposite sign, the optical axis is rotated to a 45°

---

[4]A chiral molecule cannot be superimposed upon its mirror image by translation, rotation, or reflection.

angle from the incident polarisation so that both the ordinary and extraordinary modes are excited. For the correct FLC layer thickness d, a total phase shift of π will accumulate between these two modes, and the incident light's polarisation will be rotated by 90°. The transmission I through the crossed-polarisers is given by:

$$I = I_0 \sin^2(4\phi) \sin^2(\frac{\Delta n d \pi}{\lambda}) \tag{2.1}$$

where I_0 is the input intensity, Δn is the FLC birefringence, and λ is the wavelength of the incoming light. The contrast ratio is theoretically infinite, but in real devices, it is limited by the uniformity of the director alignment and the quality of the polarisers. Gourlay [31] has measured contrast ratios of up to 500:1 for FLC test cells fabricated in-house in the Applied Optics Group facilities, at the University of Edinburgh.

Phase modulation

Two methods can be used to generate phase modulation. The first has one of the switchable optical axes parallel to the polarisation vector of the incoming light. When the FLC is switched to the other state, the phase of the light is changed by an amount dependent on both the FLC rotation angle and the birefringence. In the other method, the polarisation vector of the incident light and first polariser are set to bisect the FLC's two switchable optic axes. The propagation of light through the FLC produces elliptically polarised light from the incident linear polarisation. The sense of the polarisation depends on which side the FLC's director is switched to. The elliptically polarised light then passes through an analyser, orthogonal to the first polariser, which reduces the elliptical polarisation to linear polarisation exhibiting a 0 or π phase difference, independent of switching angle or optical path length. For 100% transmission, ϕ must be 45°, rather than 22.5°.
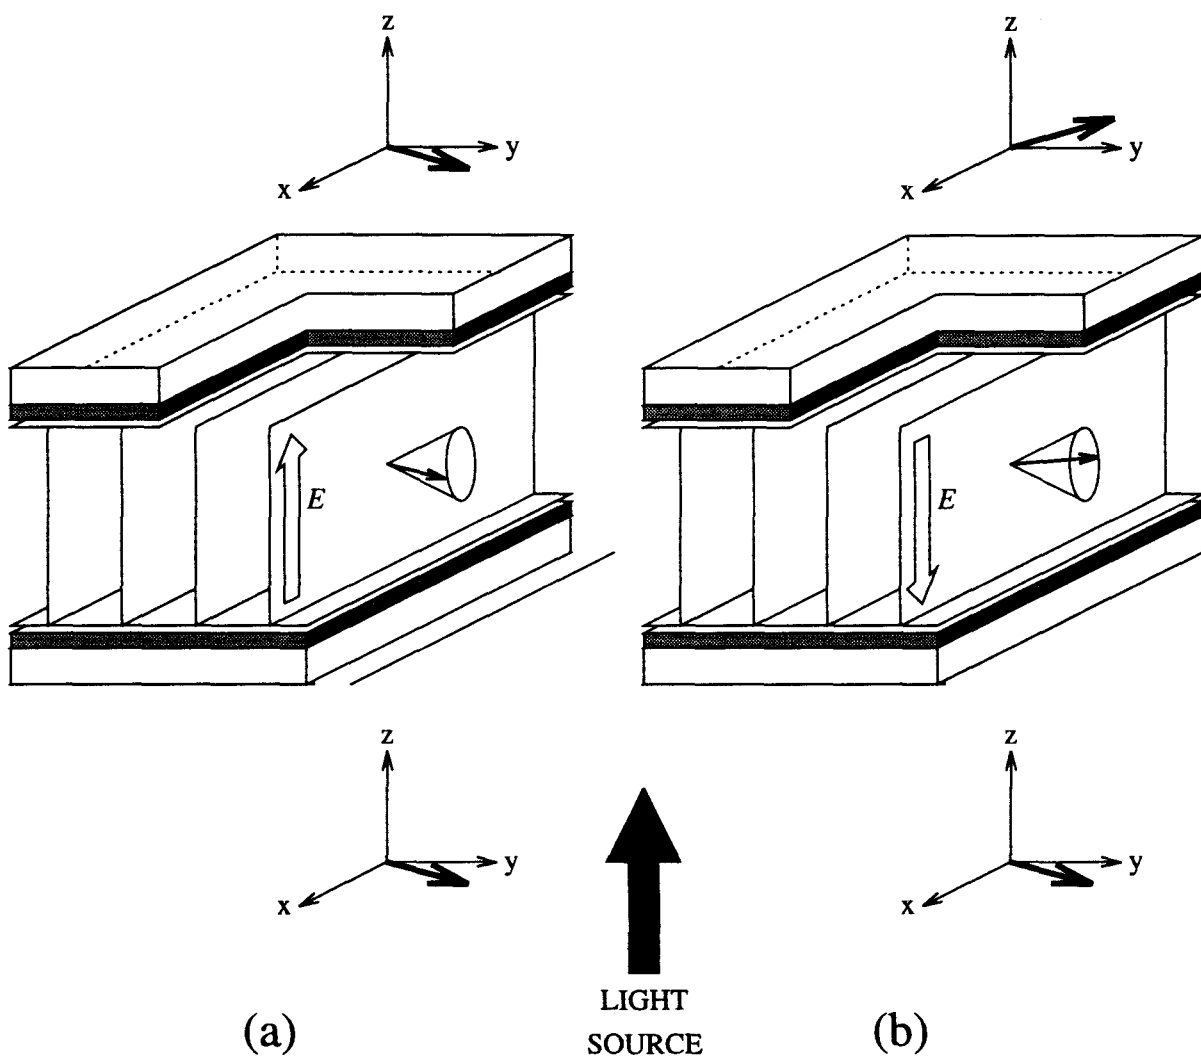
Figure 2–4: Ferroelectric liquid crystal operation. (a) FLC optic axis switched so that it is parallel to incident polarised light - no change in polarisation. (b) FLC optic axis at 45° to incident light - polarisation rotated by 90°.

Real devices – the chevron defect

The most common defect found in SSFLC S_C^* devices is the chevron defect [85, 54]. They are caused by the shrinking of the smectic layers during the final stages of cell fabrication, so that the layers are not normal to the bounding glass plates; zig-zag patterns can be observed when viewed through a polarising microscope. With devices in the chevron structure, the switched states are not fully bistable, and one of the states is usually preferred when the driving voltage is removed. Contrast ratios are also severely reduced. The chevron geometry can be changed to the more desirable bookshelf geometry by applying a low frequency alternating electric field across the FLC layer, typically 30 V/μm or more.

Switching time

To a first approximation the switching time τ can be given by,

$$\tau = \frac{\eta}{P_S E} \tag{2.2}$$

where, η and P_S are the FLC material's viscosity and spontaneous polarisation, and E is the electric field applied across it. Clearly τ is lower for low viscosities, high polarisation, and high electric fields.

Switching energy

In general, the power dissipated, P, in charging and discharging a capacitor from 0 to V and back to 0 again at frequency f is given by,

$$P = QVf \tag{2.3}$$

where Q is charge dumped onto the capacitor. For an area A of FLC material with spontaneous polarisation P_s, the charge Q associated with switching from one state to another is given by,

$$Q = 2|P_s|A \tag{2.4}$$

so,

$$P_{FLC} = 2\,|P_S|\,AV_{FLC}f \qquad (2.5)$$

A rough calculation shows that P_{FLC} for a typical test cell filled with SCE13 FLC, is about 1 mW when the following values are used:

$$
\begin{aligned}
P_S &= 26\ \mathrm{nC/cm^2} & V_{FLC} &= 10\ \mathrm{V} \\
A &= 1\ \mathrm{cm^2} & f &= 2\ \mathrm{kHz.}
\end{aligned}
$$

Charge balancing

In order to prevent chemical degradation in any liquid crystal device, care must be taken to ensure that over a period of time, the LC layer does not receive a net voltage or field across it. Practically, this means that if a positive voltage pulse is applied across the LC for a certain period of time, a negative pulse of equal duration should then be applied across it. The is usually referred to as the charge balancing requirement. The FLC layer in SSFLC devices responds to the *sign* of the electric field applied across the layer. Therefore because of charge balancing, the duty cycle of the device can be significantly affected as the FLC will be in its OFF state for 50% of the time. This is explained in more detail for specific device implementations in Section 2.4.

2.3.2   Other FLC modes

SSFLC devices are inherently binary in nature, however, some FLC modes can have an analogue response.

Electroclinic effect

The smectic A (S_A) phase, used in the electroclinic or soft mode electrooptic effect [4,16], has attracted attention because of its analogue optical response and sub-microsecond switching times. In this phase, the directors are perpendicular

to the layers and the optic axis can be rotated in proportion to the applied field. Davey and Crossland [16] have built test cells capable of generating more than 50 grey levels. However, high electric fields are required, but, the main limitation of using the electroclinic effect to generate reproducible grey levels, is that the cell's temperature must be controlled to within a few hundredths of a degree. This could be a problem for addressing an array of pixels on a silicon backplane as there are likely to be variations in temperature across the array, because of power dissipated in switching transistor circuits.

Distorted helix effect

The distorted helix effect (DHE) also has a fast, analogue response [21]. The DHE device structure is similar to the SSFLC but the helical precession of the director from layer to layer is not suppressed. This can be achieved by using FLC mixtures that possess a short helix pitch. The application of an external field distorts the helix and thus changes the cell's birefringence. However, at present DHE device research is at an early stage of development and cells tend to suffer from alignment and scattering problems.

2.3.3   Addressing schemes

Liquid crystal devices can be addressed optically [32,71] or electrically. Electrical addressing of array LC devices can be performed by direct drive, passive, or active matrix addressing methods.

Direct drive addressing

In direct drive addressing, one wire is used to address each pixel. It is useful for arrays up to 16 × 16, but above this, the area required for interconnect would dominate the pixel array. The frame rate is potentially high as all the pixels can be addressed simultaneously, and good contrast ratio can be expected as pixels can be constantly driven.

Passive addressing

The simplest way to address a large array of pixels is to use matrix addressing techniques. Row and column electrodes are patterned onto the glass substrates and the pixels are defined where the electrodes overlap. Data is simultaneously presented to all columns, and one row at a time is selected by an appropriate voltage signal. This scheme results in error voltages which affect the unselected rows [3]. The size of the array is limited by the effects of the error voltages and the characteristics of the (typically nematic) liquid crystal. The frame time is proportional to the number of rows in the display as the signal has to be present on each row long enough for the liquid crystal to switch. This limits device to about 100 lines, which is acceptable for calculators, status display panels, etc.

Multiplexing capabilities are improved with supertwist nematics (STN), which provides higher contrast, sharper turn-on characteristics, and faster response. The fast switching times and bistability of FLCs can also improve the multiplexing capabilities of passively addressed devices: CRL Smectic Technology have developed several transmissive FLC devices, including a 320 × 320 array [98], with a contrast ratio of at least 50:1.

Active matrix addressing

Active matrix addressing overcomes the problems associated with passive addressed nematic devices by including an active element in each pixel. The active element is usually a thin-film transistor (TFT), although two-terminal devices and diodes have also been used [44]. The TFT has the effect of sharpening the effective cell threshold by replacing the unfavourable switching characteristics of the TN cell with the favourable characteristics of a transistor. These TFTs serve to hold the desired voltage at each pixel, while subsequent rows are addressed with significantly reduced cross-talk.

With TFT displays, there is also the opportunity to integrate some of the drive circuitry onto the display, which can drastically reduce the number of external

connectors [24,63]. However, the presence of TFTs on the back glass plate limits both the resolution and the ability to miniaturise the displays. The opaque TFTs take up room and must be a certain physical size to attain the necessary drive. This can cut the display's light efficiency to less than 5%. The main disadvantage of TFT displays is their high cost. This is a direct consequence of the need to produce displays with a working TFT at each pixel. Defect densities of 0.01 defects/cm^2 lead to yields of only 10–20% for a 10 inch diagonal display. The defect density is much lower than that needed for Si memory chips. Fault-tolerant architectures can improve yields [97], but they are still expensive.

Matsushita manufacture a 15 inch diagonal, full colour, display with a resolution of 1152 × 900 pixels [109] is an example of a state-of-the-art active matrix device. This device is probably at the limits of what is economically feasible.

2.4 Ferroelectric liquid crystal over silicon

In recent years there has been an active interest in developing the hybrid technology of ferroelectric liquid crystal over silicon (FLCOS). The fast switching times and low switching energies of FLCs, coupled with the high addressing speeds of Si backplanes make FLCOS SLMs attractive components for compact optical systems. The Si and LC industries have each progressed rapidly over the last decade so that the performance capabilities of SLMs has increased and can be expected to continue as a consequence of improvements in the component technologies.

During the 1980's, complementary metal oxide semiconductor (CMOS) became the dominant technology for general-purpose integrated circuit (IC) applications. It has come to the forefront primarily due to low power consumption; aside from leakage currents, power is dissipated only during switching events and not when the circuit is in a stable (non-switching) state. This is a major improvement over what occurs in circuits based on bipolar and nMOS fabrication technologies. Similarly, circuits can be designed where small deviations in transistor characteristics do not perturb the operating point. Consequently, it has

proven to be a robust manufacturing technique that permits large quantities and varieties of ICs to be fabricated with high yield. The technology is very flexible – it is suitable for implementing digital logic, memory devices, photodetectors, signal processing, and various analogue applications.

An FLCOS SLM is constructed by sandwiching a thin layer of FLC between a silicon backplane IC and a block of glass coated with a transparent front electrode (see Fig. 2–5). The backplane IC normally comprises an array of pixels, each with a memory circuit and an aluminium pad that doubles as an electrode and a mirror. Peripheral circuitry, such as shift registers and/or decoders, is usually included to address the array. The transparent front electrode is common to all pixels so the voltage applied across the FLC at each pixel is the difference between the front electrode voltage and that applied to the pixel electrode mirror. The device works in a reflective rather than transmissive mode but the principle of selectively rotating the polarisation of the read-beam is exactly the same (see Section 5.4).
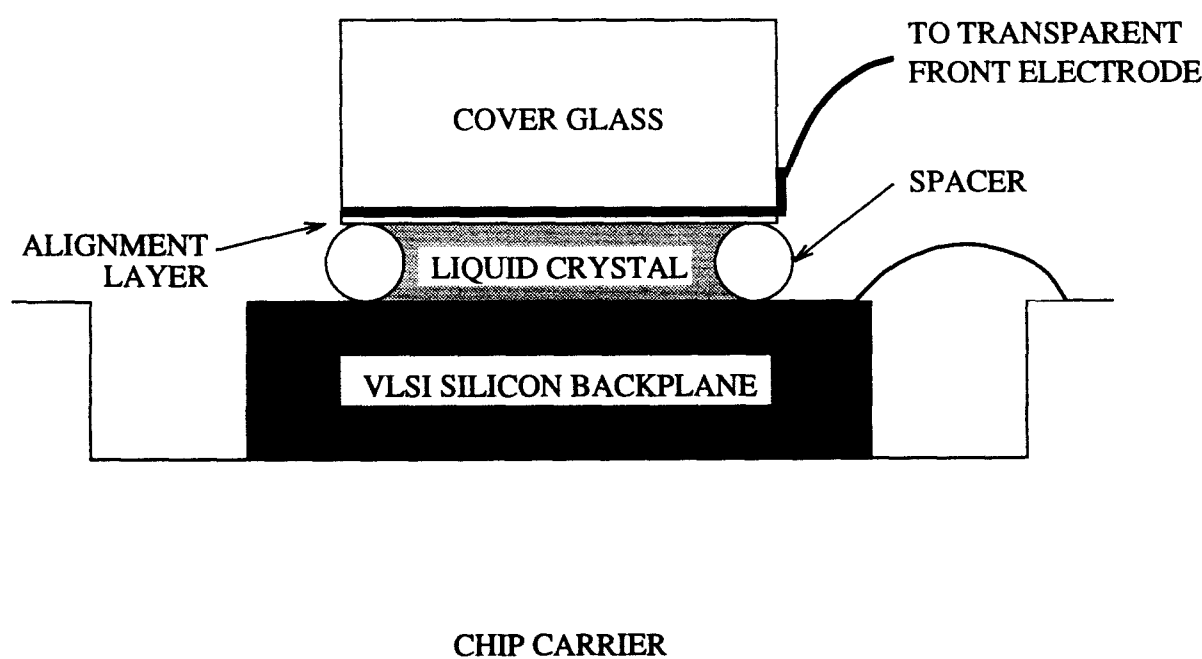
Figure 2–5: Cross-section of a FLCOS SLM.

The remainder of this chapter reviews the various pixel designs that have been used in FLCOS SLM backplanes. These range from simple single transistor pixels to so-called 'smart pixels' that incorporate optical input with some primitive local

processing. I will discuss the single transistor pixel in some detail because the SBS256 device described in Chapters 4 and 5 was developed to overcome the short-comings inherent in the single transistor design.

2.4.1 DRAM-type devices

The main thrust of research into FLCOS SLMs has focused on new device designs with increased bandwidth. As with active-matrix displays, most designs utilise a pixel circuit based on the one-transistor cell of the dynamic random access memory (DRAM) to achieve a high pixel density (see Table 2–1).

One-transistor pixel

The basic structure of the one-transistor pixel is shown in Fig. 2–6. The capacitor C_{FLC} represents the capacitance of the overlying FLC layer, while C_{MIRROR} is the storage capacitance associated with the electrode mirror. The transistor acts as a voltage controlled switch, that isolates the capacitors when open (ENABLE low), and permits the capacitors to charge to the DATA line potential when closed (ENABLE high).

The simplicity of the single transistor design offers small pixel size and thus high pixel density. However, it also has some important limitations:

- Can only provide a limited amount of charge to switch the overlying FLC.

- Suffers from photo-induced charge leakage.

- Limited FLC charge balancing capability – requires a pattern/inverse-pattern drive scheme.

These will now be considered in turn.

To achieve short frame scan times, it is usual for the row address time (the time when each ENABLE is high) to be shorter than the FLC response time.

Array Size	Backplane Technology	Pixel Area (in μm^2) (Pitch in μm)	Frame Rate (kHz)	Ref. (Year)
240 × 240*	CMOS	145 × 200 (165 × 220)	0.06	[114] (1982)
176 × 176	3 μm p-well CMOS	22 × 14.5 (30)	1	[106] (1991)
12 × 12	CMOS	70 × 100 (110)	. . .	[5] (1992)
128 × 128	2 μm n-well CMOS	22 × 22 (30)	5**	[49] (1993)
256 × 256	2 μm p-well CMOS	17 × 17 (20)	4**	[36,35] (1993)
256 × 256	1.2 μm p-well CMOS	28.4 × 28.4 (30)	4**	[36,35] (1993)
256 × 256	1.2 μm n-well CMOS	17 × 17 (21.6)	8.3**	[66,67] (1993)
512 × 512	3.0 μm p-well CMOS·	22 × 14.5 (30)	0.4	[10] (1994)

* This device has a nematic LC modulator layer.

** These frame rates do not include charge balancing (divide by two).

Table 2–1: Examples from the literature of the development of one-transistor DRAM-type devices.
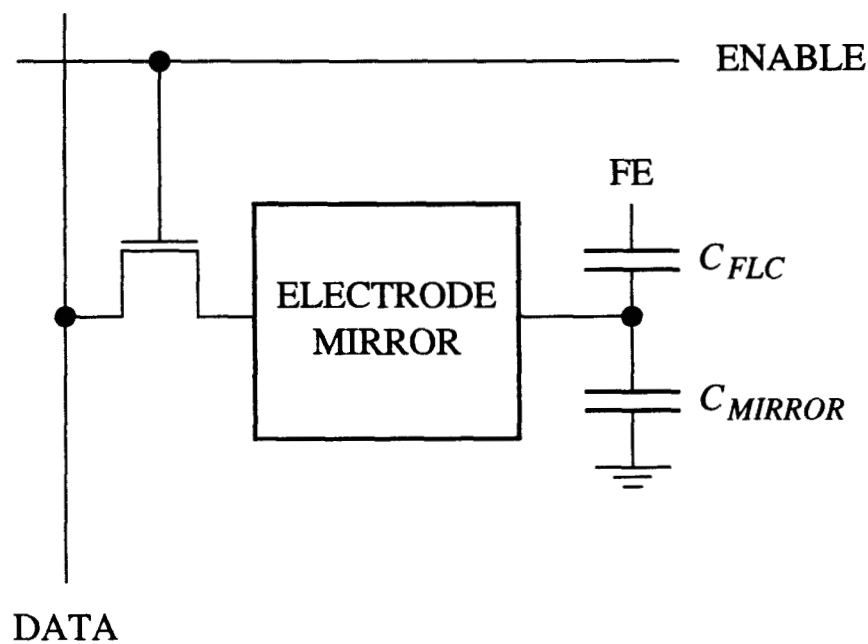
Figure 2–6: Schematic of the one-transistor DRAM-type pixel. It comprises a FET pass transistor addressed by ENABLE and DATA buslines, a charge storage capacitor, and an electrode mirror. FE denotes the transparent front electrode.

Therefore the pixel capacitors are charged and then left electrically isolated until they are readdressed in the next frame. The movement of the ferroelectric dipoles while switching the FLC results in a voltage drop ΔV across the FLC layer given by,

$$\Delta V = \frac{2 P_S A}{C} \tag{2.6}$$

where A is the pixel area, and C is the effective capacitance of the pixel. From Equation 2.2, shorter switching times can be achieved with higher P_S FLC materials but ΔV would be also be larger (assuming A and C were not changed). If ΔV approaches the voltage impressed by the DATA signal, only partial switching will occur. Each row of the array would then have to be addressed for a time approaching τ so that more charge could flow onto the pixel as the FLC switches; the frame scan time would therefore be drastically increased. This may limit the FLCs used on one-transistor DRAM devices to those that possess a P_S of less than a few tens of nC/cm^2 [15].

Photo-induced charge leakage is a problem inherent with DRAM-type pixels [35]. Once the pixel capacitor has been charged and isolated, the charge can leak away

through the transistor channel resistance R_{OFF} (10^{12} Ω), the liquid crystal resistance R_{FLC} ($10^{11\pm1}$ Ω), and the drain-substrate reverse-biased junction (D_{DB}) of the pass transistor (see Fig. 2–7). When light falls onto the backplane sub-
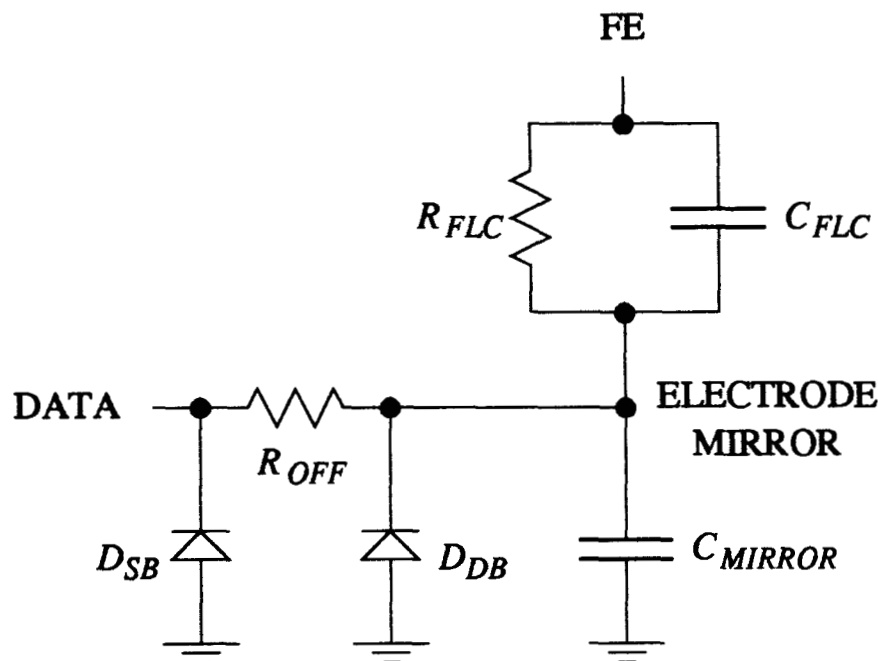
Figure 2–7: Equivalent circuit of the one transistor DRAM-type pixel.

strate, the most significant leakage path is through D_{DB}. Light falling onto the substrate generates electron-hole pairs as it is absorbed near the surface. For a p-type substrate, the additional holes do not really affect the total majority hole population, however, the minority electron population is greatly increased. The electrons can exist in the substrate for 10–100 μs, where they can travel for 100's of microns before they recombine with a hole. If some pass close to the reverse-biased diode they may be swept across its built-in electric field and so discharge the pixel capacitor. The photo-induced minority carrier population and thus the discharge rate is proportional to the light intensity. Handschy *et al* [35] measured a 50% erasure time of 100 ms with a light intensity of 10 mW/cm^2 for their second-generation 256 $\times$ 256 pixel device. To reduce the effects of charge leakage, the following steps can be taken:

- Make the pixel storage capacitor as large as possible.

- Shield as much of the substrate as possible from incident light.

- Refresh the pixel array at high frame rates.

- Use a low intensity read-beam.

- Pulse the read-beam on a low duty cycle.

It should be noted that if the FLC in FLCOS devices could be aligned into the fully bistable bookstack geometry, then charge leakage would only be significant if the capacitor discharged before the FLC had switched. However, the bookstack geometry is usually attained by treating the FLC layer with a high electric field (20–30 V/μm) – at the time of writing, this treatment has not been reported for a DRAM FLCOS device, probably because researchers do not want to risk damaging one of their limited supply of working devices.

The usual method of charge balancing the FLC layer on a one-transistor DRAM device, requires that after a pattern has been scanned onto the array and viewed, the inverse of the pattern must then be scanned on and held for the same duration. The pattern on the SLM is therefore discontinuous – it can only be interrogated for 50% of the time. Again, if the FLC layer had the bistable bookstack structure, the array would only need to be addressed when the pattern was to be altered, so a suitable drive scheme could be developed where the valid duty cycle was almost 100% (at the expense of frame rate).

Three-transistor pixel

The three-transistor DRAM pixel design shown in Fig. 2–8, permits continuous interrogation of the overlying FLC layer whilst ensuring it is fully charge balanced.

The drive scheme for this type of pixel requires the following:

- The FLC layer must be fully bistable.

- The DATA line can be set to 0, $V_{DD}/2$, and V_{DD} voltage levels.

- The front electrode must be held at $V_{DD}/2$.
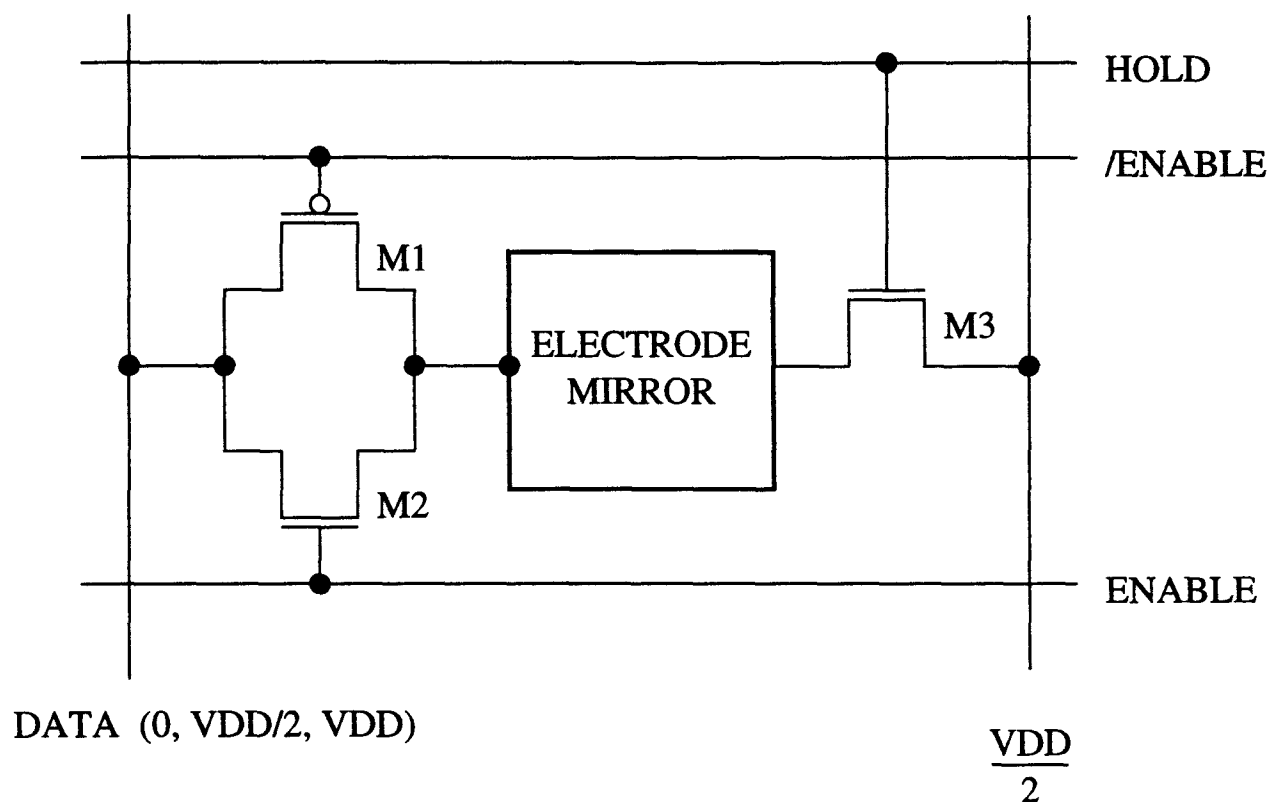
Figure 2–8: Schematic of the three-transistor DRAM-type pixel capable of providing almost continuously valid patterns. The two power rails, VDD and GND, which are required for substrate and well taps, are not shown.

Under normal conditions the pass transistor M3 is on so the electrode mirror is also held at $V_{DD}/2$; as there is no voltage across the FLC layer, it remains in the state it was last driven to. Now assume the FLC is in its off state. It is switched to its on state by turning off the pass transistor, applying 0 V to the DATA line, and turning on the transmission gate M1 and M2. This results in a positive voltage (with respect to the mirror electrode) of $V_{DD}/2$ across the FLC layer. After the FLC has switched, the transmission gate is turned off and the pass transistor is turned on again. To switch the FLC off again, the sequence is repeated with the DATA line set to V_{DD}, to generate a negative voltage across the FLC layer.

As row addressing is normally used to address pixel arrays, the DATA lines can be set to a third voltage level, $V_{DD}/2$ to permit individual pixels to be changed by setting the appropriate DATA line to 0 or V_{DD} while setting all the others to $V_{DD}/2$. A 64 × 64 device using this type of pixel has been designed by Mike Snook in the Applied Optics Group at the University of Edinburgh [94]. It is

currently being fabricated using a high voltage (35 V) CMOS process that will permit both microsecond FLC switching times and high field treatment of the FLC layer to ensure bistability. It will be used in an efficient optical cross-bar routing network [72] (see Section 3.2.1).

Four-transistor pixel

The four-transistor DRAM pixel implemented by Jared *et al* [46] is illustrated in Fig. 2–9. The inverter decouples the storage capacitor from the the pixel electrode mirror, and thus overcomes the P_S limitation of the single transistor design. However, the pixel still stores the state dynamically, and so suffers from light-induced charge leakage.
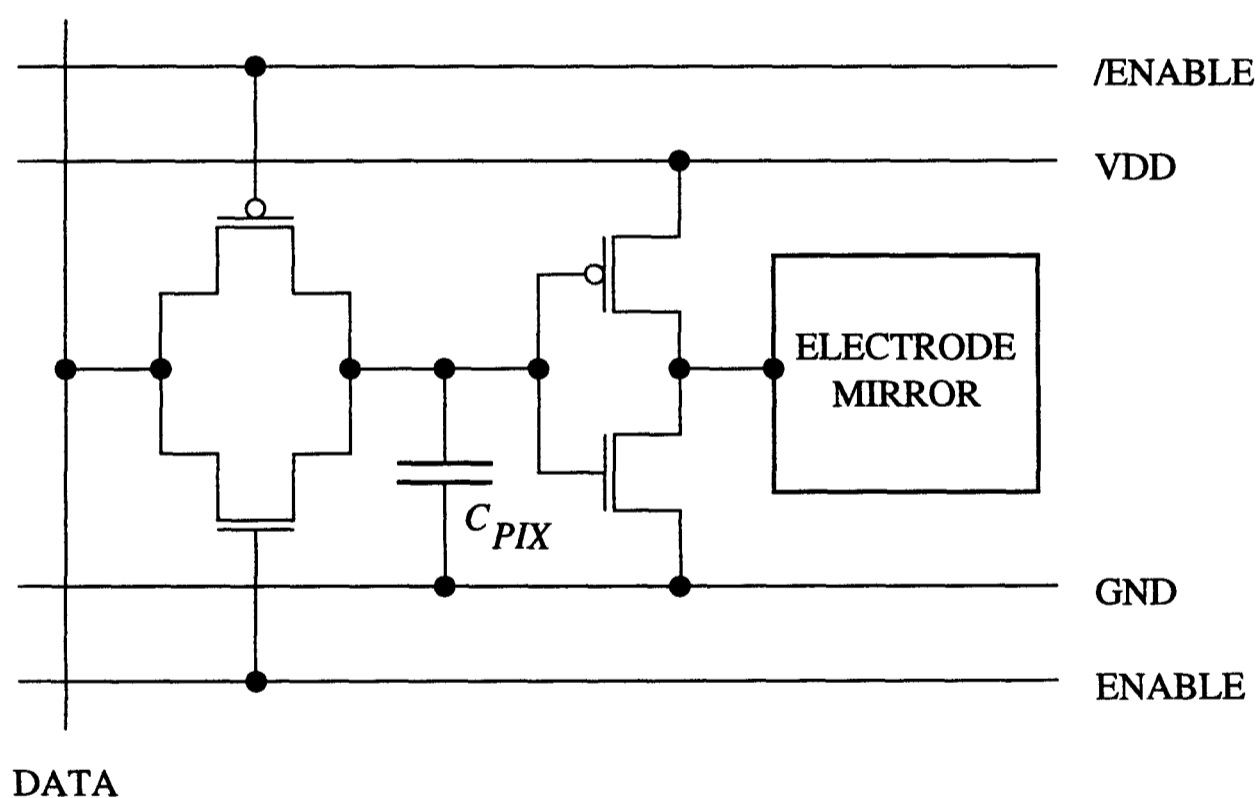
Figure 2–9: Schematic of the four-transistor DRAM-type pixel implemented by Jared *et al* [46].

2.4.2 SRAM-type devices

The one-transistor DRAM-type pixel offers high density at the expense of limited storage and liquid crystal drive capability. For a given process, a typical

one-transistor DRAM-type pixel has a pitch 10–20 times the length of the minimum size transistor, whereas a SRAM-type pixel may be 30–40 times the minimum transistor length. The SRAM-XNOR pixel was originally designed by I. Underwood [103] to provide a stable drive waveform for nematic liquid crystal modulator materials. Here stable implies there is no charge leakage voltage 'droop', a problem associated with the one-transistor DRAM pixel. My preliminary investigations [105] with the Mk II 16 × 16 test-bed device showed that the SRAM-XNOR pixel could also be used to drive FLC modulators (see Fig. 2–10).
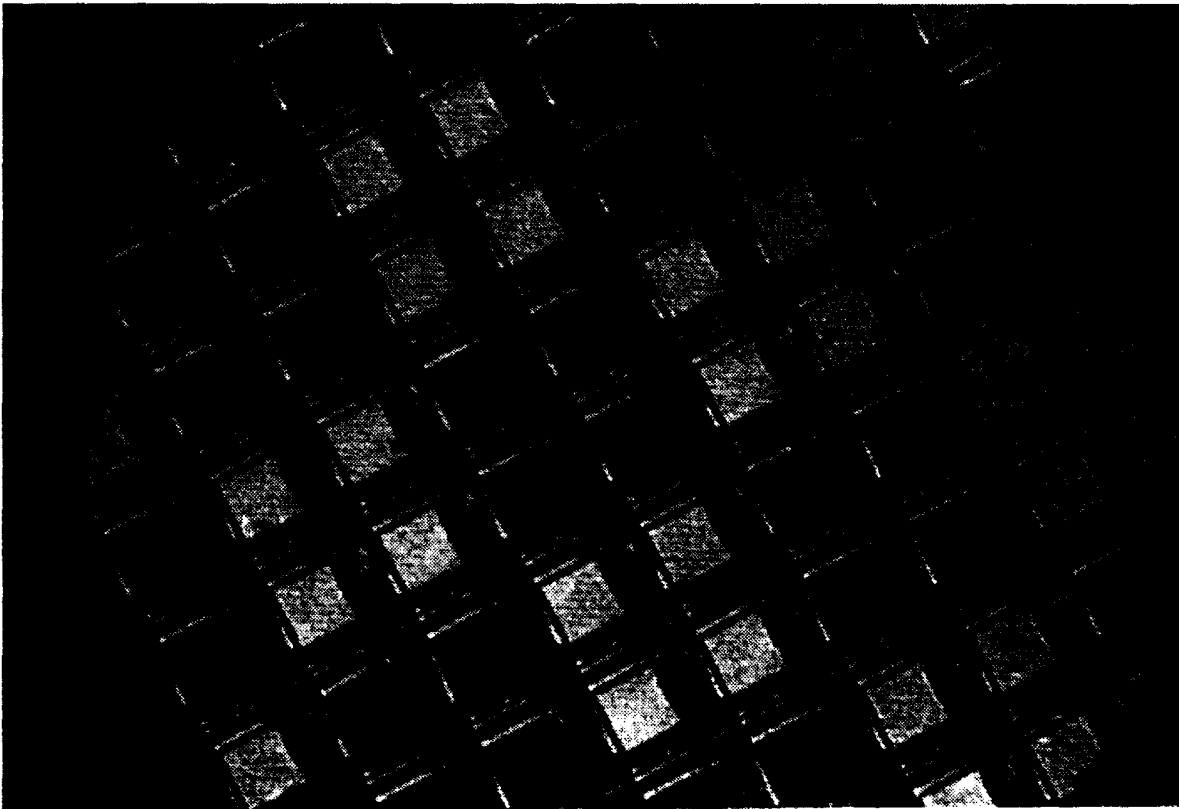
Figure 2–10: Test pattern on the 16 × 16 SRAM-XNOR device from my preliminary investigations into using FLC modulators.

Although requiring more area, the SRAM-XNOR pixel shown in Fig. 2–11, has significant performance advantages over the DRAM-type. These include:

- Robust data storage under high read-beam intensities. By storing the data value in a static latch, the pixel does not have to be continually refreshed because of the latch's internal positive feedback loop.

- Easy FLC charge balancing to prevent chemical degradation. The inclusion of the XNOR gates permits the FLC layer to be charge balanced simply by toggling the global clock signal CK on a 50% duty cycle.

- Synchronous FLC switching across the whole array (Section 4.6 describes a suitable drive scheme.)

- The ability to switch high spontaneous polarisation FLC mixtures. As in the case of the inverter in the four-transistor DRAM pixel, the XNOR gate also continually drives the electrode mirror via the power rails, resulting in an effectively unlimited amount of charge being available to switch the FLC. This can permit high P_S materials to be used.

- An important side effect of the robust data storage and the continuous drive capability of the XNOR gate is that although the bistable SSFLC bookshelf geometry is desirable (for its high contrast), it is not required.
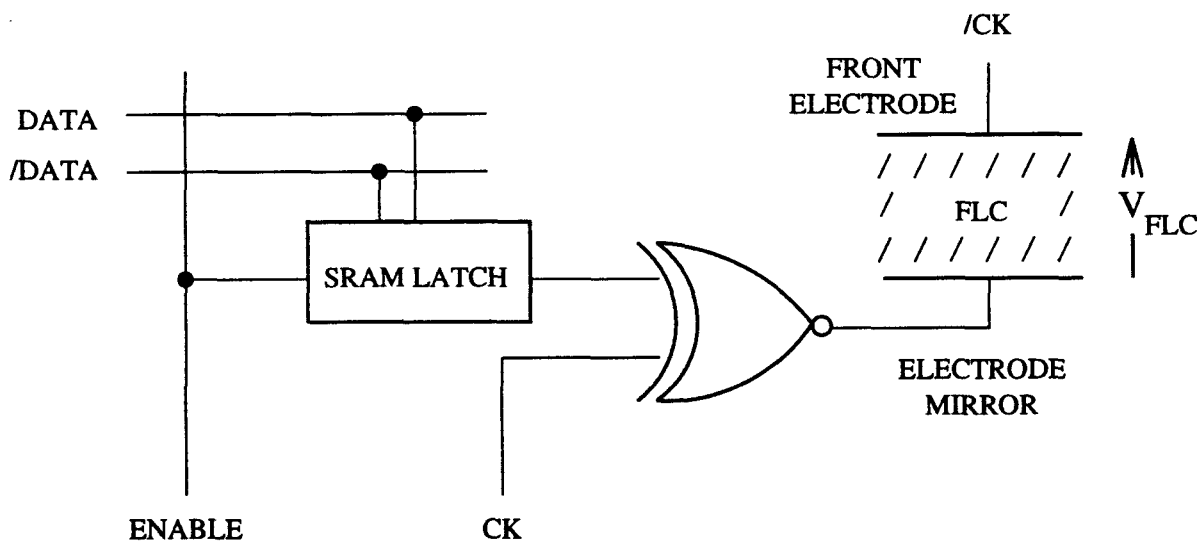
Figure 2–11: Schematic of the SRAM-XNOR pixel used in the 16 × 16 device [103].

Table 2–2 shows examples from the literature of the development of SRAM-type devices. McKnight [65] developed a 50 × 50 device, using nMOS technology, to drive a nematic liquid crystal working in the hybrid field effect mode. The 64 × 64 device by Jared *et al* [46] uses an inverter rather than an XNOR/XOR

Array Size	Backplane Technology	Pixel Circuit	Liquid Crystal	Frame Rate	Ref. (Year)
16 × 16	6 μm nMOS	SRAM-XNOR	GHLC	5 Hz	[103](1987)
50 × 50	1.5 μm nMOS	SRAM-XNOR	HFE E7	60 Hz	[65](1989)
64 × 64	3 μm CMOS n-well	SRAM-Inv.	ZLI-4003	4.5 kHz	[46](1991)
128 × 128[a]	1.5 μm nMOS	SRAM-XNOR	...	...	[65](1989)
256 × 256	1.2 μm CMOS n-well	SRAM-XOR	SCE13	2 kHz	[10] [b] (1994)

[a]Probe testing revealed a fatal design flaw on the backplane.

[b]Chapters 4 and 5 describe the development of this device in detail.

Table 2–2: Examples from the literature of the development of SRAM-type devices.

function so charge balancing must be accomplished using a pattern/inverse pattern drive scheme. A derivative of the SRAM-XNOR pixel has been used in the 256 × 256 pixel FLCOS SLM (the SBS256 device), and is described in detail in Chapters 4 and 5.

2.4.3 Smart pixel devices

As SLM technology has progressed, designers have tended to either develop higher resolution arrays or incorporate more functionality at each pixel. So-called 'smart pixels' [5] with photodetectors for optical input, can perform primitive thresholding and processing operations on incoming signals. Many smart pixel devices have been developed (see the review paper by Johnson *et al* [49]) – here I consider just two smart pixel designs to illustrate the level of functionality available in current devices.

[5]I can only assume if anyone designs a pixel incorporating something like a 4-bit microcontroller, it will be referred to as a 'Mensa' pixel!

Isophote pixel

The SASLMIII [6] isophote chip designed by Snook *et al* [95] performs a variable threshold window edge enhancement function on an input image. It has a 64 × 64 pixel array, where each pixel determines whether it lies on an edge of the input image by computing the following nearest neighbour logic function:

$$Y = ((A \oplus I) + (B \oplus H) + (C \oplus G) + (D \oplus F)) \bullet E, \qquad (2.7)$$

where $A - I$ correspond to the pixel positions illustrated in Fig. 2–12, and $\oplus$ is the XOR function. The pixel's own state E is included to ensure that only one line is activated along an input image edge. Each pixel contains 92 transistors, making it one of the most complex designs yet implemented in a large pixel array. The device must be operated using non-overlapping, and pulsed input and read beams. Its frame rate is limited only by the FLC switching time, and the usual charge balancing requirements.

A	B	C
D	E	F
G	H	I
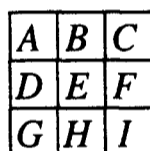
Figure 2–12: Nearest neighbour designations for the Isophote pixel.

Winner-take-all pixel

As its name implies, the winner-take-all (WTA) device [55] finds the maximum of a series of (analogue) values. An optical WTA device that can find the largest of an array of optical intensities would be useful for applications such as correlation peak detection, and certain optoelectronic neural network architectures that use competitive learning algorithms.

---

[6]SASLM – Smart and Advanced Spatial Light Modulators LINK programme.

Slagle and Wagner [93] have built a number of WTA devices including a 32 × 32 array [93] and a 9 × 64 array broken into 31 competitive patches [110]. In each unit of the competitive patch device (see Fig. 2–13), a phototransistor detects the incoming light and converts it to a current. This is copied by the current mirror and fed into the two transistor competitive circuit [55]. The unit that generates the most photocurrent wins the WTA competition to supply current to the global inhibition bus. The inhibition bus effectively turns off the supply from all the other units. The current flowing in the winning unit is mirrored onto the electrode mirror to turn the overlying FLC to its on state.
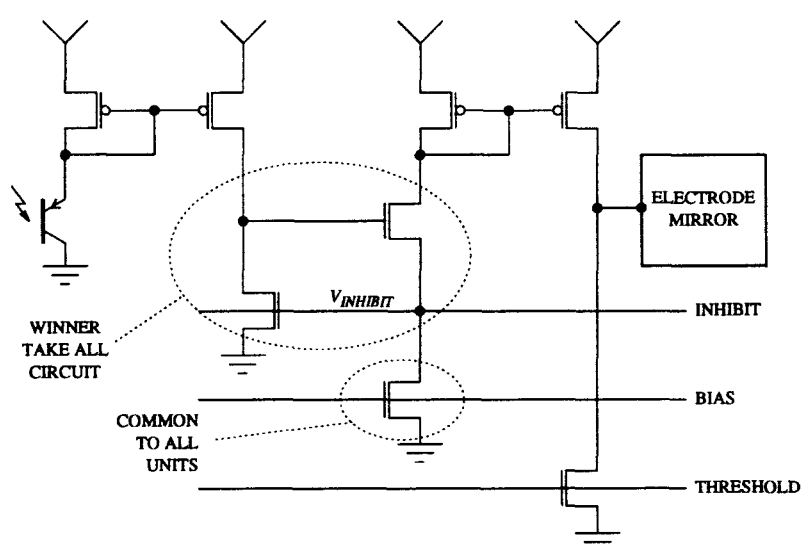
Figure 2–13: Schematic of the winner take all (WTA) pixel implemented by Wagner and Slagle [110]. The winning unit supplies current to the global inhibition bus and increases $V_{INHIBIT}$, which, by positive feedback, ensures that no other unit supplies current to the bus.

As an alternative to generating an optical output, Turner and Johnson [102] have developed an electronic output WTA array to sequentially locate peaks at the output plane of an optical correlator. The winning pixel on the 12 × 12 array reports its position by a combination of digital (course) and analogue (fine) signals, and is then disabled so that the next peak can be found. Once all the peaks have been reported, a global reset prepares the array for the next correlation frame.

Technology	Array Size	Pixel Pitch (μ m)	Frame Rate (Analogue/ Binary)	Contrast Ratio	Operating Voltage (V)	Reference
Liquid crystal devices						
AMLCD	1024 × 1440	50	50 Hz (A)	200:1	15	[63]
FLCD	320 × 320	80	200 Hz (B)	50:1	60	[98]
FLCOS-DRAM	256 × 256	20	2 kHz (B)	29:1	5	[35]
FLCOS-SRAM	256 × 256	40	2 kHz (B)	8:1	5	[10]
FLCOS-Smart	64 × 64*	160	2 kHz (B)	8:1	5	[94]
Self electrooptic effect devices						
S-SEED	256 × 128	...	...(B)	2:1	10	[11]
Si-SEED	16 × 16*	120	70 MHz (B)	2:1	10	[30]
FET-SEED	6 × 6	80	2 GHz (B)	2:1	10	[56]
Others						
DMD	768 × 576	17	1.5 kHz (B)	50:1	5	[69]
Si-PLZT	8 × 8*	250	10 MHz (A)	>20:1	>20	[48]

* Devices currently being fabricated – expected performance values are given.

Table 2–3: Typical characteristics of the various SLM technologies described in this chapter.

WTA implementations tend to suffer from the device parameter variations inherent in all CMOS fabrication technologies which can result in certain units being more likely to win the competition. Another problem is that the circuit will still attempt to generate a winner even when there is no input signals.

2.5 Summary

This chapter has given an overview of some SLM technologies, including digital micromirror devices (DMDs), Si-PLZT devices, SEEDs, LCDs, and ferroelectric liquid crystal over silicon (FLCOS) SLM devices. Table 2-3 provides a brief summary of the typical characteristics of each of the technologies.

Texas Instrument's proprietary DMD technology is very promising, with much of their effort being directed at trying to make it a viable technology for high-definition projection display systems. Both Si-PLZT and SEED technologies offer very fast switching times, and are suitable for smart pixel and optical computing architectures. However, they will probably be limited to systems with 10's to 100's of channels because of, (1) power dissipation, and (2) significant area requirements of high frequency drive transistors. FLCOS SLMs are attractive because of their small size, low power, high frame rates, versatility, and because they are based on two well established technologies, namely integrated circuit technology and liquid crystal displays (LCDs). In the next chapter, I review some applications for which FLCOS SLMs are particularly well suited.

Chapter 3

FLCOS SLM applications

I begin this chapter with a brief discussion of why FLCOS SLMs are attractive components for optoelectronic systems, and then I review some applications that have already benefited, or could benefit, from the use of FLCOS SLMs. Some of these only require binary (phase or amplitude) modulation, but for others, grey-scale modulation can greatly enhance their capabilities, so some methods for generating grey-scale are also considered. The remainder of the chapter focuses on the emerging field of artificial neural networks and the suitability of FLCOS SLMs for their optoelectronic implementation.

3.1   Why use FLCOS SLMs ?

The technology of FLCOS SLMs has advanced rapidly in the past decade or so. Nevertheless, researchers in the field have tended to concentrate their efforts on designing, fabricating, and characterising new devices with higher pixel count and performance rather than actually using them. As FLCOS SLMs become available to other researchers 'off-the-shelf', the balance of effort should shift over to applications-led research.

The main strengths of FLCOS SLMs (see Table 3–1) can be attributed to combining two well-matched and rapidly advancing technologies – CMOS VLSI

Backplane characteristics	Optical characteristics
Mature, advancing technology	Small – compact, cheap optical systems
Versatile functionality	Illumination – coherent or incoherent
Fast addressing	Fast switching times and frame rates
Low power	High contrast
Low cost in bulk quantities	Grey scale – temporal/spatial multiplexing
Post-processing available	

Table 3–1: Strengths of FLCOS SLM technology.

electronics and the FLC modulating materials being developed for the flat-panel display industry. The low voltage requirements and fast switching times of FLC materials combine well with CMOS VLSI technology which offers versatile functionality with high density integration, low power dissipation, and potentially low cost for volume production. The resolution, frame rate, and performance of FLCOS SLMs are likely to improve purely as a consequence of advances made in the component technologies, but also because of progress in SLM-specific research, such as more efficient circuit and layout design, post-processing, and device assembly. For example, by adding of an extra layer of metal after post-processing planarisation procedures [79], the optical quality and fill-factor of a device can be significantly improved. This can also lead to a reduction in the area set aside for the mirror electrodes in each pixel, and so reduces the pixel size, which in turn, can decrease the device size or allow more pixels to be incorporated onto a device.

3.2 Binary mode applications

Most electrically-addressed FLCOS SLMs work in a binary mode; the Si backplanes are usually composed of digital (binary) circuits, and the FLC layers are normally configured in the binary SSFLC geometry. This binary mode operation makes them suitable for many applications.

3.2.1   Interconnect/crossbar switching

Some optical systems require reconfigurable optical-space switches to connect two arrays of optoelectronic devices.

Vector-matrix crossbar switch

A commonly used optical-space switch architecture is the vector-matrix multiplier crossbar originally developed by Goodman *et al* in 1978 [29] (see Fig. 3–1). A number of systems using this architecture have been described [17,50]. However, the cylindrical lenses required for fan-out and fan-in tend to smear the light non-uniformly which limits potential scalability.

| INPUT
VECTOR | FAN-OUT
OPTICS | INTERCONNECT
SHUTTERS
(SLM) | FAN-IN
OPTICS | OUTPUT
VECTOR |
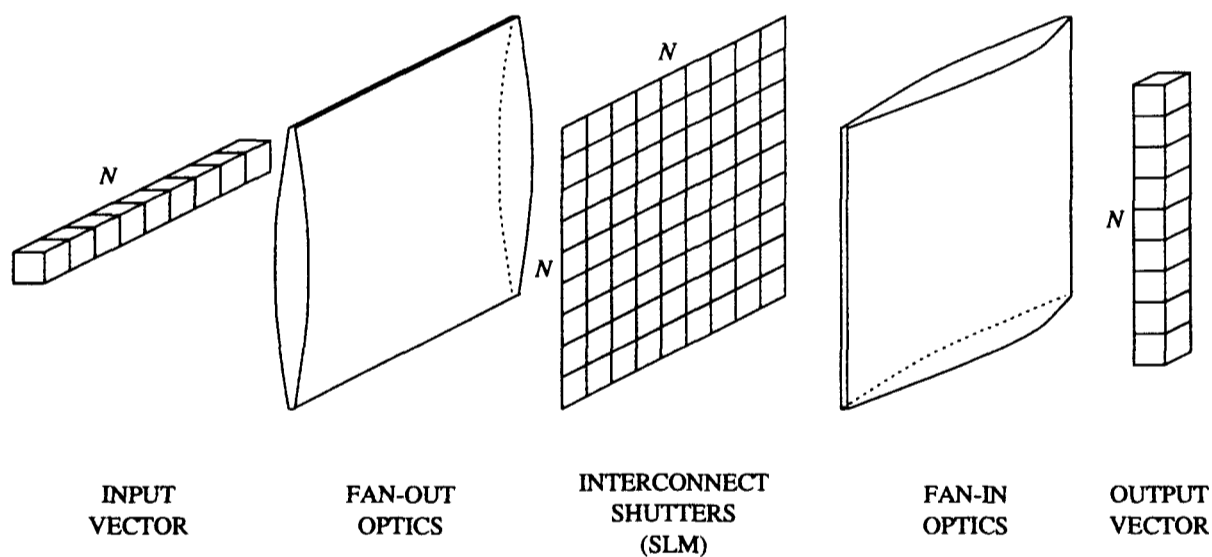
Figure 3–1: The vector-matrix multiplier crossbar switch architecture. Each optical input channel is imaged onto a column of the SLM and the transmitted light from a row is focused onto an output channel.

Matrix-matrix crossbar switch

The scalability of the crossbar switch can be greatly improved by redesigning it as a two-dimensional matrix-matrix crossbar (MMC) [52,53]. The input and output vectors must be rearranged into two-dimensional arrays. The two $\sqrt{N} \times \sqrt{N}$

matrices are interconnected through an N^2 interconnection matrix. The N^2 matrix is partitioned into sub-matrices so that there is one sub-matrix for each output channel. A copy of the input matrix is multiplied by each sub-matrix and then summed into the appropriate output channel.

| INPUT ARRAY | PASSIVE HOLOGRAPHIC REPLICATION | INTERCONNECT SHUTTERS (SLM) | FAN-IN LENSLET ARRAY | OUTPUT ARRAY |
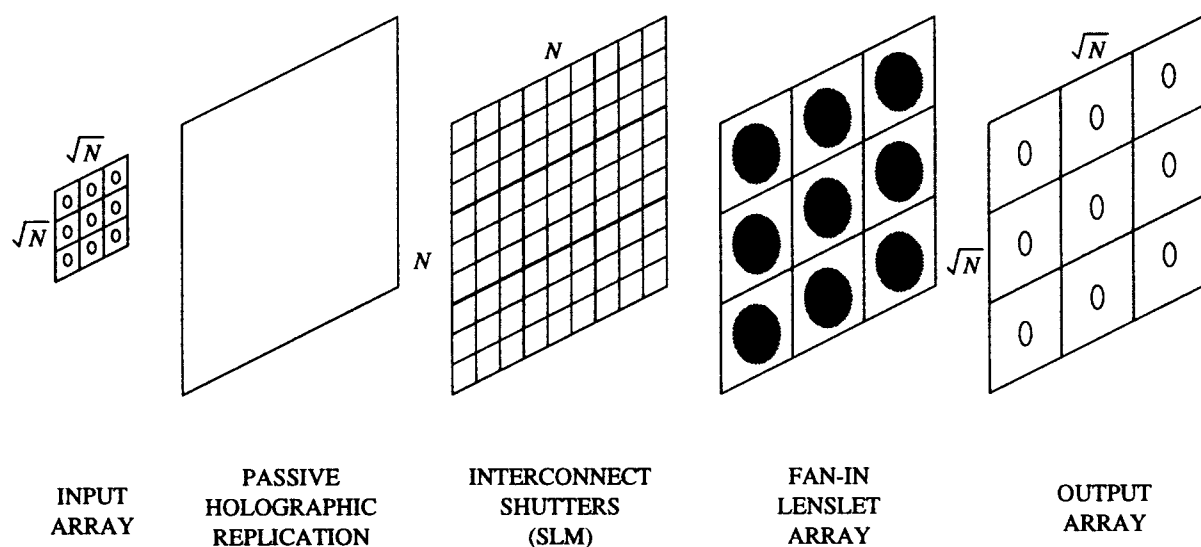
Figure 3-2: The matrix-matrix crossbar (MMC) architecture. The input matrix is replicated using holographic fan-out optics. Each copy is multiplied by an interconnection sub-matrix on the SLM, and then focused onto output channel matrix.

FLCOS SLMs are particularly well suited for implementing crossbar switches because when in the 'on' state, pixels simply act as passive reflectors. Once an optical path has been configured between an input and output channel, the signal bandwidth is only limited by the modulating frequency of the input source. The "Optically Connected Parallel Machines" [72] LINK project plans to build a 64 channel MMC using a FLCOS SLM. It is intended to operate at up to 640 Mbits/sec per channel. The SLM backplane will utilise the three transistor pixel design described in Section 2.4.1 to permit fast reconfiguration times ($\tau_{FLC} \leq 10$ μs) and pixels to be switched on or off individually to open or close a particular path to an output channel.

Holographic crossbar switch

For both VMM and MMC architectures, each of the input channels are each split N ways, creating a loss that increases dramatically as the number of channels increases. The binary phase-only holographic cross-bar switch [75] offers an alternative method for signal routing that does not have the fan-out loss problem. In the binary phase-only holographic cross-bar switch, almost all the available input light from an input channel is directed only to the desired output channel (or channels) by displaying the appropriate routing hologram on the SLM.

O'Brien *et al* [75] implemented a 2×2 crossbar with single-mode optical fibre inputs and outputs using a transmissive THORN EMI 128×128 pixel FLC SLM [8] operating in a binary-phase mode. Transferring this technique to FLCOS SLM technology would permit faster reconfiguration times and smaller, cheaper optical systems to be implemented. The rather limited space bandwidth product of presently available FLCOS SLMs can effectively be increased by combining them with fixed holographic elements and optically addressed SLMs, to permit large fan-out systems [76]. The main limitation of the holographic crossbar switch is the computational overhead involved in calculating the required holographic pattern for each interconnection.

3.2.2   Coherent optical processing

Coherent optical processing systems exploit the Fourier transforming properties of simple lenses [33]. When an object is illuminated with collimated, coherent illumination, and imaged through a lens, the 2-dimensional Fourier transform of the object is formed at the back focal plane of the lens.

Spatial filtering

Just as the Fourier transform of a time-domain electrical signal contains its frequency spectrum, the two-dimensional Fourier transform of an image contains

its spatial frequency spectrum. The Fourier transform can be manipulated to perform image processing functions such as image enhancement and pattern recognition [61]. Fig. 3–3 illustrates the classical 6–f coherent optical processor. Three planes are of interest – the input plane, the Fourier plane, and the output plane. The transform lens forms the Fourier transform of the input image on its back focal plane – the Fourier plane. By including an appropriate filter at the Fourier plane, certain spatial frequencies can be passed or blocked. For example, edge enhancement can be performed by including a filter to block the low spatial frequencies. The retransforming lens then forms the filtered version of the input image at its back focal plane.

A photographic transparency can be used as the Fourier plane filter, but this requires mechanical removal and replacement if a different filter function is needed. Obviously, if an SLM is used, adaptive spatial filtering can be performed. Underwood [103] and McKnight [65] have successfully used the 16 × 16 and 50 × 50 LCOS SLMs respectively, as adaptive spatial filters.
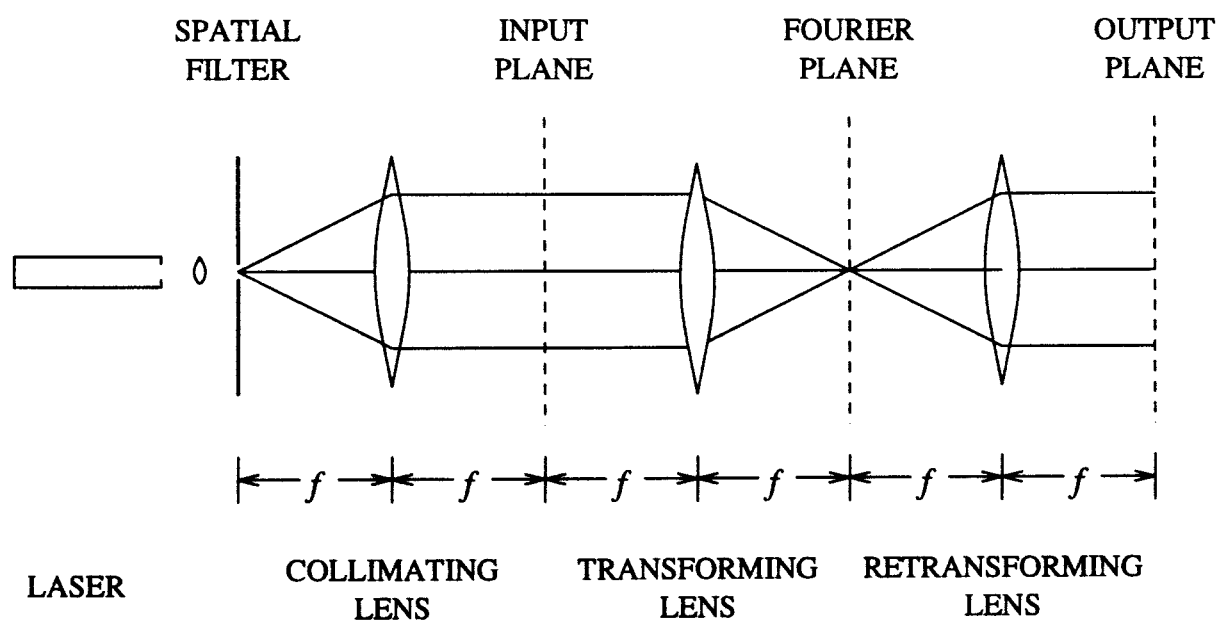
Figure 3–3: The 6–f coherent optical processor. The transform lens forms the Fourier transform of the input image on its back focal plane – the Fourier plane. Certain spatial frequencies can be passed or blocked by including an appropriate filter at the Fourier plane. The filtered version of the input image is then reconstructed by the retransforming lens.

Correlation

A great deal of research has been done in the area of optical correlation for image recognition and classification. Several correlator architectures have been developed. Turner *et al* [101] implemented the classical VanderLugt correlator [60] by using two 64 × 64 FLCOS SLMs [46]. The input plane SLM operates in a binary amplitude mode, while the Fourier plane filter SLM operates in a binary phase-only mode. The input images are Fourier transformed by a lens, and the Fourier filters are computer generated. They identified warpage of the Si backplanes, introduced in the fabrication processes, as a major performance-limiting parameter of their system. McKnight *et al* [68] used their higher resolution 256 × 256 DRAM-type devices in a correlator capable of working at 1 kHz.

In the joint transform correlator (JTC) architecture, both the input and reference images are Fourier transformed simultaneously, and interference between the transforms is achieved in one step. Both the input and reference images can be displayed on the same SLM, thus simplifying optical system alignment at the expense of resolution. The SBS256 SLM described in Chapter 4 is particularly suitable for this type of correlator. Its static memory pixels do not require continual refreshing so the input image can be scanned onto one half of the pixel array, and then a sequence of reference images can be scanned into the other half. Using this drive scheme effectively increases the SBS256 frame rate because only half the array has to be addressed while the reference images are being cycled through. The JTC performance can be improved by including a nonlinear thresholding function at the Fourier plane. The non-linear JTC can exhibit higher correlation peaks and better discrimination [47]. It can be efficiently realised by adding a sharp-thresholding optically addressed FLC SLM at the Fourier plane.

SLM pixellation effects

When a pixellated SLM is used in a coherent optical system, replicated images occur in the output due to the regular structure of the SLM. These replicas arise

through convolution with distinct spectral orders in the SLM's Fourier spectrum and in the case of a correlator can lead to false correlation signals. The effects of replication can be significantly reduced by randomising the pixel positions on the SLM [39]. For FLCOS SLMs, this could be incorporated into the layout of the pixel array during the design phase, but would create many problems for the designer, including signal routing through the array, and placement of the pixels. A more realistic solution would involve post-processing the backplane wafers to add an extra layer of randomised electrode mirrors over the existing regular array of pixels.

The pixel flat fill factor can seriously affect the optical throughput of a correlator system. The power in the zeroth-order replica of the transformed output of the SLM is proportional to the flat fill factor squared. Cascading two devices in a correlator results in the output power at the detector being proportional to the flat-fill factor to the fourth power. For an unplanarised SBS256 device the flat fill factor is 23%, so the throughput would be less than 0.3%, while for a planarised SBS256 (see Section 5.3) with a 84% flat fill factor, the throughput is almost 50%.

3.2.3 Optical computing and distributed processing

For optical computing systems with up to a few thousand channels, SEED-type devices appear to be the most attractive because of their potential to work above 100 MHz (see Ref. [83] for a selection of several SEED-type devices and systems designed for optical computing). FLCOS SLMs are probably too slow to be competitive with these, however, they may be suitable for systems that do not need to operate as fast but do require 10^4–10^5 channels.

3.3 Grey-scale generation

Nematic liquid crystals have been widely used in flat panel displays because of their grey-scale capability [80,114]. Used in combination with active-matrix addressing, a reasonable number of grey levels (4–16) can be generated in real-time. Electroclinic and distorted helix effect FLCs can also exhibit analogue response and have much faster switching times than nematics. However, research into electroclinics is at a very early stage; mixtures currently being investigated require high fields and very accurate control of the the operating temperature (< 0.025 °C for 100 grey levels) [16]. It would therefore be very difficult to incorporate electroclinics with Si backplanes, unless care is taken to ensure that switching in the electronic circuitry did not cause any local temperature variations.

McKnight *et al* [68] have recently reported on an analogue 128×128 DRAM-type LCOS SLM that uses a distorted helix effect FLC. However, as this FLC geometry is not bistable, the DRAM charge leakage problem significantly degrades the uniformity across the array when the frame scan time is comparable to the discharge time of the pixel storage capacitors.

Although the SSFLC geometry is generally binary [1], several techniques can be used to generate grey-scale in SSFLC devices [38], including charge-metering, spatial multiplexing, and temporal multiplexing. These will now be considered.

3.3.1 Analogue charge metering techniques

Certain cell configurations [37,51] can be used to generate grey-scale on SSFLC devices using charge metering drive techniques. By using a well controlled rubbing process on one of the cell's alignment layers, the SSFLC can be broken down

---

[1] I will not consider the possibilities of multi-state switching: see Ref. [14] for more information.

into a multidomain structure with a fairly homogeneous domain size. To switch the whole surface A_{MAX} of a cell from one bistable state to the other, a charge $Q_{MAX} = 2P_S A_{MAX}$ must be applied, where P_S is the FLC spontaneous polarisation. Applying a charge $Q < Q_{MAX}$ will consequently only switch an area $A < A_{MAX}$, leading to the formation of a mixture of black and white domains. The spatial distribution of these domains is determined by the more or less random distribution of defects in the FLC layer which act as nucleation sites for the domains.

The optical transmission of the cell can be varied between 0 and 100% by adjusting the amplitude or duration of the charge-dumping switching pulse. Either voltage or current mode switching pulses can be used. The voltage pulse technique has been used successfully by Hartmann [38] in a 96 × 108 active-matrix-addressed transmissive SSFLC device, while Killinger *et al* [51] have investigated the current mode approach in test cells. They found that a sequence of current pulses can provide a more linear grey-scale response than the voltage mode approach. To ensure grey level reproducibility, it is essential to incorporate a blanking step into the electronic addressing scheme, before the writing step is performed.

Domain sizes can be made as small as 2 μm, so there is the potential to generate large numbers of grey levels on relatively small pixels. If charge metering circuitry was included on a one-transistor DRAM-type FLCOS SLM, hundreds of grey levels could be generated on each pixel (for a typical 30 × 30 μm^2 pixel). Of course, it should be noted that the FLC must be truly bistable to overcome the DRAM charge leakage effects.

3.3.2   Multiplexing techniques

Multiplexing techniques are often denoted by the name dither and can be spatial or temporal in character, or a combination of both.

Spatial multiplexing

With spatial multiplexing the pixel is subdivided into parts or subpixels. Using an SLM where every pixel is the same size, a group of 16 pixels would be required to generate a 'super' pixel with 17 grey levels. This has the disadvantage of loss of effective resolution and redundancy; there are many different ways of achieving each grey level. A more efficient method would involve binary weighted pixels. To achieve 16 linear grey scales, the subpixels could be binary weighted as in Fig. 3–4. With this technique, the number of greys obtained is 2^S, where S is the number of subpixels.
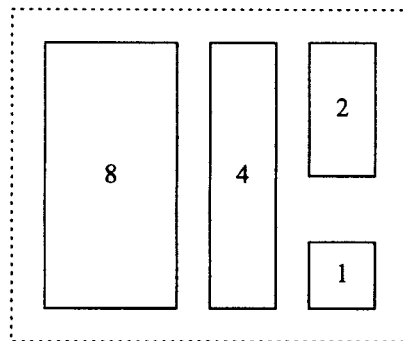
Figure 3–4: Binary weighted subpixel division to generate 16 grey levels. The number of greys obtained is 2^S, where S is the number of subpixels.

Temporal multiplexing

Temporal multiplexing (TM) involves splitting each grey-scale image into a number of time-sequential subframes. Each subframe has an associated bit-plane generated by an appropriate encoding algorithm. These bit-planes are then sequentially scanned into the SLM so that by integrating at the detector plane over a number of subframes, the desired grey-scale image is produced. In the case of display applications, the integration is carried out by the human eye.

Fig. 3–5 shows how a 4-grey level image is built up using a simple linear encoding algorithm to create the bit-planes. I used this technique to generate grey-scale on the 16 × 16 device [105]. For the six grey levels shown in Fig. 3–6(a), five subframes were required. Each subframe lasts 500 μs, resulting in a
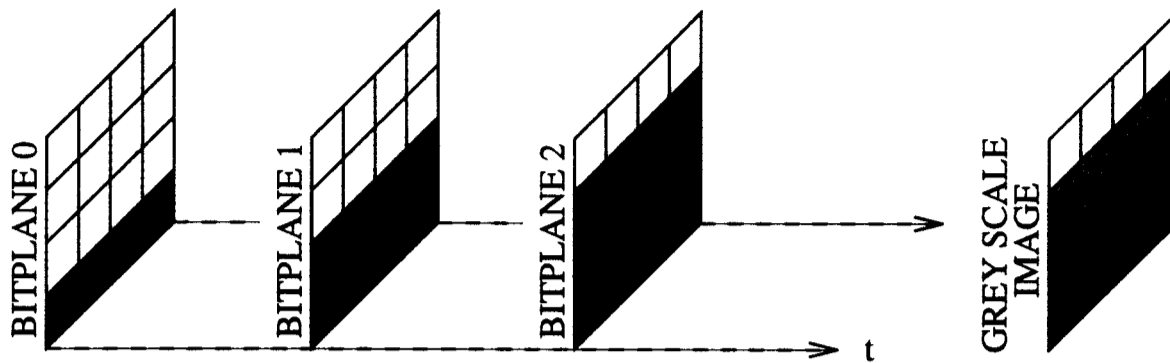
Figure 3–5: Linear encoding algorithm to generate four grey levels.

frame rate of just above 133 Hz. The linearity is satisfactory (see Fig. 3–6(b)), however, notice the variation across each pixel; this is caused partly by 'hillocks'[2] on the aluminium pads, and partly by defects in the FLC layer.

The linear encoding algorithm is satisfactory for generating a small number of grey levels, however, each extra grey level requires an extra subframe: the frame rate is inversely proportional to the number of grey levels. A more efficient binary weighted encoding algorithm using a pulsed light source is illustrated in Fig. 3–7. In this case, only one extra subframe is required to double the number of greys. For example, 256 grey levels can be generated at 125 Hz using eight binary encoded bit-planes, with 1 ms subframes. In Chapter 6, I describe how this technique can also be used if the detector system, rather than the light source, is pulsed in a binary weighted fashion. The pulsed light source technique could also be extended to colour rendition by using three coloured (red, green, and blue) light emitting diodes (LEDs) as the light sources. These multiplexing techniques permit FLCOS SLMs to be used in helmet mounted display systems for virtual reality, night vision, and other military applications [113].

---

[2]Stress-relief 'hillocks' form when the metal layer is sintered during the MOS fabrication sequence.

(a)

DISTANCE (μm)
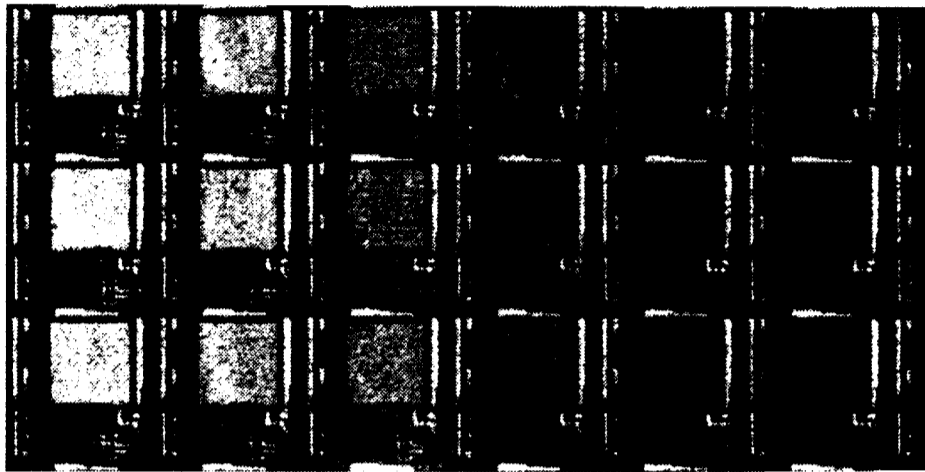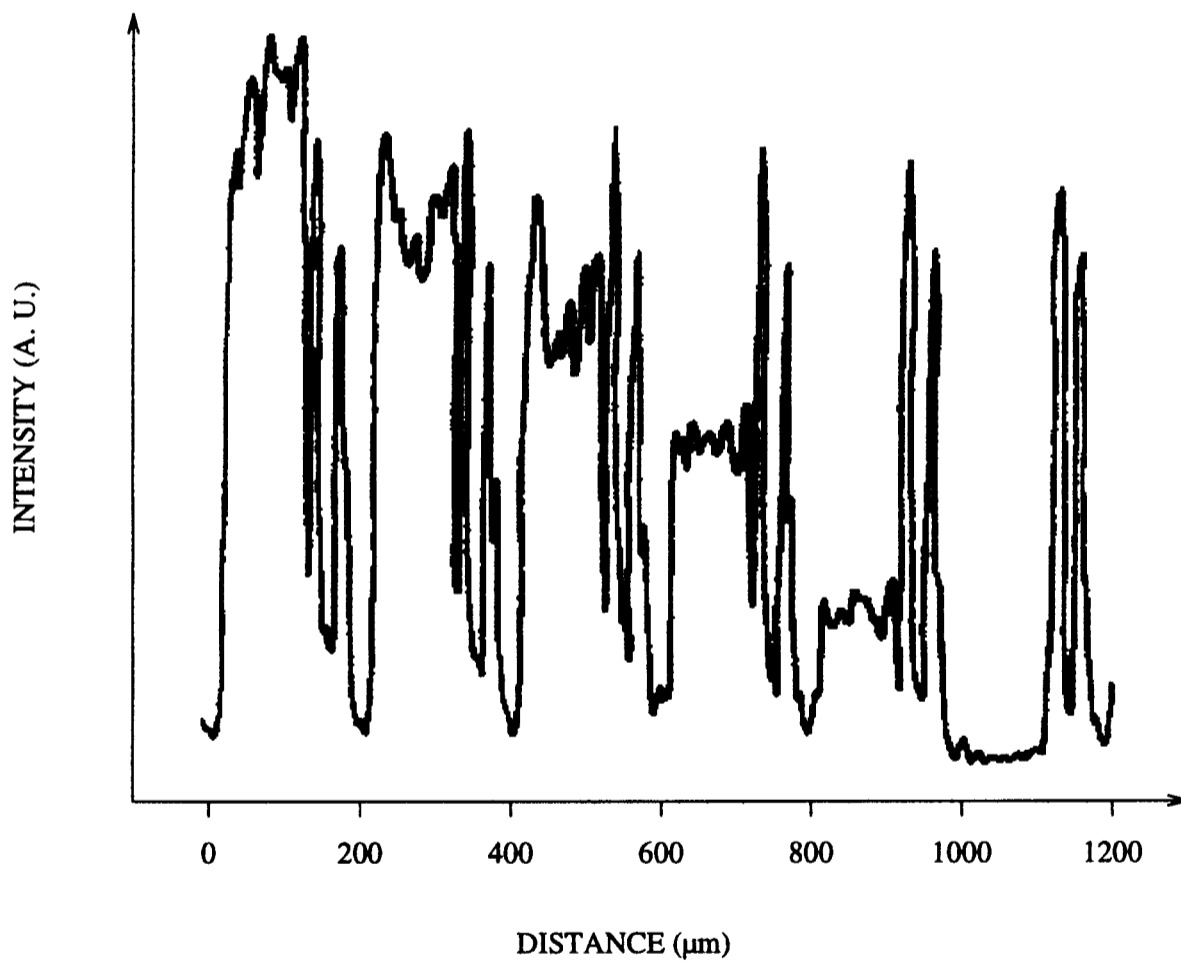
(b)

Figure 3–6: (a) Six linearly encoded temporal multiplexed grey levels on the 16 × 16 device, and (b) corresponding microdensitometer trace through a section of the bottom row of pixels.
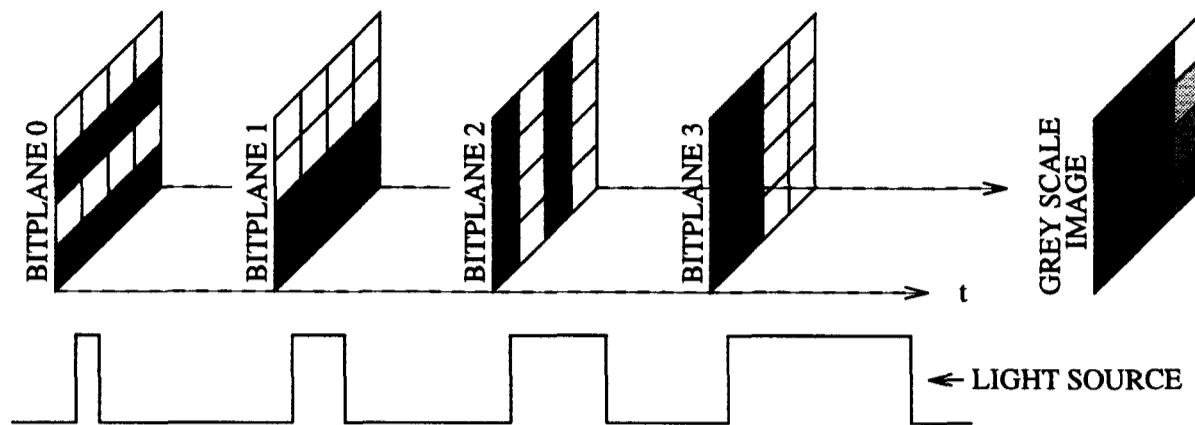
Figure 3-7: Binary encoded bit-planes to generate eight grey levels.

3.4 Artificial neural networks

In recent years artificial neural networks (ANNs) have emerged as powerful tools capable of solving many pattern recognition, classification, and prediction problems. ANNs are highly parallel arrangements of many simple processing elements, or neurons, connected together. Every connection between neurons has an adaptive coefficient, or weight, assigned to it, which determines the connection strength between them. A positive weight can help excite a neuron, while a negative one can help to inhibit it.

Many different neural network architectures and models have been developed and are well documented elsewhere (see Refs. [40,88] for a thorough study of neural algorithms and architectures, or Lippmann's two papers [57,58] for a more condensed review). I will briefly consider the multi-layer perceptron [86] because it illustrates the basic concepts of many neural architectures and also because it is probably the most widely studied.

3.4.1 The multi-layer perceptron

The multi-layer perceptron (MLP) consists of two or more layers of neurons where each neuron in a layer is connected to all the neurons in the next layer, but not

connected to any in its own layer, as shown in Fig. 3–8. The input layer distributes the input pattern, through appropriate weights, to the hidden layer.

The neurons in the hidden layer perform a nonlinear transfer function on the sum of their weighted input signals. This can be represented mathematically by a vector-matrix product and a nonlinear operator $f\{\}$:

$$J_j = f\left\{\sum_{i=1}^{m} W_{ij}I_i\right\} \tag{3.1}$$

where J_j represent a neuron in the hidden layer, the vector I represents the states of the m input neurons, and W is the interconnection weight matrix between the two layers. Similarly the output neurons can be described by:

$$K_k = f\left\{\sum_{j=1}^{n} W_{jk}J_j\right\} \tag{3.2}$$

An error back-propagation algorithm [86] is used to adjust the interconnection weights between the neurons. Briefly, an input pattern is presented to the network and it generates an output pattern dependent on the network's connection weights. If the output pattern does not correspond to a specified target pattern, then the weights are adjusted by back-propagating the errors so that the next time the input pattern is presented, the output will be closer to the target. This is repeated until the output corresponds to the target.
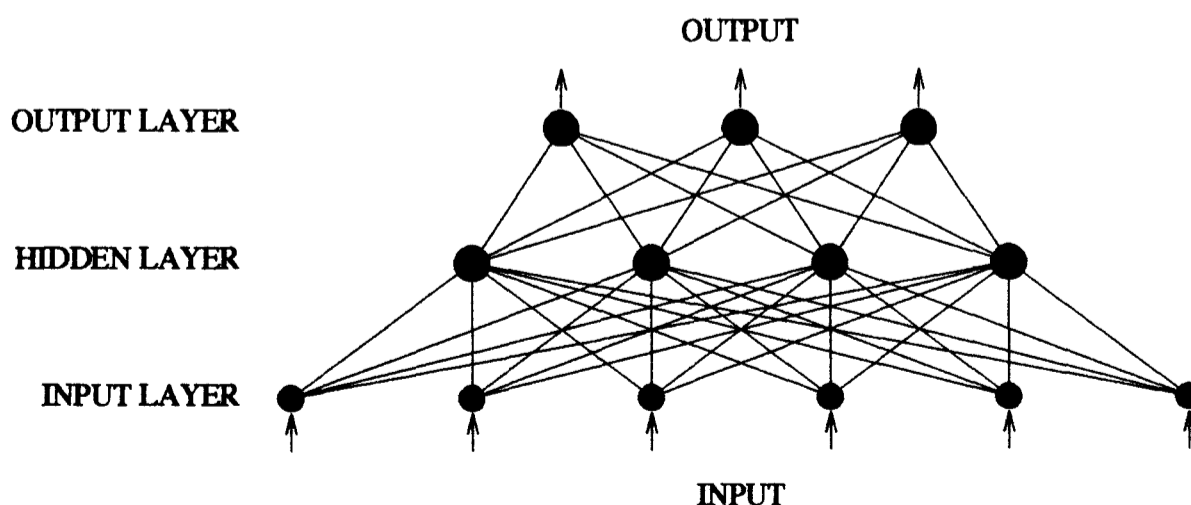
Figure 3–8: The multi-layer perceptron.

3.4.2 Electronic implementations

Much progress has been made in understanding the behaviour and limitations of many different architectures, largely through software simulations and theoretical analysis. However, efficient hardware implementations are crucially important in the long term if the full potential of neural networks is to be realised.

While, most ANN software simulations can use single or double precision floating point representations for the weights and neuron states, most hardware designs, whether electronic or optoelectronic, are restricted to limited precision or even binary weight and state representations. For electronic implementations, designers must make a tradeoff between the mathematical precision required and the silicon area, which will in turn affect the chip size, cost, and attainable neuron density. For example, in a digital design the size of the multiplication hardware increases as the square of the word size of the multiplicands. In analogue implementations, circuits can be made smaller but tend to suffer from noise, process variation, and limited precision. References [115,73] provide a description of many digital and analogue implementations, and consider their inherent advantages and disadvantages.

The intrinsic vector-matrix multiplication function of a layer of neurons can be mapped onto the two-dimensions of a silicon chip. The one-dimensional input vector is fed via the edge of the chip, and each neuron is fanned-out by running a signal wire along the appropriate column of the two-dimensional interconnection matrix. If the output of each $I_i W_{ij}$ component is a current, as in analogue [42] and pulse-stream implementations [34], the summing function for each neuron J_j can be achieved simply by running a wire along each row of the array.

3.4.3 Optoelectronic implementations

The two-dimensional architecture used in electronic implementations can be extended to the optoelectronic domain. Stearns [96] has implemented a 49-36-10 MLP with electronic inputs and a optically presented weight matrix using a two-

dimensional array of amorphous silicon (a-Si) photoconductors programmed by an AMLCD placed directly over the array.

A number of researchers have implemented neural network architectures using the vector matrix multiplier (see Fig. 3-1) with optical input signals [6,22,81,117]. In practice, some alterations must be made to the system to permit representation of negative weights. A common method of implementing negative weight values is to spatially separate the weight matrix into two sections – one for the positive values, the other for the negative values. Farhat *et al* [22] implemented a 32 neuron Hopfield model [43] network that used an array of LEDs to represent the input vector, and two photographic transparencies to represent the positive and negative sections of the binary bipolar interconnection weight matrix. The signals generated by an array of photodetectors at the end of each arm of the system were then electronically combined and thresholded to give the neuron states which were then fed back to the input LEDs. As photographic masks were used, this system was obviously not adaptable in real-time.

Zhang *et al* [117] built a 128-32 neuron, single layer perceptron (SLP) for character recognition, using a polarisation-based optical system. A 1 × 128 transmissive stripe FLC SLM is used for the input vector, obviating the need for a cylindrical fan-out lens, while a 128 × 128 FLC SLM is used for the weight matrix (see Fig. 3-9). Each weight comprises a strip of four pixels to permit five spatially multiplexed grey levels corresponding to weight values: -4, -2, 0 , +2, and +4. A polarising beam splitter separates the vertical and horizontal polarisation components (representing positive and negative weight components respectively), which are then fanned-in onto photodetectors using cylindrical lenses. Signals from appropriate pairs of photodetectors are electronically amplified, compared (subtracted), and thresholded to give an output signal for each of the 32 output neurons. A personal computer (PC) is used to control the system calculate the weight adjustments during the learning cycle.

Zhang *et al* [118] also implemented a second-order neural network by rearranging the system to pass the optical beam through the input SLM twice.
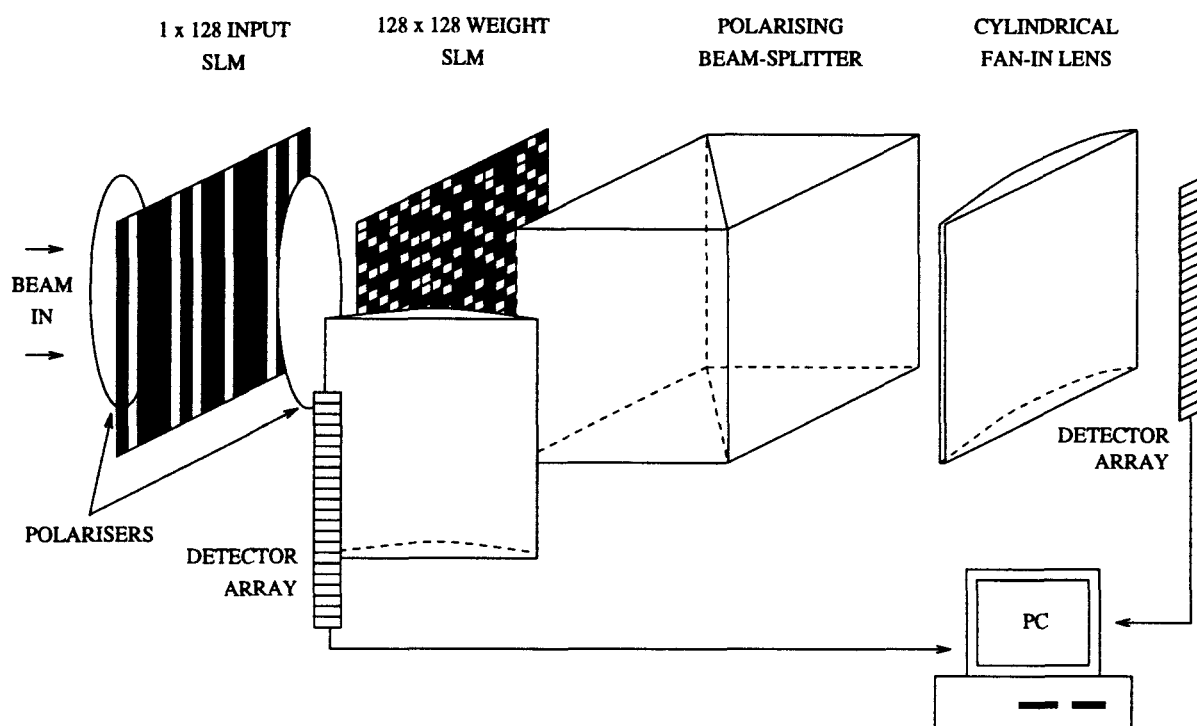
Figure 3–9: The polarisation-based single layer perceptron implemented by Zhang *et al* [117].

This introduced quadratic product interconnections between the neurons which enhanced the classification ability of the network, by permitting translation and scale invariance with regard to the input vectors [26]. The weighting scheme was also changed to permit 17 grey scales by grouping the pixels on the weight SLM into 4 × 4 blocks.

Oita *et al* [81] developed an MLP for character recognition using a two stage VMM architecture that included programmable binary mode LCDs for the two weight matrices. The system only represents binary unipolar weights so an input dependent thresholding operation is required.

Some optoelectronic ANN implementations incorporate a learning algorithm to adjust the weighted interconnections to the required values. Mao and Johnson [62] have designed a 2 × 32 array of integrated neurons that compute the error and weight modification for a bipolar supervised Delta rule [64] architecture. Each neuron compares two optical input signals. The difference is thresholded and compared with a target value. If there is a difference between the thresholded and target value, one of a pair of modulator pads is turned on to increase or

decrease the corresponding weight on an optically-addressed SLM. Wagner and Slagle [110] have proposed a similar system to implement a self-aligning unsupervised competitive learning architecture. They plan to use a dynamic volume hologram for the weighted interconnections, and their sparse-topology winner-take-all FLCOS SLM [93] as the neuron array.

Many of the optoelectronic ANN implementations described above use the vector-matrix crossbar architecture. They also require two optical pathways to implement both positive and negative activations. The MMC architecture is more suitable for applications where the input is two-dimensional. Later in Chapter 6, I will present a scalable, single optical-path optoelectronic ANN system that uses the MMC architecture, and a 2's complement variant of the pulsed temporal multiplexing scheme described in Section 3.3.2

3.5   Summary

In this chapter I have reviewed some applications for which FLCOS SLMs are particularly suitable. They can be used in crossbar switching systems, as input and filter planes in correlation and Fourier processing systems, in display applications, and as weight planes in optoelectronic ANNs. I also demonstrated temporally-multiplexed grey-scale on the 16 × 16 FLCOS SLM.

Chapter 4

The SBS256 - backplane design issues

In this and the next chapter, I cover the successful development of the SBS256 –
a FLCOS SLM that incorporates a static random access memory (SRAM) latch
and an exclusive-OR (XOR) gate at each pixel. In this chapter I concentrate
on backplane design issues such as process selection, pixel circuit functionality,
backplane architecture, and power dissipation.

4.1   Introduction

Most of the medium-resolution FLCOS SLMs designed elsewhere have been based
on the single transistor dynamic random access memory (DRAM) type pixel [35,
66,106]. The DRAM-type pixel offers high density at the expense of limited
liquid crystal drive capability. Preliminary investigations [105] with the Mk II
16×16 test-bed device illustrated the potential of the SRAM-type pixel. The
pixel provides a stable drive waveform for the overlying FLC layer, permits easy
charge balancing, and is robust under high read-beam intensities. The SRAM-
type device is also easy to control with a simple interface, and can be used in a
constant illumination optical system. However, the prototype device did have a
low resolution: for many optical processing applications, higher resolution would
obviously be desirable.

56

To compete with DRAM-type devices that were being developed elsewhere, a 256 × 256 SRAM-type pixel array was considered both desirable and feasible. The electronic characteristics of the backplane determine many of the optical characteristics of the SLM. For example, the time it takes to address the pixel array is an important factor in determining the optical frame rate of the device. Similarly, the pixel layout affects the fill factor and thus the device's optical efficiency. For a 256 × 256 array, the inherent parallel architecture of the silicon backplane imposed several constraints on the potential performance of the device. Decisions made throughout the design cycle resulted in many compromises between the functionality, speed, power dissipation, reliability, and yield of the device. I will discuss these issues throughout this chapter.

4.2   Process selection

The prototype 16 × 16 and 50 × 50 SRAM-XNOR backplanes were implemented in n-channel metal oxide semiconductor (nMOS) technology. However, complementary metal oxide semiconductor (CMOS) has since become the dominant technology offered by silicon foundries for general purpose IC applications. This dominance is primarily due to its low power dissipation. Low power dissipation permits extremely dense integration; some state-of-the-art microprocessors contain several million transistors [82] on a single chip. Compared with their bipolar and nMOS equivalents, digital CMOS circuits dissipate power only during switching events, and not when they are in stable or non-switching states. For these reasons, I chose CMOS as the most appropriate technology for the backplane.

4.2.1   Which CMOS process ?

From a survey of possible silicon foundries, the Applied Optics Group at the University of Edinburgh selected the Austria Mikro Systeme [28] (AMS) silicon foundry to fabricate a series of SLM backplanes designed by the group [1]. The

Process Name	Min. Trans. Geom. (μm)	Operating Voltage (V)	Max. Die Size (mm^2)	Comments
CYB	0.8	2.2–5.5	7 $\times$ 7	epitaxial wafer, n-well, double poly, double metal
CAE	1.2	2.2–5.5	14 $\times$ 14	epitaxial wafer, n-well, double poly, double metal
CBH	2.0	2.2–11	14 $\times$ 14	n-well, double poly, double metal
CBK	2.0	2.2–50	14 $\times$ 14	n-well, LDD, double poly, double metal
CCD	3.0	2.5–5.5	18 $\times$ 22	p-well, single poly, double metal

Table 4–1: A selection of CMOS processes available from the AMS silicon foundry.

foundry offers a number of CMOS processes with various integration densities, operating voltages, and die sizes (see Table 4–1).

The 3 μm CCD[1] process, one of their older processes, is used for the 512 $\times$ 512 pixel one-transistor DRAM device [10]. The device is an expansion of the 176 $\times$ 176 device, and many of the layout cells have been re-used with minimum redesign. It has a 30 μm pixel pitch, which results in a large, and consequently, relatively slow backplane capable of a data transfer rate of less than 10 MHz. As mentioned in Section 2.4.2, a SRAM-XOR type pixel should fit into a pixel pitch of 30–40 times the minimum transistor gate length for a given process. For the 3 μm process, a SRAM-XOR pixel with ten transistors, would probably occupy at least 100 $\times$ 100 μm^2, so a 256 $\times$ 256 pixel array would not fit onto the maximum available 18 $\times$ 22 mm^2 die size for the CCD process.

The 2.0 μm CBH process was attractive because it has a maximum operating voltage 11 V, but a 256 $\times$ 256 pixel array would probably not fit onto the 14 $\times$ 14 mm^2 die available for the process. The high voltage CBK version, utilises lightly doped drain (LDD) regions to help limit the electric field strength at the drain junctions. However, each high voltage transistor requires approximately five times the area of a normal transistor so a 256 $\times$ 256 pixel array would

---

[1]This process designation should not to be confused with charge coupled device technology.

almost certainly not be possible. The 3-transistor, 64 × 64 pixel device [94] is implemented in this process with a 120 μm pixel pitch.

The CAE 1.2 μm process was considered the most suitable technology for fabricating the 256 × 256 device because it combines small geometry with a 14 × 14 mm^2 die size. It is also an epitaxial process and is therefore much less susceptible to latch-up [77]. However, it has a lower operating voltage (5.5 V) than the 2 μm CBH process, so FLC switching times can be expected to be slower. The 0.8 μm CYB process became available while the backplane was being designed but has the disadvantage of small die size (7 × 7 mm^2), a restriction imposed by AMS process engineers to ensure reasonable yield figures during the early stages of commissioning the process[2].

4.2.2   Process variation

For every CMOS process, transistor characteristics such as threshold voltage, carrier mobility, and parasitic capacitances, vary from batch to batch, wafer to wafer, chip to chip, and even from transistor to transistor. The silicon foundry can only guarantee that the characteristics will fall within certain limits and usually provides the circuit designer with three sets of parameters for each type of transistor. One set contains the typical or average parameters while the other two, known as the fast and slow models, contain the extreme values that can be expected for the process. These models are used with a circuit simulator such as HSPICE [108] to model the effects of process variation on circuit functionality and performance. The circuit designer must ensure that his circuit will perform satisfactorily for any combination of transistor models. For example, a shift register may be required to operate up to 40 MHz, so the worst-case combination for high frequency operation would be when all the pMOS and nMOS transistors

---

[2]At the time of writing, AMS engineers were considering relaxing the die size restriction to 10 × 10 mm^2 for SLM backplanes, on the condition that the customer takes full responsibility for low yield results.

in the circuit on a fabricated chip are slow (the combination is referred to as SS). Similarly the worst-case combination for power dissipation in a circuit or chip is when all the transistors are fast (FF).

4.3 Transient current spikes

Even though CMOS offers low static power dissipation, the switching power dissipation can often limit the operating frequency of many ICs. However, the frame rates of SLM backplanes are generally very slow compared to typical electronic switching frequencies, so the switching power dissipation is less significant than in, say, a microprocessor (see Section 4.7). Nevertheless, the highly parallel architecture of an SLM, and the potential magnitude of the transient switching current spike associated with setting or resetting the whole array simultaneously, could seriously affect the performance of a device by causing excessive power supply noise, latch-up, or electromigration. Before considering these effects, I will briefly explain why the current spikes occur.

4.3.1 The mechanism: switching an inverter

The basic building block of many digital circuits is the inverter (see Fig. 4-1(a)). During the edge of an input signal V_{IN}, there will always be a short-circuit current flowing from the supply rail to ground (see Fig. 4-1(b)). This current flows as long as V_{IN} is higher than V_{Tn} and lower than $V_{DD} - V_{Tp}$, where V_{Tn} and V_{Tp} are the threshold voltages of the nMOS and pMOS transistors respectively. For simplicity, assume the inverter has no capacitive load and is symmetrical, so

$$V_{Tn} = -V_{Tp} = V_T, \tag{4.1}$$

and

$$\beta_n = \beta_p = \beta, \tag{4.2}$$

where β is the common gain factor of the transistors (μA/V^2). Using the simple first-order MOS transistor formula, this lead to

$$I = \frac{\beta}{2}(V_{IN} - V_T)^2 \quad \text{for} \quad 0 \leq I \leq I_{max} \tag{4.3}$$

As the inverter is assumed to be symmetrical, the current reaches a maximum when V_{IN} equals half the supply voltage ($V_{IN} = V_{DD}/2$), so

$$I_{max} = \frac{\beta}{2}\left(\frac{V_{DD}}{2} - V_T\right)^2 \tag{4.4}$$

When a load capacitance is added to the output node V_{OUT}, the falling edge of the spike may become broader as the capacitor must be charged or discharged.

(a)   (b)
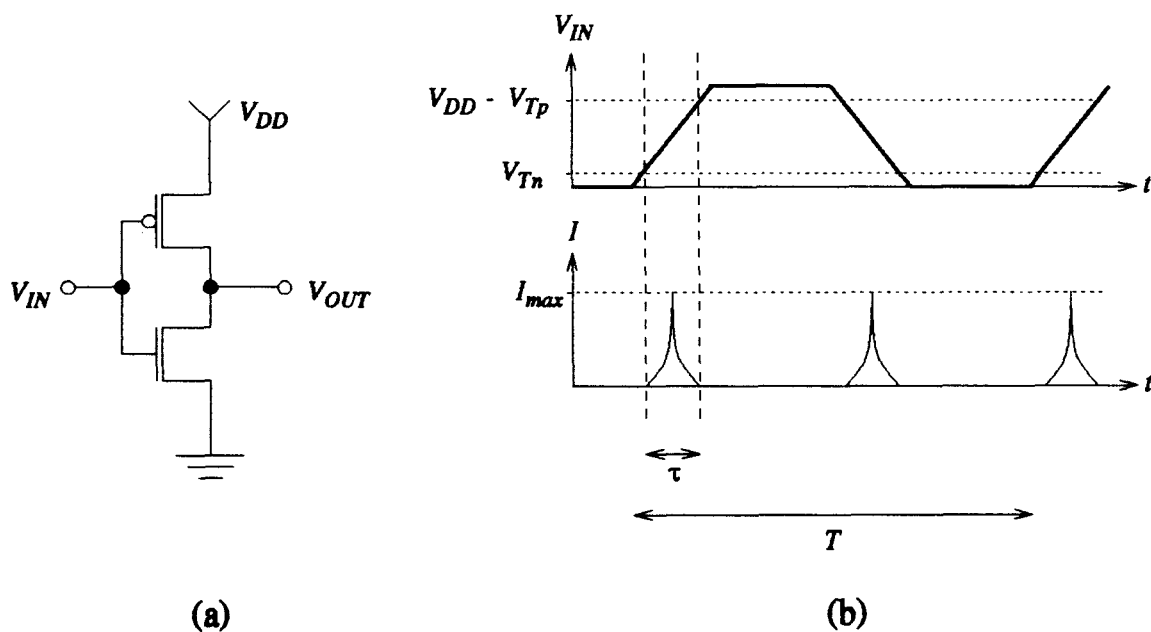
Figure 4–1: (a) Basic CMOS inverter, and (b) the transient current spike behaviour on switching.

HSPICE simulations predicted a worst-case (FF) I_{max} current spike of 100 μA for a single minimum size symmetric inverter with a 50 fF load. If every pixel in an array had an inverter, and they were all switched at once, a 256 × 256 array could generate a combined current spike of 6.5 A! Even if this simple extrapolation is an order of magnitude too high because process variation and signal skew are not taken into account, the effects of the spike could serious affect the functionality and performance of the device.

4.3.2 Power supply noise

Power is usually supplied to a die via bonding wires. These can have an inductance of 0.2–2 nH (0.5–1 nH/mm). The back emf generated in the bonding wires by a large current surge can add volts of noise to the power supply. Sensitive circuitry may be affected and could even switch to an erroneous state. The best way to reduce power supply noise is to ensure that large current spikes do not occur in the first place. At the design stage, careful circuit and timing analysis should show up any potential problems, that can then be reduced or eliminated.

4.3.3 Latch-up

The potential for latch-up, a pnpn self-sustaining low impedance state, is inherent in standard bulk CMOS circuit structures [77,99]. The parasitic pnpn circuit contains a pair of cross-coupled pnp and npn transistors that can be triggered into a latched state by some disturbance (see Ref. [27, pages 122–130] for a description of the mechanism and some protection strategies). The large transient current spike could affect the power and ground rails enough to turn on one of the bipolar transistors in a pnpn latch path, resulting in latch-up. Latch-up is a positive feedback effect, so the chip's power supply must be cut to reset the chip. For an SLM backplane, the current spike would probably be pattern dependent, therefore the possibility of latch-up would also be pattern dependent: this would obviously be unsatisfactory for a general purpose SLM.

4.3.4 Electromigration

High current densities in the metal buslines can result in electromigration effects [7]. The main force that acts on the lattice atoms in conductors arises from the momentum exchange between the lattice atoms and the electron wind being driven by the electric field. The collisions exert a force in the same direction as the electron flow. Vacancies move upstream where they condense to form voids,

while ions move downstream where they form crystals and hillocks. The voids can coalesce to cause a wire to open circuit, whereas the accumulation of ions downstream can short wires together. Electromigration effects are roughly inversely proportional to the square of the current density but are also dependent on temperature and the crystal structure of the conductor.

For the aluminium interconnect used in typical CMOS processes, it is normally recommended to keep the mean current density below 1 mA/μm^2, and the peak current density below 10 mA/μm^2. Again using simple extrapolation for the minimum size inverter mentioned above, the 100 μA current spike per inverter could require the main power and ground rails to be at least 650 μm wide for a 256 × 256 array.

During the design stage I was not sure whether the simple linear superposition of the individual current spikes was a valid assumption to make, but I felt it was better to err on the side of caution to ensure that the backplane could operate for this worst-case scenario. This decision affected both the pixel design and also the backplane addressing architecture: these are now discussed.

4.4 Pixel circuit

Each pixel comprises a SRAM latch, an XOR gate, an electrode mirror, and six control and supply buslines (see Fig. 4–2). The latch is for storage, while the XOR gate is used to charge-balance the FLC layer above the electrode mirror.

4.4.1 SRAM latch

The SRAM latch is based on the standard six-transistor circuit (see Ref. [27, page 390]). It consists of a pair of cross-coupled CMOS inverters connected by pass transistors to the DATA and /DATA[3] lines. The pass transistors are

---

[3]The slash (/) refers to the complement or inverse of the signal.

(a)

(b)
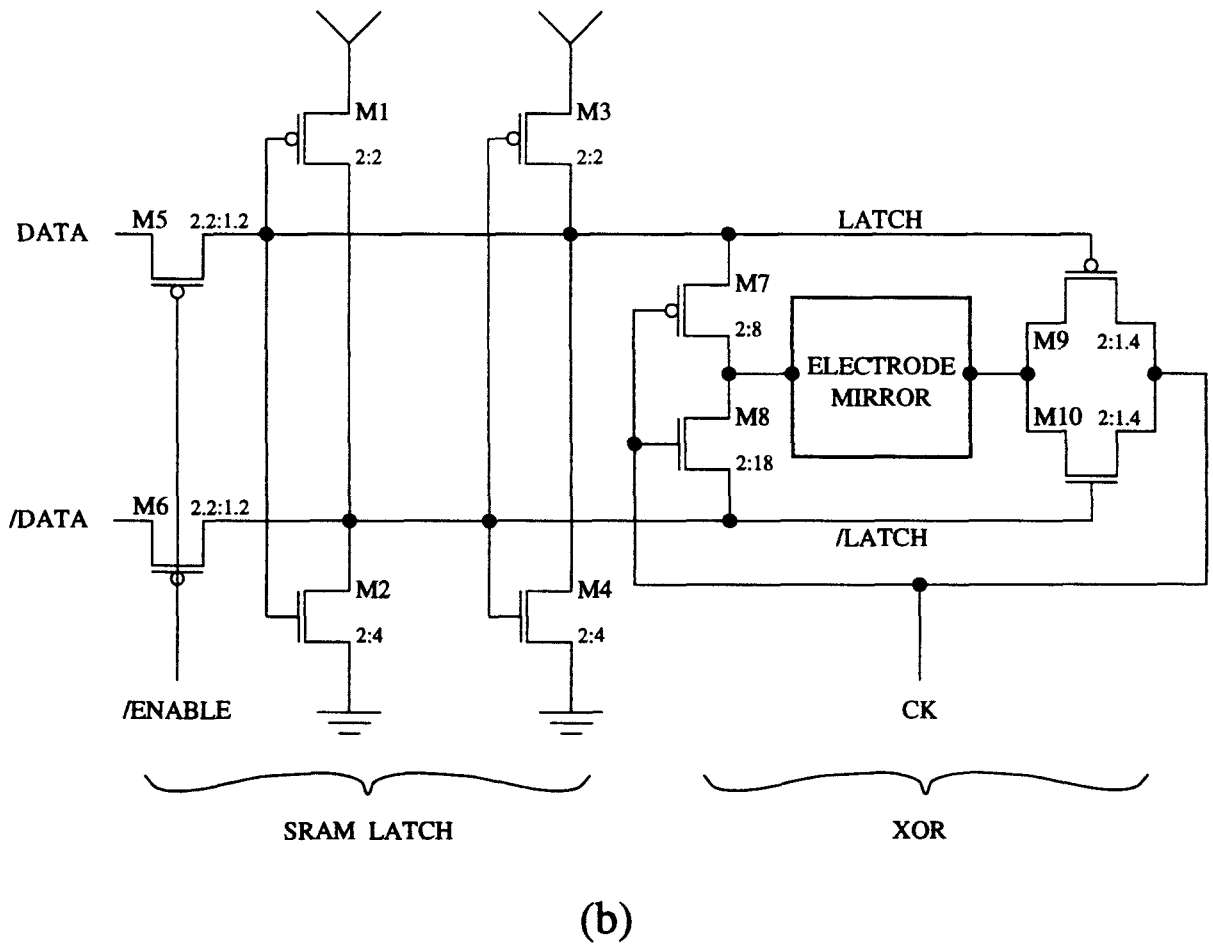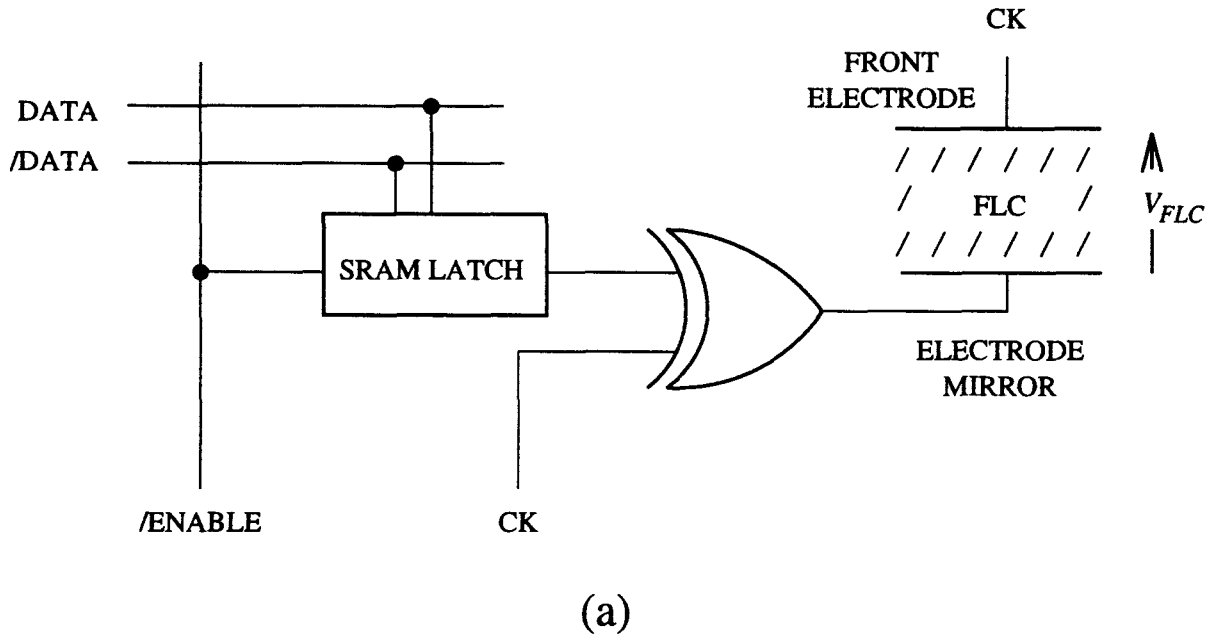
Figure 4–2: (a) Schematic diagram and (b) transistor-level diagram of the SRAM-XOR pixel. The transistor dimensions are given as W:L in μm. The slash (/) refers to an active low signal or complement of a signal.

pMOS rather than nMOS, so that with the appropriate drive scheme, the FLC directly above the /ENABLE access line tends to be switched to its OFF state (see Section 4.4.3).

To consider the write operation, assume a LOW (0 V) is stored at node LATCH and a HIGH (V_{DD}) at /LATCH, and we wish to set LATCH to HIGH and /LATCH to LOW. When the /ENABLE line is asserted, DATA is HIGH and /DATA is LOW. Node LATCH rises towards V_{DD}. The ratio of M5 to M4 must ensure that LATCH can rise to at least 2/3 V_{DD} to turn off M1. Node /LATCH falls to a voltage determined by the ratio of transistors M1 and M6, and the voltage on the gate of M1. In this manner the cross-coupled inverter pair is unbalanced and positive feedback then causes the latch to switch state.

While switching the latch from one state to another, the most difficult combination occurs for slow pMOS and fast nMOS (SF) transistors. Extensive HSPICE simulations were performed to ensure that the latch would switch reliably for this case. HSPICE simulations predicted a worst-case (FF) current spike of 60 μA whenever the pixel was set or reset. The spike could not be reduced without compromising the size of the pixel as all four transistors in the pair of cross-coupled inverters would have had to be lengthened. The effects of this current spike on the array addressing scheme are considered in the Section 4.5.

4.4.2   XOR gate

Standard CMOS implementations of XOR gates can require up to 16 transistors, obviously too many to fit into a small pixel. However, if a low specification gate is acceptable, and the true and complement of one of the signals is available, as is the case with the SRAM latch, the XOR function can be performed with just four transistors (see Ref. [27, page 31]). This cut-down XOR-gate consists of an inverter and transmission gate (see Fig. 4–2(b)). As mentioned above, HSPICE simulations showed that if minimum size transistors were used to implement the inverter, a worst-case (FF) 100 μA transient current spike occurs in each pixel when the input signal CK is toggled. In this case both inverter

Inputs to XOR gate		ELECTRODE	FRONT		Optical response
LATCH	CK	MIRROR	ELECTRODE	V_{FLC}	of FLC
0	0	0	0	0	OFF
0	1	1	1	0	OFF
1	0	1	0	$-V_{DD}$	OFF
1	1	0	1	$+V_{DD}$	ON

Table 4–2: Truth table and optical response of the FLC layer for the SRAM-XOR pixel. When a LOW is stored in the latch, the voltage across the FLC layer, V_{FLC}, is zero. In this case, the FLC tends to be in its OFF state for the monostable behaviour seen in the devices assembled in-house. (see Section 5.5 for more details). When a HIGH is stored in the latch, $V_{FLC} = \pm V_{DD}$, depending on the state of the global clock signal CK.

transistors are made as long and thin as space would permit to limit the current spike, and therefore minimise the potential latch-up and electromigration problems considered earlier. By doing this, the worst-case current spike is reduced to 7 μA per pixel for the rising edge of CK and 10 μA for the falling edge.

4.4.3 Pixel drive scheme

In each pixel, the data value stored in the SRAM latch and the global clock signal CK are the inputs to the XOR gate, whose output is connected to the second-level metal electrode mirror. The CK signal is also applied to the transparent front electrode. In this arrangement, the voltage signal on the electrode mirror is in phase or in anti-phase with the front electrode depending on whether a LOW or HIGH is stored in the latch (see Table 4-2).

As mentioned in Section 2.3.1, the main operating limitation of FLC is that it must be charge balanced to prevent chemical degradation; if a positive voltage pulse is applied across the FLC, it should be followed by a negative voltage pulse

of equal magnitude and duration. Inclusion of the XOR gate permits the FLC layer to be charge balanced simply by toggling the global CK signal on a 50% duty cycle (see Table 4–2 and Fig. 4–3). Consider the case where a HIGH is stored in the latch: if the CK signal is toggled, the voltage across the FLC layer, V_{FLC}, is $\pm V_{DD}$. Because of the charge balancing requirement, CK should be toggled on a 50% duty cycle, so the pixel is in its ON or reflecting state for 50% of the time and in its OFF or non-reflecting state for the other 50%.

Now consider the case where a LOW is stored in the latch: here the FLC state depends critically on the FLC cell configuration. If the FLC is in the fully bistable configuration, it remains in the state it was last switched to, however, as is described in Section 5.5, the SBS256 SLM devices assembled in-house seem to exhibit a monostable behaviour. For these devices, the FLC tends to favour one of the two possible optical states, corresponding to the OFF state of the pixel. So when $+V_{DD}$ is applied across the FLC (LATCH = 1 and CK = 1), it flips to its ON state, but when 0 V is applied, the FLC relaxes to (or stays in) the more favoured OFF state. This is fortunate because if the FLC was not fully bistable and each of the optical states were equally favourable, the FLC could relax to an intermediate tilt angle state when 0 V is applied across it. For SRAM-type pixels that can apply 0 V across the FLC, this would result in a significant reduction in the contrast of OFF pixels. The monostable behaviour permits an efficient SLM drive scheme to be used that takes full advantage of the SRAM-XOR pixel functionality and drive capabilities (see Section 4.6).

DRAM-type devices are often charge balanced using a pattern/inverse pattern drive scheme [106]. A synchronised, pulsed light source is then required to interrogate the device. The extra degree of freedom provided by having an XOR-gate in each pixel permits the device to be charge balanced without an inverse pattern appearing, so a continuous rather than pulsed light source can be used to interrogate it. A further advantage of the SRAM-XOR pixel over the DRAM-type is that the electrode mirror is continually driven by the power rails via latch and/or XOR gate transistors (see Fig. 4–2). This results in an effect-
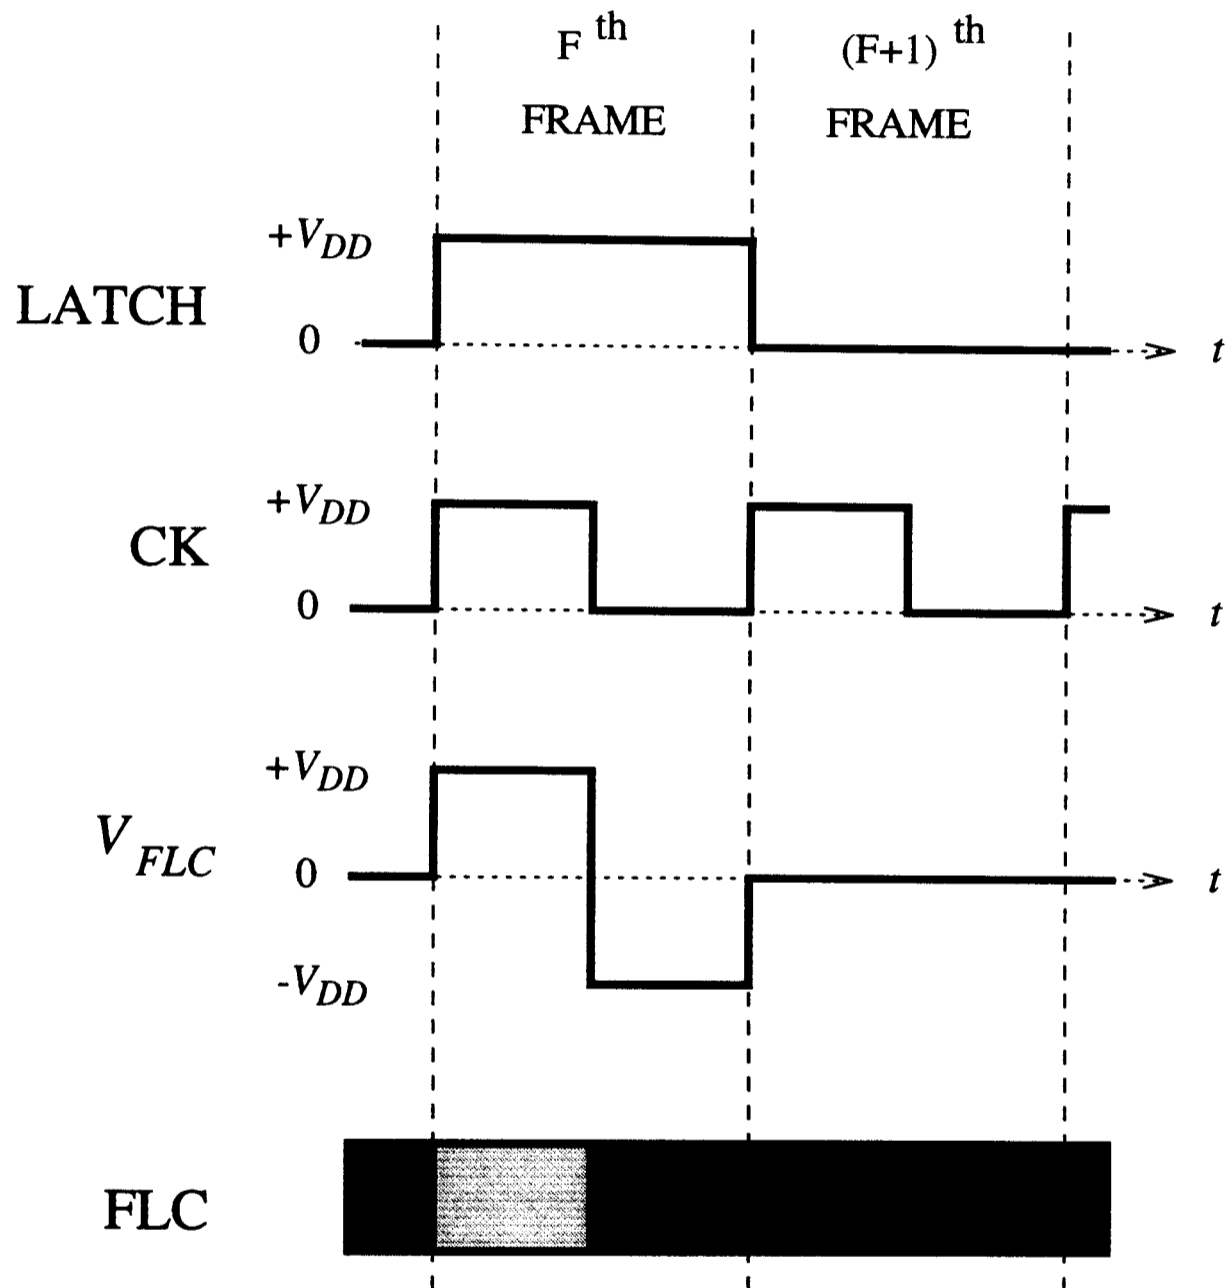
Figure 4-3: Pixel drive scheme. In frame F, the pixel has a logic one in its latch; it is on for 50% of the time, and for the other 50% it is off while the negative charge balancing voltage pulse is applied across the FLC. The pixel is off in frame (F+1) as a logic zero is stored in its latch.

ively unlimited amount of charge being available to switch the FLC, thus making the backplane suitable for use with current and future FLC mixtures possessing spontaneous polarisations greater than 30 nC/cm^2.

4.4.4 Pixel layout

The pixel layout (see Fig. 4–4) evolved through an iterative process. Throughout the design cycle, electronic functionality, current spike problems, potential yield, and fill factor all imposed conflicting constraints on the pixel layout. For example, to limit the magnitude of the transient current spikes, most of the transistors are longer than the minimum permitted by the design rules. These obviously take up more space, and leave less area for the electrode mirror.

The pixel has a 40 μm pitch. Each pixel requires two power buslines and four control signal buslines. Two levels of metal are available with the CAE process, so three signals are routed on each level. The GND, DATA, and /DATA signals are routed in first level metal (METAL1), and the VDD, CK, and /ENABLE buslines are routed in second level metal (METAL2) as they tend to switch the overlying FLC to its off state (see Table 4–3). The two power rails and the CK signal buslines are wider than the minimum design rules permit as they have to cope with significant transient current spikes, especially at the edges of the array where the current for a whole column or row enters (or exits). The others are minimum width (2.4 μm for METAL1, 2.8 μm for METAL2). When scaled up to a 256 $\times$ 256 pixel array, the busline spacings are likely to be the main limiting factor on backplane yield, so the spacing between all metal buslines within the pixel are at least 0.2 μm above the minimum design rules (1.4 μm for METAL1, 1.8 μm for METAL2).

The previous 16 $\times$ 16 and 50 $\times$ 50 SRAM-XNOR devices had flat fill factors of 30–35%. For this design, the fill factor was not as critical as other design constraints because post-processing planarisation was planned (see Section 5.3.1). However, it was also important that the device should be usable without planar-

VDD   CK
↓   ↓

DATA →

/DATA →

ELECTRODE
MIRROR

← GND

↑
/ENABLE
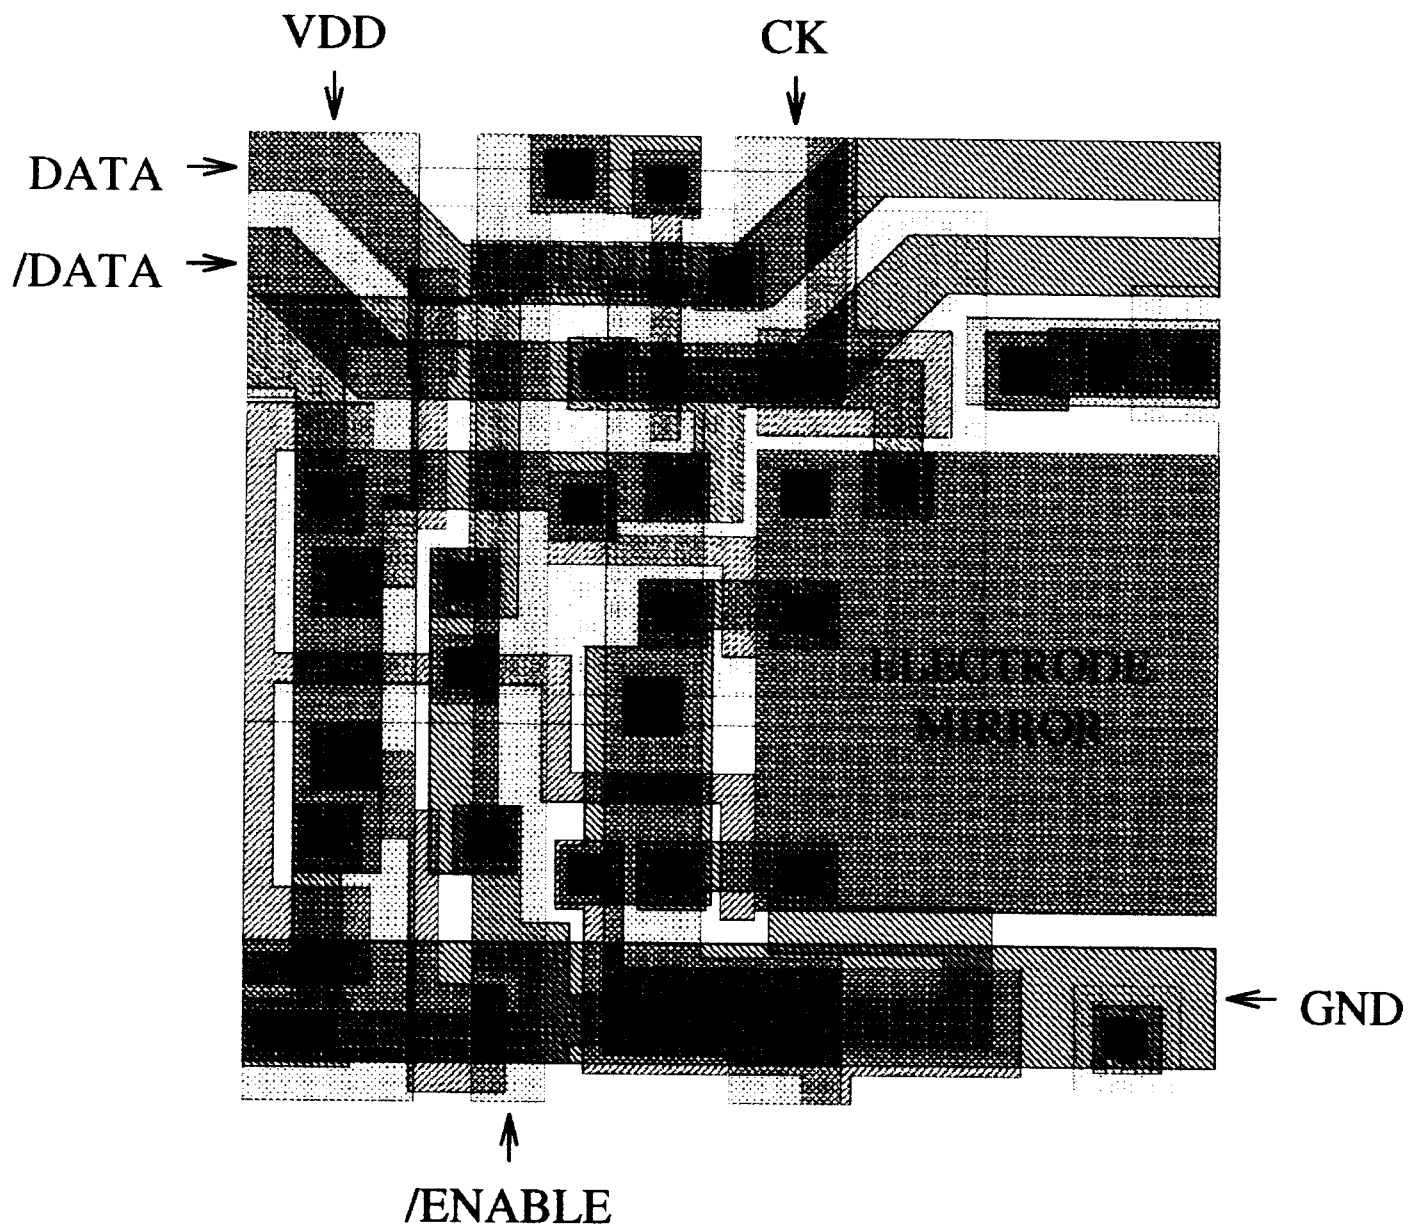
LAYER KEY:

☐	N-WELL	■	CONTACT
▨	DIFFUSION	▨	METAL1
▨	POLY	■	VIA
▨	BORON	▨	METAL2

Figure 4–4: Pixel layout. The pixel occupies an area of 40 × 40 μm^2.

Busline		FRONT ELECTRODE	
		0	V_{DD}
VDD	(HIGH)	$-V_{DD}$	0
CK	(HIGH)	0	0
	(LOW)	0	0
/ENABLE	(HIGH)	$-V_{DD}$	0

Table 4–3: The voltage across FLC overlying the buslines routed in METAL2 – VDD, CK, and /ENABLE. Under normal operating conditions, these buslines tend to switch the overlying FLC to the off state or leave it in the off state.

isation, so although the 23% pixel fill factor was lower than the previous devices, it was considered acceptable.

4.5 Backplane architecture

Figure 4-5 shows a block diagram of the main functional units on the backplane. The 256 × 256 pixels with a 40 μm pitch, result in a 10.24 mm × 10.24 mm active array. For the 14 × 14 mm^2 die available on the CAE process, this leaves a 1.5 mm border between the edge of the array and the bonding pads for placing glue spots during SLM assembly. The array is addressed using a column-at-a-time addressing scheme. Early in the design cycle I decided to use a decoder to select each column as it ensured that a maximum of 256 latches can change state at once. This helps limit the current spike generated by setting or resetting the pixel latches, to 4 mA; that is, 256 times 60 μA, as mentioned earlier. If a shift register with a rolling enable signal had been used, a potentially damaging 1 A

current spike could occur if a drive scheme with frame setting/blanking was used, as in the 176 × 176 device.
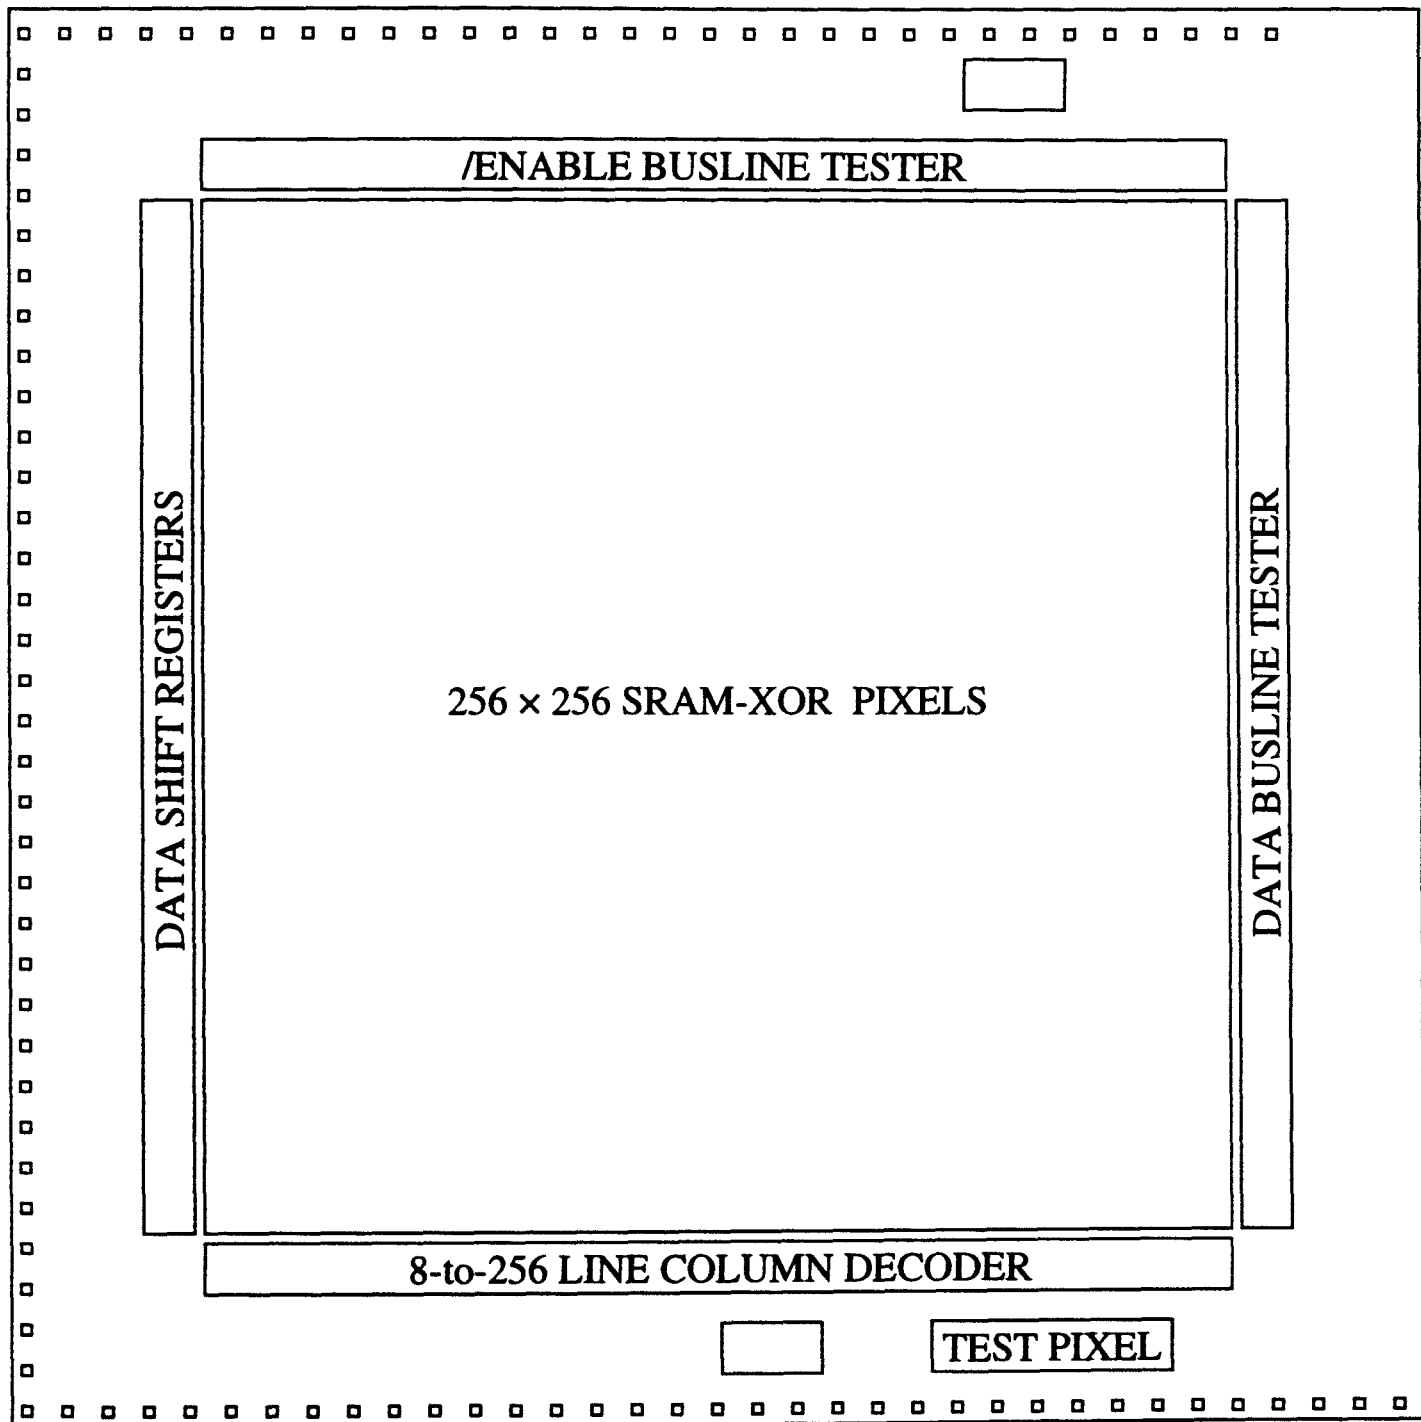
Figure 4–5: Block diagram of the main functional units on the backplane.

Data enters the backplane via a 32-bit bus. The data bits are loaded into thirty-two parallel 8-bit shift registers; eight clock cycles are required to assemble the 256-bit data word for a column. The data word is then latched to the auxiliary 256-bit register-buffer which drives the DATA and /DATA buslines. The latch-buffer circuit, previously implemented on the 176 × 176 backplane [106], helps optimise the addressing scheme as it permits the $(n+1)^{th}$ data word to be

assembled in the shift registers while the n^{th} word is written to its appropriate column. As both operations are performed simultaneously, the frame scan time can be kept to a minimum. The appropriate column /ENABLE signal is selected using the 8-to-256 line column decoder.

4.5.1   Data shift registers

Each cell in the thirty-two 8-bit shift registers consists of two latches in the D-type configuration. The data shift registers are operated using a two phase non-overlapping clocking scheme (see Fig. 4-6). To prevent possible race conditions, the non-overlapping clock signals, PHI1 and PHI2, propagate in the opposite direction to the data flow. The maximum clock frequency at which the shift registers can operate at is limited by a combination of the distributed resistance and capacitance of the clock buslines and the ON resistance of the input pad inverter driver. At high frequencies, these impedances smear the edges of the non-overlapping clocks enough to make them overlap.
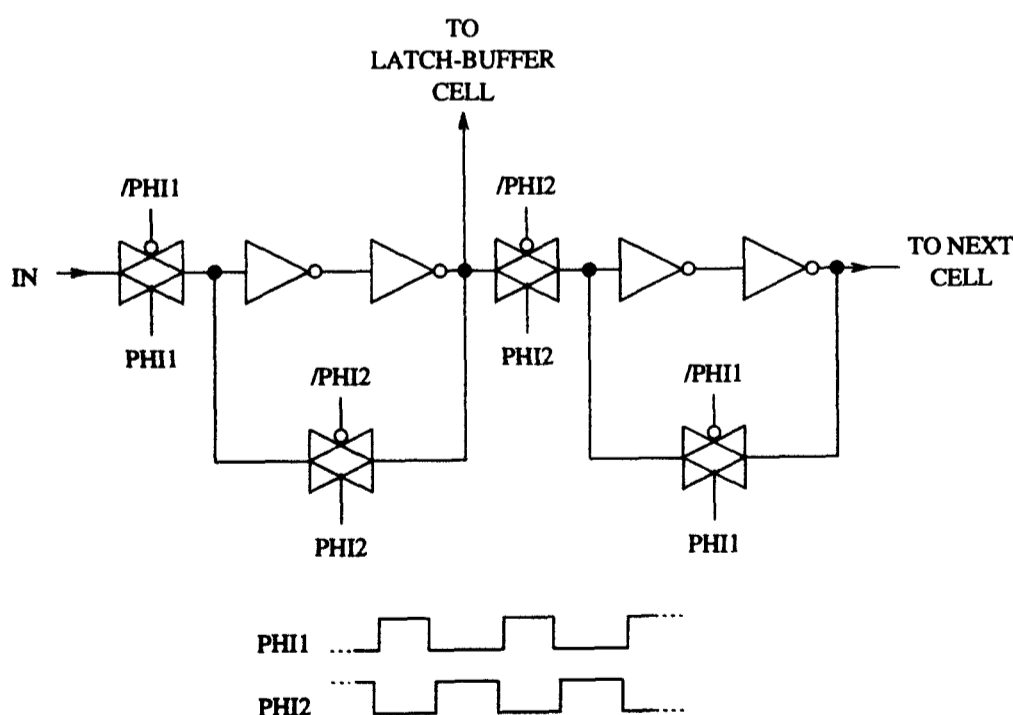
Figure 4-6: Schematic of a single data shift register cell.

The capacitance per unit area of METAL1 (0.029 fF/μm^2) is almost twice that of METAL2 (0.016 fF/μm^2). However, PHI1 and PHI2 are routed in METAL1

rather than METAL2, to help reduce the effect of FLC capacitance when the backplane is assembled into an SLM. For the buslines running the length of the shift register the distributed capacitance is about 7 pF while the resistance is about 500 Ω. In HSPICE simulations, each clock busline is modelled using a 3-π network, rather than a simple lumped RC model. Using 3-π network to represent a distributed RC busline results in a rise/fall-time error of less than 3% [87]. From HSPICE simulations, the worst-case (SS) operating frequency is 25 Mhz (see Fig. 4-7), which corresponds to a frame scan time of just less than 82 μs. Of course, from the inherent process variation, some die should work at higher frequencies.
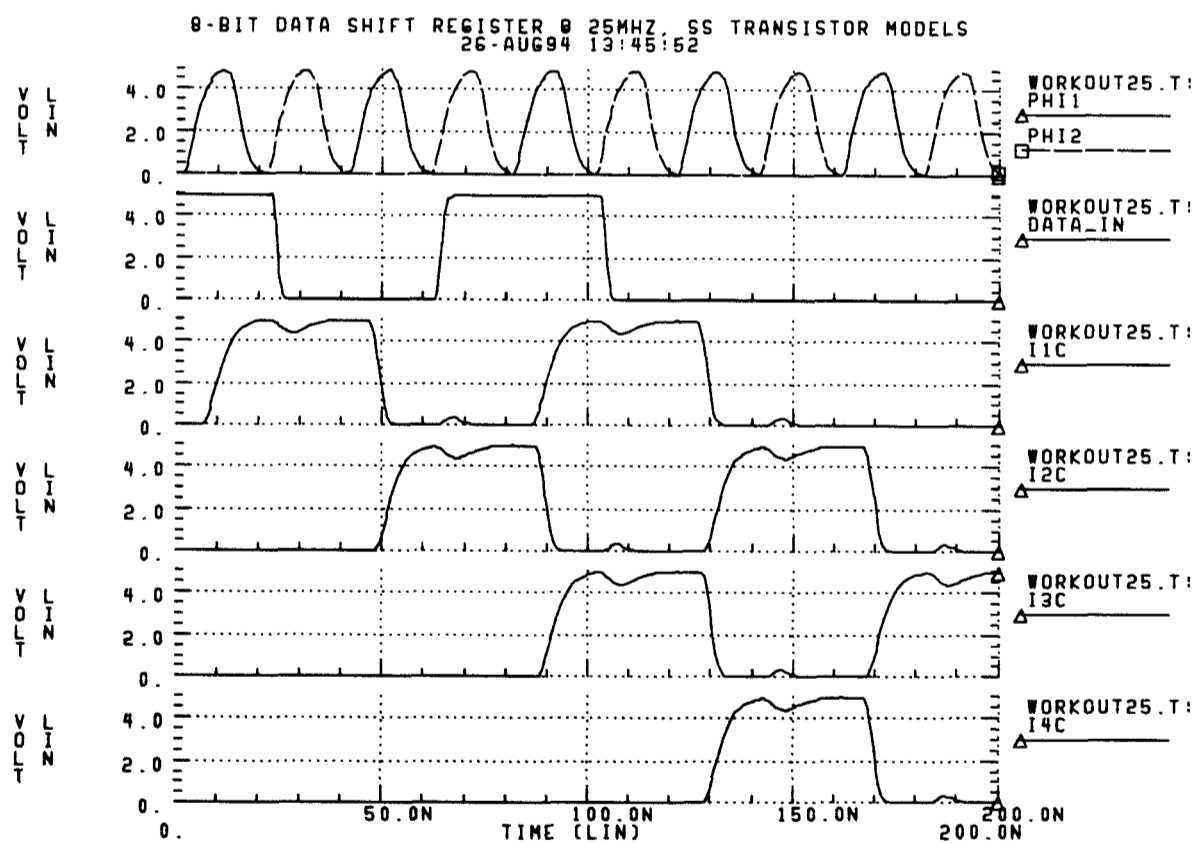
Figure 4-7: Worst-case (SS) HSPICE simulation of a section of the data shift register operating at a frequency of 25 MHz. It shows part of a 1010000 data train propagating along the shift register. The simulation includes the effects of the input pads' circuitry, and the distributed capacitances and resistances of the clock signal lines.

4.5.2 Auxiliary 256-bit register-buffer

When the 256-bit data word is assembled, it is transferred to the register-buffer on the rising edge of the signal LDLAT. Each register-buffer cell comprises a single SRAM latch with two inverters to drive the DATA and /DATA buslines. The inverters have four of the eight clock cycles per column to charge or discharge the buslines, therefore they could be designed to switch quite slowly ($4 \times 62.5 = 250$ ns for a 16 MHz shift register clock frequency) to help limit short-circuit power dissipation.
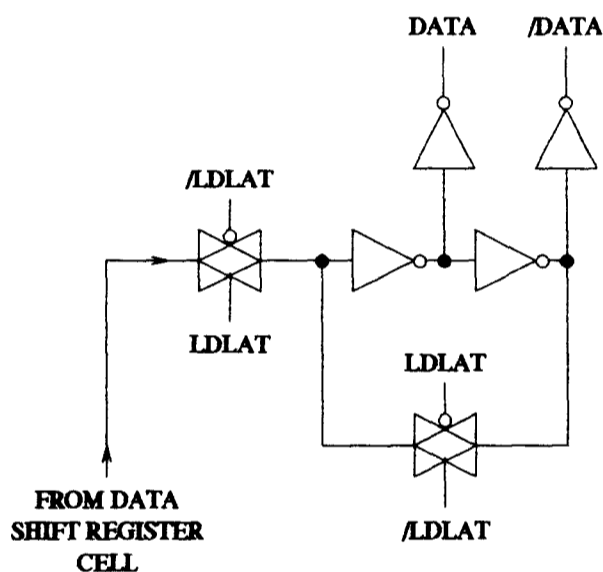
Figure 4–8: Schematic of a single latch-buffer cell.

4.5.3 8-to-256 line column decoder

Each cell in the 8-to-256 line column decoder consists of two 4-input NAND gates which are fed into a 3-input NOR gate, and a inverter buffer to drive an /ENABLE busline (see Fig. 4–9). The other signal to the NOR gate is the control signal /WRITE. The /ENABLE drivers are bigger than the DATA and /DATA drivers as only one busline is charged or discharged at a time. When all eight decoder lines to a cell are high, the /ENABLE busline goes low on the falling edge of /WRITE, and data on all the DATA and /DATA lines is written to the appropriate column of pixel latches.
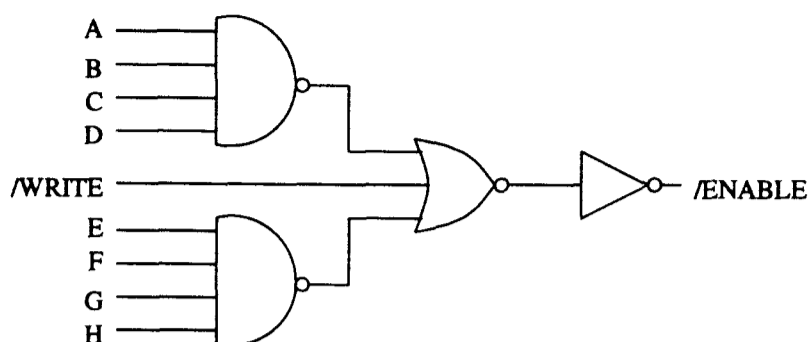
Figure 4–9: Schematic of a single column decoder cell. The /ENABLE line remains HIGH unless all the input lines A→H are HIGH and /WRITE is LOW.

Fig. 4–10 illustrates the timing sequence of the control signals for the data shift register, auxiliary latch-buffer, and column decoder circuits. Note that the sequence is continuous: the PHI's do not have to stop while the other signals are pulsed, a consequence of including the auxiliary latch-buffer.
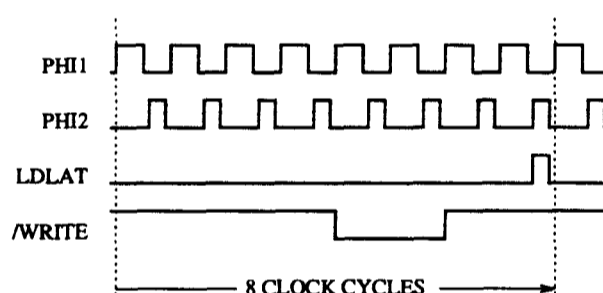
Figure 4–10: Control signal timing sequence to load a single column of the array. For clarity, the data bytes and column decoder byte are not included. The sequence is continuous: the PHI's do not have to stop while the other signals are pulsed, a consequence of including the auxiliary latch-buffer in the addressing circuitry. Note that when /WRITE is pulsed LOW, the data assembled in the previous eight clock cycles is written to the appropriate column.

4.5.4   Busline distribution

Six VDD and four ground pads supply power to the die. The VDD power rail is routed in METAL2 onto the array from both the top and bottom edges, while the ground rail is routed in METAL1 and comes in from both the left and right sides.

The power buslines leading up to the array are 60–80 μm wide to ensure they can handle the transient current spikes without voltage drop or electromigration problems. The arrows in Fig. 4–4 indicate how the DATA, /DATA, /ENABLE, and CK control signals come into a pixel, and so onto the pixel array.

The global clock signal CK is actually split up into eight separate signals, each one supplying a vertical block of 32 × 256 pixels (see Fig. 4–11). If the CK transient current spike was found to cause problems, these could be temporally staggered off-chip so that the spike would occur as eight smaller spikes. The separate clock signals also permit a device with a partially non-functioning array to be properly charge-balanced. From yield considerations, it was expected that some die would have shorts between the pixel clock lines and neighbouring buslines. If CK was a single global signal, and was shorted by a defect to another busline somewhere on the array, it would probably not toggle between 0 V and V_{DD}, and so the FLC layer could not be properly charge balanced. This would lead to gradual degradation of the FLC layer, and the device would eventually become unusable. As CK is routed as eight separate signals, the working sections of the device can be fully charge-balanced while the CK signal for the section where the defect occurs can be disconnected to reduce shorting current. The device can then be used in applications where a fully working array is not required.
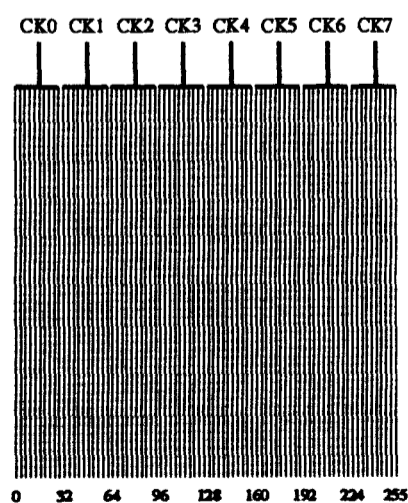
Figure 4–11: Schematic of the split global clock signal CK.

4.5.5 Control signal busline test circuits

To help estimate backplane yield, some control signal busline test circuits to check the integrity of the DATA and /ENABLE buslines are included along the edge of the array. These are discussed in detail in Section 5.2.4.

4.6 SLM drive scheme

For a large array operating at high frame rates, the bit-plane scan time, T_{BPS}, can take up a significant proportion of the frame time. For example, if the data shift register is clocked at 16 MHz, $T_{BPS} = 128$ μs as it takes 2048 clock cycles to scan in a bit-plane. For a 2 kHz frame rate, this corresponds to 25% of the frame time. I will refer to this as the duty cycle factor D, given by

$$D = 2048 \frac{f_{CK}}{f_{DSR}} \tag{4.5}$$

where f_{DSR} is the clocking frequency of the data shift register; and f_{CK} is the SLM frame rate. The time elapsed between addressing the first and last columns can cause problems with charge balancing. However, the inclusion of the XOR-gate at each pixel permits the use of a drive scheme that can fully charge balance and synchronously switch the FLC layer across the whole array (see Fig. 4–12). Assume that pixel A is in the right-most column, pixel Z is in the left-most column, and that both their latches are set to logic high in the F^{th} frame and logic zero in the $(F+1)^{th}$ frame. The pattern is scanned in from right to left with the data shift register operating at 16 MHz.

By toggling the global clock signal, CK, at appropriate times within the frame, pixels with a logic one in their latch switch on and off simultaneously - there is no addressing skew. They do however, receive their charge balancing pulses for different portions of the frame, with pixels in columns other than the right-most actually receiving part of theirs during the bit-plane scan of the next frame.
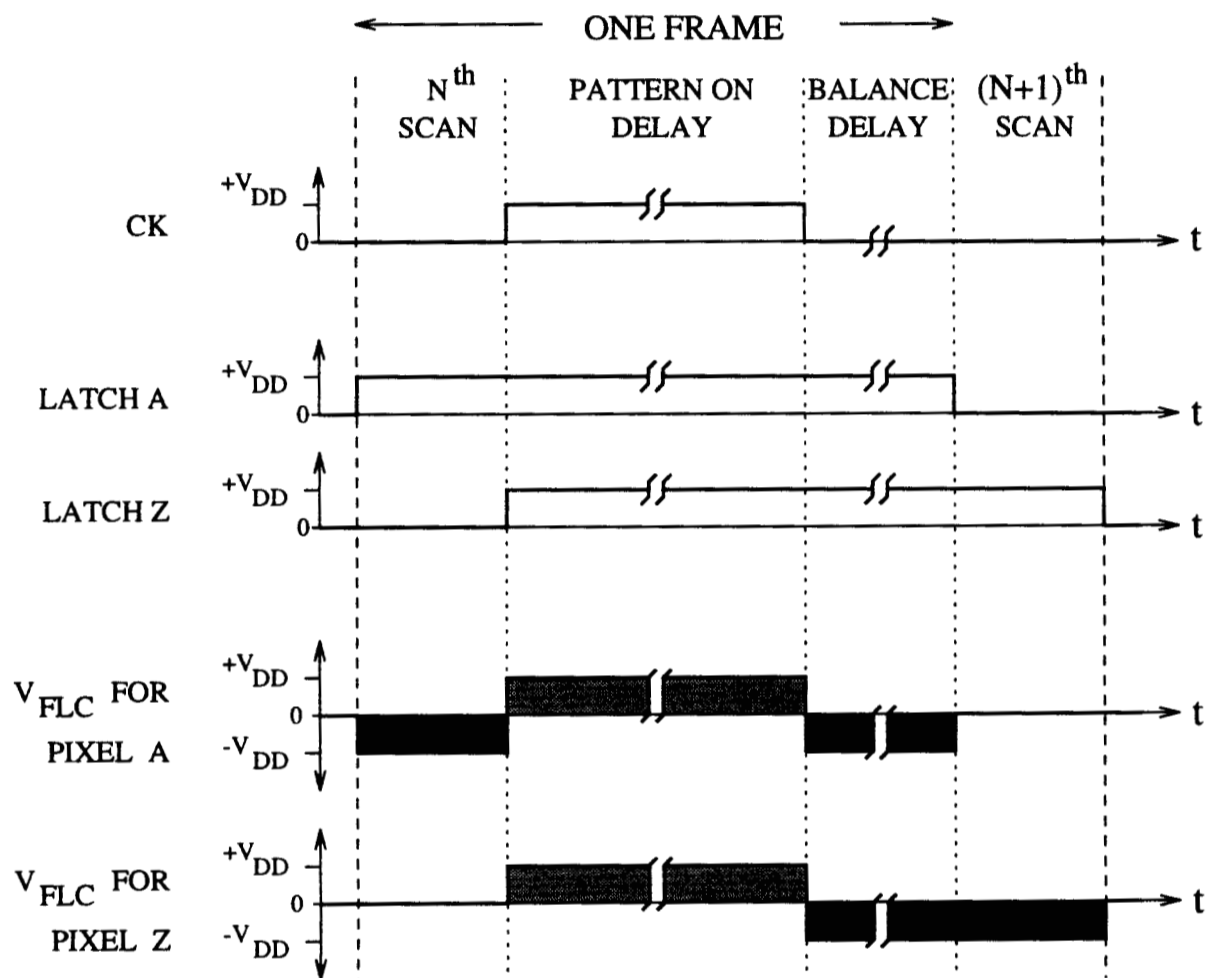
Figure 4–12: Charge balanced SLM drive scheme. Pixel A is in the right-most column and pixel Z is in the left-most column. Assume that both their latches are set to logic high in the F^{th} frame and logic zero in the $(F+1)^{th}$ frame. The array is addressed from right to left. By toggling the global clock signal, CK, at appropriate times within the frame, the FLC above pixels with a logic one in their latch switches on and off simultaneously - there is no skew across the array.

The pattern on delay, T_{PO}, is given by,

$$T_{PO} = T_{BPS} + T_{FLC} + T_{EX} \tag{4.6}$$

and the balance delay, T_B, by,

$$T_B = T_{FLC} + T_{EX} \tag{4.7}$$

where, T_{BPS} is the bit-plane scan time; T_{FLC} is the liquid crystal settling time; and T_{EX} is extra time inserted to give the appropriate frame rate. Notice that a T_{BPS} term must be included in T_{PO} to ensure proper charge balancing. Therefore, from Equations 4.6, 4.7, and adding T_{BPS} for the bit-plane scan, the frame rate f_{CK} is given by

$$f_{CK} = \frac{1}{2(T_{BPS} + T_{FLC} + T_{EX})} \tag{4.8}$$

From this, the maximum frame rate, $f_{CK_{MAX}}$ is given by,

$$f_{CK_{MAX}} = \frac{1}{2(T_{BPS} + T_{FLC})} \tag{4.9}$$

that is, when $T_{EX} = 0$.

4.7 Power dissipation

In normal operation the backplane dissipates dynamic power. The following circuits or switching events were identified as potential contributors to power dissipation:

1. Input pad buffers for the control signals.

2. Input pad buffers for the 32-bit data signals.

3. Data shift register circuitry.

4. DATA and /DATA busline drivers.

5. Column decoder circuitry.

6. Setting/resetting pixels.

7. /ENABLE busline drivers.

8. FLC switching

The amount of power these circuits dissipate depends both on the operating frequency of the data shift registers and on the charge balanced frame rate of the device. Before discussing the pattern dependent contributions from 2, 3, 4, 6, and 8, I will the consider pattern independent dissipation (1, 5, and 7) in the various control signal circuits.

4.7.1 Pattern independent power dissipation

Simple inverters are used to drive all the control signal buslines. The control buslines are all at least 10.24 mm long so their capacitances' significantly load the driver inverters. The dynamic and short-circuit power dissipation components for a simple inverter with a load capacitance C_L, are:

$$P_{DY} = C_L V_{DD}^2 f D \qquad (4.10)$$

$$P_{SC} = I_{mean} V_{DD} D \qquad (4.11)$$

where V_{DD} is the typical voltage swing in a CMOS inverter; and f is the switching frequency. Here I have also included the duty cycle factor D because most of the control signals on the backplane work in a 'burst' mode. For a symmetrical inverter with no load, Veendrick [107] gives I_{mean} as

$$I_{mean} = \frac{1}{12} \frac{\beta}{V_{DD}} (V_{DD} - 2V_T)^3 \frac{\tau}{T} \qquad (4.12)$$

where τ is the rise or fall time of the input signal, and $T = 1/f$. Analytical analysis for an inverter with a capacitive load is much more difficult so, even though it over-estimates I_{mean}, Equation 4.12 is substituted into Equation 4.11 to give

$$P_{SC} = \frac{\beta \tau f D}{12} (V_{DD} - 2V_T)^3 \qquad (4.13)$$

Input	Duty	Dissipation per busline		Number of	Contribution
signal	cycle	P_{DY} (mW)	P_{SC} (mW)	buslines	(mW)
PHI1	D	1.03	0.15	2	2.36
PHI2	D	1.03	0.15	2	2.36
LDLAT	$D/8$	0.13	0.02	2	0.30
/WRITE	$D/8$	0.13	0.02	2	0.30
Dec-A[a]	$D/16$	0.06	0.01	4[b]	0.28
/ENABLE	$D/256$	0.01	0.01	256	2.82
				Total	8.42

[a]Least significant bit of column decoder byte.

[b]To take account of other the column decoder bit buslines.

Table 4–4: Estimated power dissipation in the pattern independent control signal buslines operating at 16 MHz and a frame rate f_{CK} of 2 kHz ($D = 0.256$).

Using Equations 4.10 and 4.13, the PHI1 busline driver dissipates up to 1.03 mW dynamic and 0.15 mW short-circuit power when the following values are used:

$$
\begin{aligned}
C_L &= 7\text{ pF} & V_{DD} &= 6\text{ V} \\
V_T &= 0.62\text{ V} & \beta &= 840\ \mu\text{A}/\text{V}^2 \\
f &= 16\text{ MHz} & \tau &= 5\text{ ns} \\
D &= 0.256
\end{aligned}
$$

The results of similar calculations for the other control signal busline drivers are shown in Table 4-4.

4.7.2 Pattern dependent power dissipation

The power dissipated in scanning in a bit-plane is strongly data dependent. For example, a pattern that requires all the array DATA lines to be changed between each two succeeding columns, such as a fine vertical stripe pattern, will cause significant power to be dissipated in the DATA busline drivers but very little in

the data shift register, while a fine horizontal stripe pattern will dissipate very little in the DATA drivers but a significant amount in the shift register.

Generally, the DATA and /DATA busline drivers make the most significant contributions to power dissipation because there are two buslines running across each row of the pixel array that can be toggled at up to 1/16 the frequency of the data shift register clock.

Using Equation 4.10 as a starting point, the dynamic power dissipation in all the DATA busline circuitry can be written as

$$P_{DY_{DBL}} = \frac{1}{8} N C_{\text{DATA}} V_{DD}^2 f_{\text{DSR}} D G_{COL}$$ (4.14)

where,

N is the number of pixels in a column;

C_{DATA} is the capacitance of a DATA busline;

f_{DSR} is the operating frequency of the data shift register; and

G_{COL} is the transition factor for data across the pixel array, i.e. for a thin vertical stripe pattern $G_{COL} = 1$, while for a thin horizontal stripe $G_{COL} = 0$.

Similarly, the DATA buslines short circuit power dissipation can be written as

$$P_{SC_{DBL}} = \frac{1}{96} N \beta_{\text{DATA}} \tau_{\text{DATA}} f_{\text{DSR}} D G_{COL} (V_{DD} - 2V_T)^3$$ (4.15)

where,

β_{DATA} is the transistor common gain factor in the DATA busline driver; and

τ_{DATA} is the rise/fall time of the DATA busline signal.

From Equations 4.14 and 4.15 there can be up to 24 mW dynamic and 8.4 mW short-circuit power dissipation in the DATA busline drivers when the following values are used:

N	=	256	C_{DATA}	=	5 pF	V_{DD}	=	6 V
V_T	=	0.62 V	f_{DSR}	=	16 MHz	D	=	0.256
β_{DATA}	=	71 μA/V^2	τ_{DATA}	=	100 ns	G_{COL}	=	1.0

However, note that the figures are a worst-case calculation: a fine vertical stripe pattern with $G_{COL} = 1$: McKnight *et al* [67] estimate that $G_{COL} = 0.20$ for

a binary thresholded version of a 'natural' scene, so lower dissipation can be expected in practice. Similar calculations can be performed for the 32-bit data input pads and the data shift registers.

The pixel array also dissipates power in changing the state of the pixel latches and overlying FLC. From Equation 2.5 in Chapter 2, FLC switching dissipates about 0.64 mW for a 1 cm area (the area of the pixel array) switching at 2 kHz, with $V_{DD} = 6$ V. The most significant array contribution comes from the short circuit dissipation in the inverters of the XOR gates $P_{SC_{\text{XOR}}}$, given by

$$P_{SC_{\text{XOR}}} = N^2 \beta_{\text{XOR}} \tau_{\text{CK}} f_{\text{CK}} (V_{DD} - 2V_T)^3 \qquad (4.16)$$

As mentioned in Section 4.4.2, the inverter transistors are quite long and thin ($\beta_{\text{XOR}} = 4.67$ $\mu A/V^2$), but with $\tau_{\text{CK}} = 1$ μs for the CK driver buffers used on the SLM interface, $P_{SC_{\text{XOR}}} = 5.5$ mW assuming values from above for the other variables.

4.7.3 Total power dissipation

Table 4–5 summarises the contributions for a selection of different patterns. The worst-case 'Checkerboard' pattern can dissipate up to 64 mW, when all the pattern independent and dependent contributions are added up. The SBS256 backplane has an area of almost 2 cm^2, so the dissipation is unlikely to increase the backplane temperature by more than a few degrees when it is mounted into a standard pin grid array chip package.

| Pattern | Data transition factors | | Input pads | | DSR[a] | | DATA[b] buslines | | Other | Total Power |
	G_{ROW}	G_{COL}	P_{DY} (mW)	P_{SC} (mW)	P_{DY} (mW)	P_{SC} (mW)	P_{DY} (mW)	P_{SC} (mW)	(mW)	(mW)
All OFF	0.0	0.0	~0	~0	~0	~0	~0	~0	08.42	08.42
All ON	0.0	0.0	~0	~0	~0	~0	~0	~0	14.83	14.83
Horizontal stripe	1.0	0.0	2.36	2.47	12.30	3.10	~0	~0	11.63	31.86
Vertical stripe	0.0	1.0	~0	~0	~0	~0	24.01	8.36	11.63	44.00
Checkerboard	1.0	1.0	2.36	2.47	12.30	3.10	24.01	8.36	11.63	64.23
'Natural' scene	0.2	0.2	0.47	0.49	2.46	0.62	4.82	1.67	09.70	20.23

[a]Data shift register.

[b]And /DATA buslines.

Table 4–5: Estimated power dissipation for a selection of patterns with the device operating at a 2 kHz frame rate and $f_{DSR} = 16$ MHz. The 'Other' column takes into account the dynamic and short circuit power dissipation in switching the pixels and overlying FLC, and the pattern independent (~8.42 mW) dissipation of the other address circuitry.

4.8 Summary

In this chapter I have discussed the backplane design issues of the SBS256 – a FLCOS SLM that incorporates a SRAM latch and a XOR-gate at each pixel. The latch provides robust data storage, while the XOR-gate permits a simple charge balancing scheme to be used. The array addressing scheme was constrained by the need to limit transient current spikes that occur because of the device's parallel architecture. The pixel circuit and addressing circuitry were optimised to ensure the device could be addressed quickly and efficiently, while also ensuring that current spikes and power dissipation would not affect functionality.

Chapter 5

The SBS256 - testing and characterisation

In this chapter I present test and characterisation results for the SBS256 SLM. I begin by describing the two interfaces I designed and built to drive the SLM backplane. Then I present backplane electrical test results. These include current spike measurements, and yield estimates. I follow this by describing post-processing planarisation and device assembly. I then present optical characterisation results including FLC switching time, contrast ratio, and device uniformity. Finally, I show that the device can be used to generate multiplexed grey-scale.

5.1 SLM controller interfaces

I designed and built two computer controlled interfaces to test and drive the SBS256 backplane.

5.1.1 SimpleInt interface

The first 'SimpleInt' interface (see Fig. 5-1), simply buffers and conditions computer generated control and data signals. These signals are generated by an

IBM-compatible personal computer, and passed to the interface via an Amplicon-Liveline CIO-DIO96 card. After conditioning and buffering, the signals are passed on to the backplane. The interface board is quite compact so that it can be mounted easily into an optical system. It has several test-points for monitoring the backplane functionality and also includes components to complete the DATA and /ENABLE busline test circuitry.
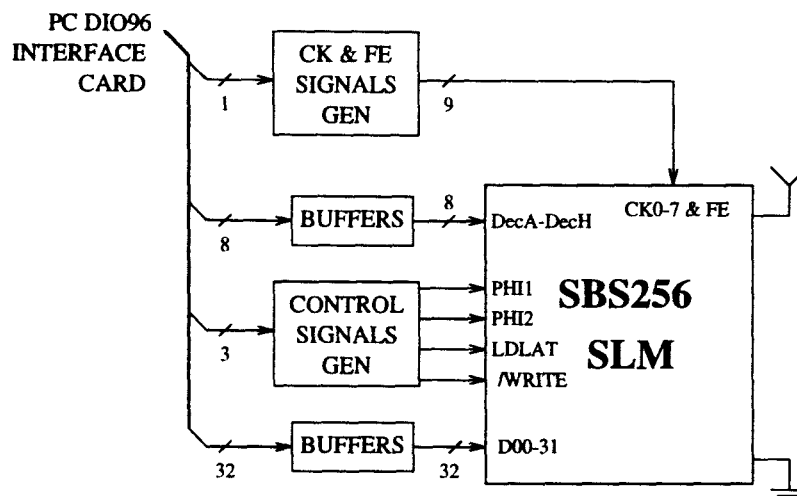
Figure 5–1: Block diagram of the simple interface (SimpleInt).

The control and data signals are synthesised by software written in the C programming language, so the frame scan time is rather slow; it takes 50 ms to scan in a single bit-plane. The frame scan time could probably be reduced by a factor of three or four if the critical output functions were rewritten in an Assembly programming language. This interface is suitable for driving the SLM in applications where high frame rates are not important. It should be noted that, once a bit-plane has been scanned into the pixel array, the global clock and front electrode signals can be toggled (with a 50% duty cycle for charge balancing) at almost any frequency below the maximum switching frequency of the FLC material.

5.1.2 FastInt interface

The second, more sophisticated, 'FastInt' interface (see Fig. 5–2) includes a 16 bit-plane frame store to permit high speed addressing of the backplane for temporal

multiplexing applications. The drive scheme sequence shown in Fig. 4–12, is controlled by a 4-bit algorithmic state machine. The control signals are generated by on-board counters and combinational logic. It was constructed on a wire-wrap prototype board and is capable of operating at up to 24 MHz: that is, data can be scanned onto the backplane at up to 24 MHz on the 32-bit bus. The interface addressing circuitry is designed so that the frame store can be expanded to 32 bit-planes by adding four extra memory chips (each 32k × 8-bits). The state machine can be configured under software control to cycle through a variable number of bit-planes. This is achieved by setting the bit-plane counter start value to the appropriate number. Similarly, the frame rate can be adjusted using the programmable delay counter.
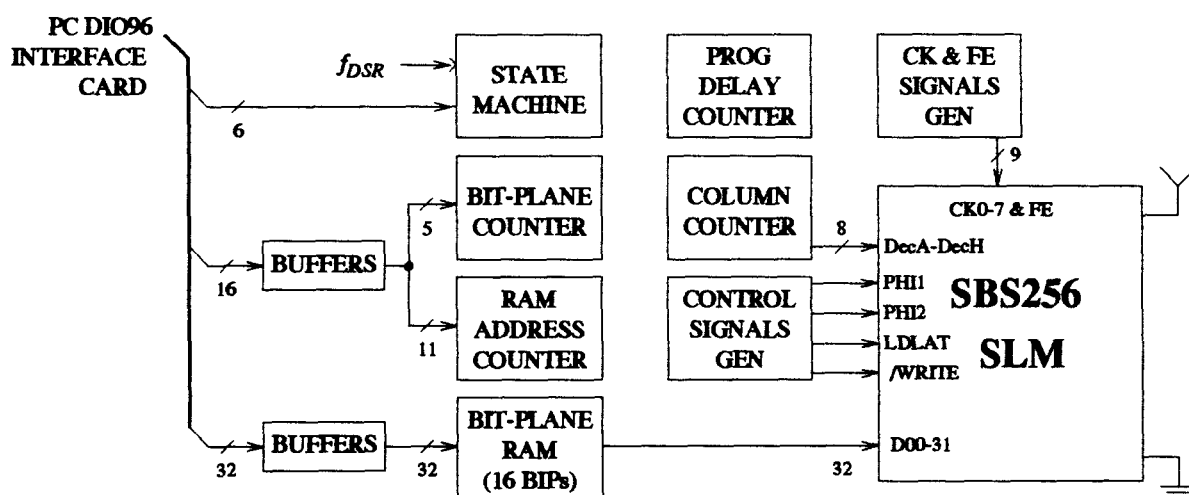
Figure 5–2: Block diagram of the fast interface (FastInt). It includes an on-board RAM store to hold 16 image bit-planes.

The interface board has a socket for a SBS256 device but, as the board is rather bulky, I also constructed a small SLM mounting board to enable the device to be easily mounted in an optical system. The extra capacitances of the connecting cables between the interface and mounting board increase the rise/fall-times of the control and data signals. Capacitive coupling within the cables can also add significant noise to the signals. Therefore, the FastInt is normally operated at 16 MHz (rather than 24 MHz) to ensure the SLM receives relatively 'clean' control and data signals.

5.2 Backplane electrical testing

A batch of ten 4 inch wafers, each with 21 candidate die, was fabricated at the AMS silicon foundry. When the fabricated wafers came back from AMS, wafer 10 was sawn up and three die (10#1–10#3[1]) bonded into 144 pin grid array packages for preliminary electrical testing using the SimpleInt interface.

5.2.1 Preliminary testing

Under quiescent conditions, a fully working die should sink a negligible amount of current from the power supply as it is a digital CMOS design with no current biassing circuitry. A simple power-up electrical test verified that the design does not have a major short-circuit failure mode; dies 10#1 and 10#3 sink 0 mA but 10#2 sinks 11 mA. Probing the CK signals on 10#2 revealed that the interface could not set CK5 to 0 V. When CK5 is set to V_{DD} the current is reduced to 0 mA. This is probably due to a short between CK5 and a neighbouring busline within the pixel array, and illustrates one of the yield limiting effects that were expected with this densely packed backplane design.

To aid with preliminary testing, the following on-chip signals are buffered and routed to output test pads

- Output of the last 8-bit data shift register.

- $DATA_{255}$ busline.

- $/ENABLE_0$ and $/ENABLE_7$ buslines.

- Electrode mirrors of the top-right and bottom-right pixels.

---

[1]The die are labelled W#D, where W is the wafer number and D is the candidate die number.

Probing of these signals on the three test chips confirmed that the functionality of the design was correct.

5.2.2 Current spike measurements

Throughout the design cycle I was concerned about the current spike that would be generated on the transitions of the global clock signal CK. I assumed a worst-case scenario where all the pixels switched simultaneously, so all the current spikes would be superimposed to produce a significant and potentially damaging spike. To minimise the effect, I made the transistors in the XOR-gate as long as space would permit, and split the CK into eight separate signals. The test configuration shown in Fig. 5-3 was used to investigate the current spikes that occur in practice. The resistor R_P (10 Ω) connects the power supply to the SLM backplane[2]. Under quiescent conditions, it has no effect as no current flows. However, when a CK transition occurs, the voltage drop on VDD is proportional to the magnitude of the current spike. This can be observed on an oscilloscope (CRO).
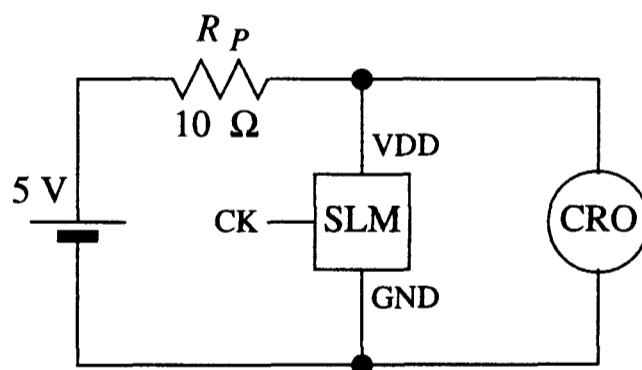
Figure 5–3: Experimental set-up used to investigate the magnitude of the CK transition current spike.

Table 5-1 shows that the current spike for a 64 × 256 block of pixels is about 60% of that expected by simple superposition of the current spike predicted by HSPICE (TT case) for a single pixel. This could be accounted for by a

---

[2]Note that R_P is not present under normal SLM operating conditions.

combination of temporal skew on the CK signal and process variation across the chip: not all the XOR inverters switch at exactly the same time. Although lower than predicted by simple superposition, the spikes are significant and vindicated the architectural and design decisions made during the backplane design cycle.

CK transition	Predicted (mA)	Measured (mA)	Ratio (%)
0 to V_{DD}	65	37.5	58
V_{DD} to 0	98	62.5	64

Table 5–1: Comparison of the predicted and measured CK transition current spikes for a 64 × 256 block of pixels on die 10#1, all with a HIGH in their latches to enable the XOR inverters. The predicted values were based on the superposition of values obtained from HSPICE simulations (TT case) for a single pixel.

Figure 5–4 shows the effect of the current spike flowing through R_P on VDD when all eight CK signals are toggled simultaneously. The voltage drop on VDD during the rising edge of CK corresponds to a 125 mA current spike or 49% of the predicted value, while the voltage drop during the falling edge corresponds to 150 mA or 38% of the predicted value. In this case the effects of CK skew and process variation reduced the peak of the spike even further. These results show that the backplane can be operated safely without the need to stagger the eight CK signals off-chip.

5.2.3 Data shift register operating frequency

HSPICE simulations predicted that the backplane data shift register can be clocked at 25 MHz (see Section 4.5.1). The output of the last data shift register cell is buffered and routed to an output test pad and can be used to verify that data can be clocked through at high frequencies.
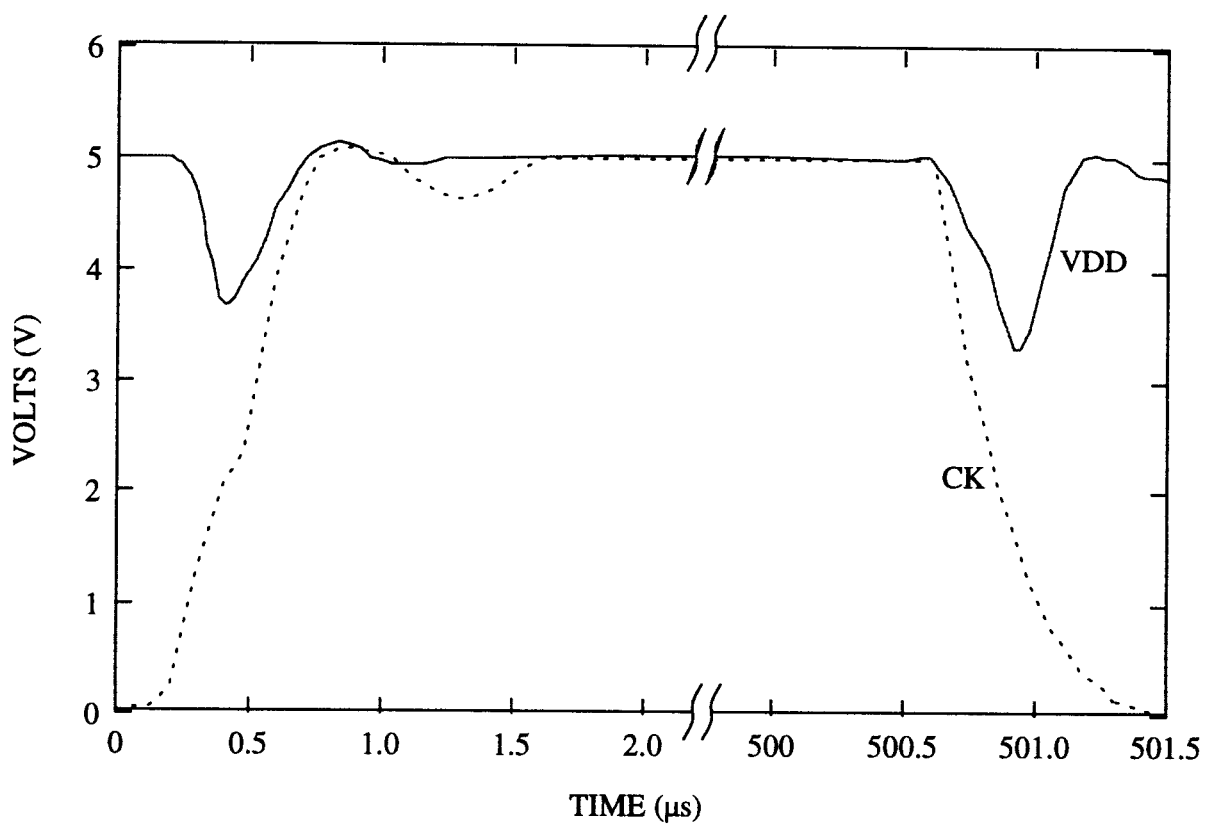
Figure 5–4: The effect of the current spike flowing through R_p on VDD when all eight CK signals were toggled simultaneously. Although the measured spikes were 40–50% lower than predicted by simple superposition of the HSPICE simulations for a single pixel, they were still significant and vindicated the architectural and design decisions made during the backplane design cycle.

The FastInt interface can operate at up to 24 MHz – limited by the programmable logic array ICs used to implement the control sequence state machine, and the fact that it is built on a wire-wrap prototype board. Eight die from different wafers were used to verify that the shift registers could operate at 24 MHz. All devices were then assembled into SLMs and re-tested to check that the extra capacitance associated with the FLC layer did not reduce the shift registers operational frequency: all eight SLMs still operated at 24 MHz. This was expected because all of the control signals and interconnect are routed in METAL1 to ensure there is a relatively thick layer of inter-level oxide (~ 1 μm between METAL1 and METAL2) to reduce the capacitance introduced by adding the front electrode.
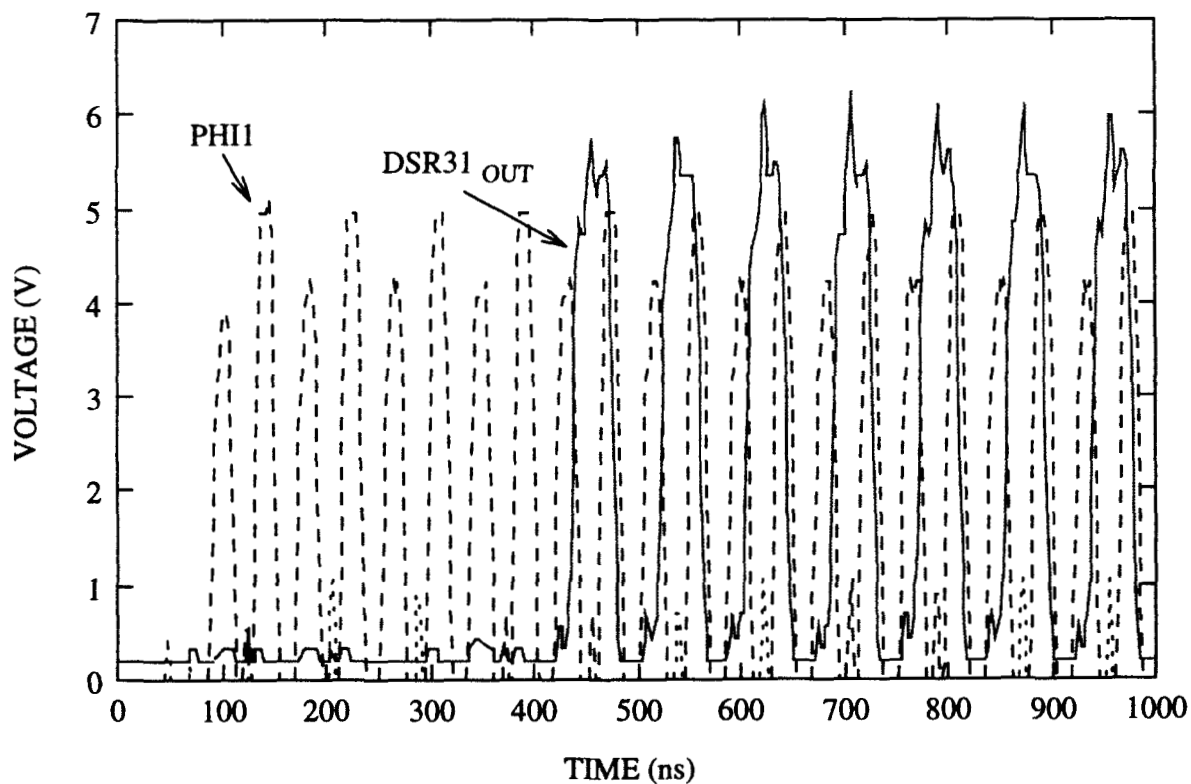
Figure 5–5: Oscilloscope trace of the data shift register test-point on die 5#6, with the non-overlapping clocks operating at 24 MHz. The data input is a fine horizontal stripe pattern which results in the test-point signal toggling at half the clock frequency.

5.2.4   Backplane yield test structures

The yield of the backplane is limited by the densely packed circuitry and interconnection within the pixel array. As mentioned earlier, the spacings of the buslines running across the array are just 0.2 μm greater than the minimum design rules permit, so shorts between them can be expected on some die.

Some of the signals, such as the power rails and global clock signals, are connected directly to external circuitry and so can be probed with a volt meter or oscilloscope to check their integrity, as described above for die 10#3. However, others signals such as the DATAs, /DATAs, and /ENABLEs are generated and distributed entirely on-chip and so are not as readily accessible. To help test the integrity of the DATA and /ENABLE buslines, two simple test structures [104] are included along the edges of the pixel array.

DATA busline test circuits

Fig. 5–6 shows the basic structure of the DATA busline test circuits. Along the edge of the array, each DATA busline is connected to the gates of two transistors – one nMOS, the other pMOS. The drains of all the nMOS transistors are ganged together and routed to the test pad TDB2. This circuit can show whether the DATA lines can be set individually to logic level HIGH. Similarly, the ganged pMOS circuit connected to the TDB1 pad can show whether each line can be set LOW. To illustrate how they function, consider the ganged pMOS structure. A load resistor is connected between the TDB1 pad and 0 V; its resistance is chosen to be comparable with the on resistance of a single pMOS transistor (about 40 kΩ for the transistor size used in the test circuit). If all the DATA lines are set to HIGH, all the pMOS transistors are off, so the test pad voltage V_{TDB1} is pulled to 0 V as no current flows through the load resistor. Then if one of the DATA buslines is switched to LOW while the rest remain HIGH, the single on transistor and the load resistor act as a potential divider so V_{TDB1} rises to about $V_{DD}/2$. If this sequence is repeated for a number of the DATA lines, a trace similar to

that shown in Fig. 5–7(a) results. Any faulty buslines show up as missing pulses in the pulse train.
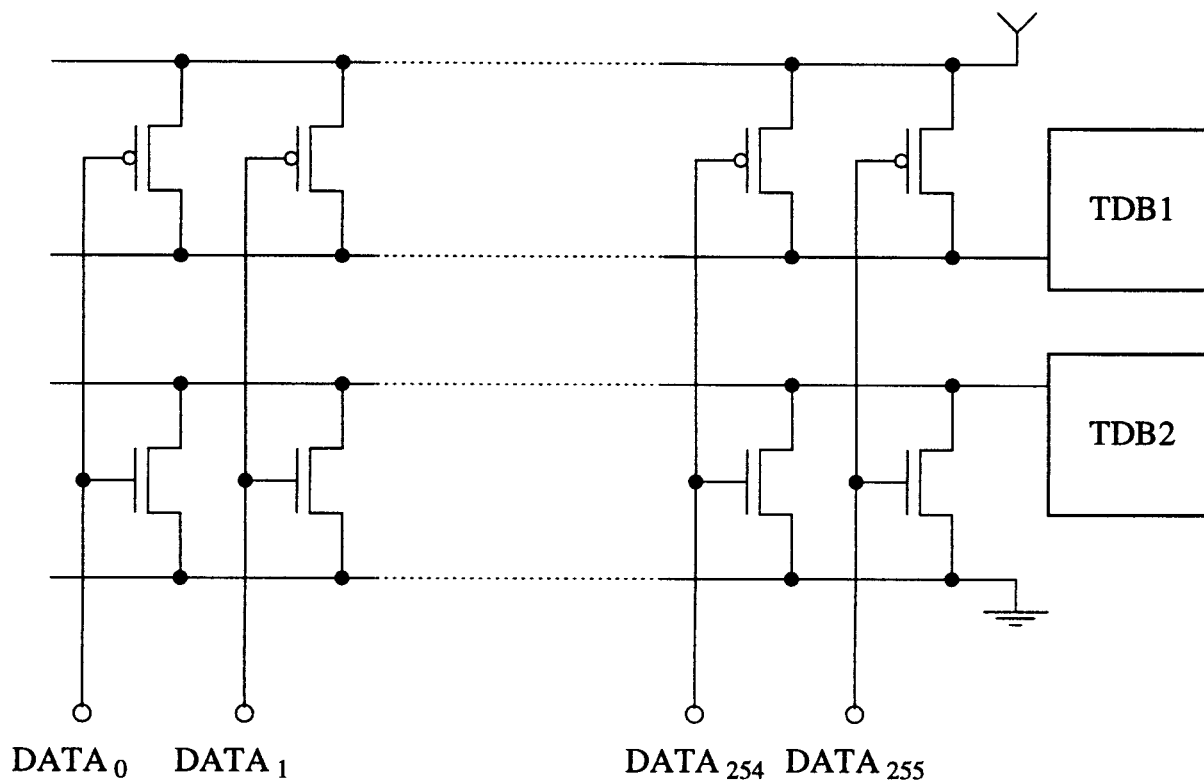
Figure 5–6: Schematic of the DATA busline tester. The circuit can help determine whether a silicon backplane die will fully or partially work when assembled into an SLM.

/ENABLE busline test circuits

A similar structure is used to test the /ENABLE buslines. Fig. 5–7 shows how the circuit can be used to detect which /ENABLE busline is broken even when the defective line switches one of the test transistors permanently on. The pulses correspond to the buslines /ENABLE$_{63}$–/ENABLE$_{32}$ (left to right) on die 8#21. On this die /ENABLE$_{47}$ cannot be set to HIGH because of a defect, so the appropriate test transistor cannot be turned off, therefore the V_{TEB1} signal cannot return to 0 V. However, the sequence of pulses can still be distinguished easily because when any other /ENABLE is set to LOW, there are then two transistors in parallel rather than one, so there is still a significant voltage change on V_{TEB1}.
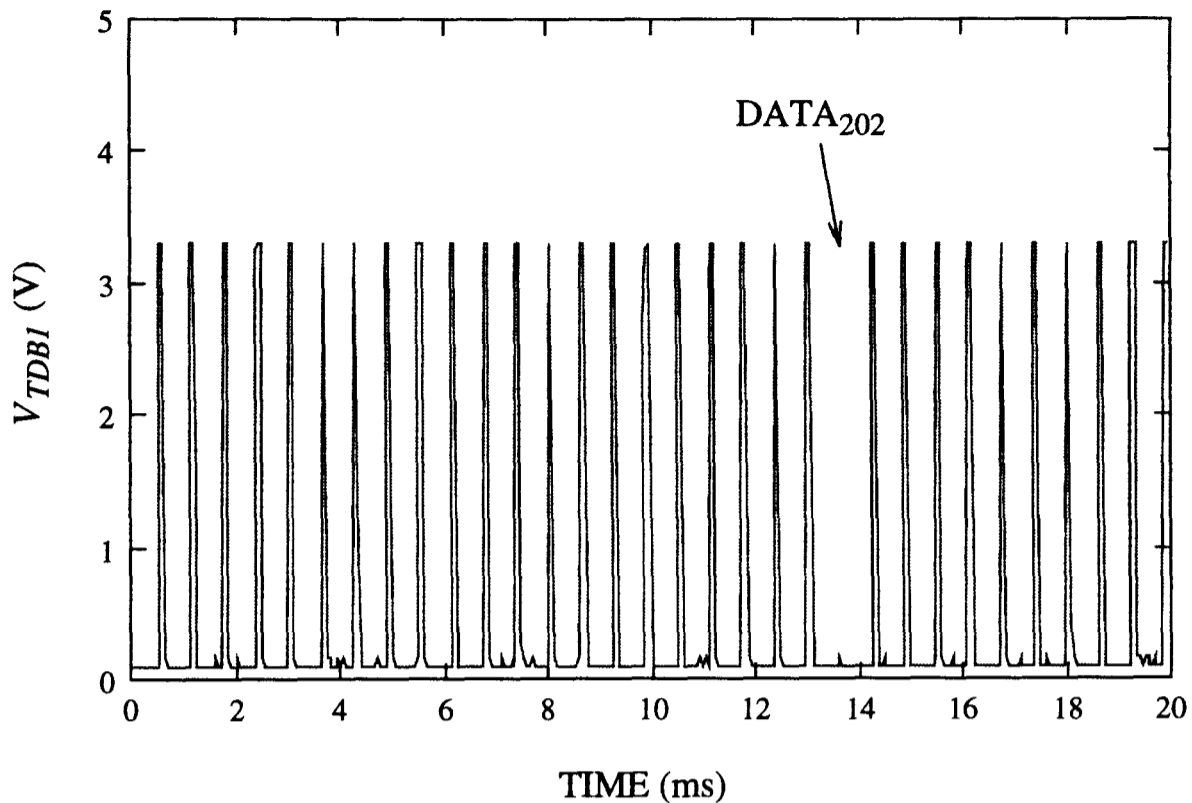
(a)
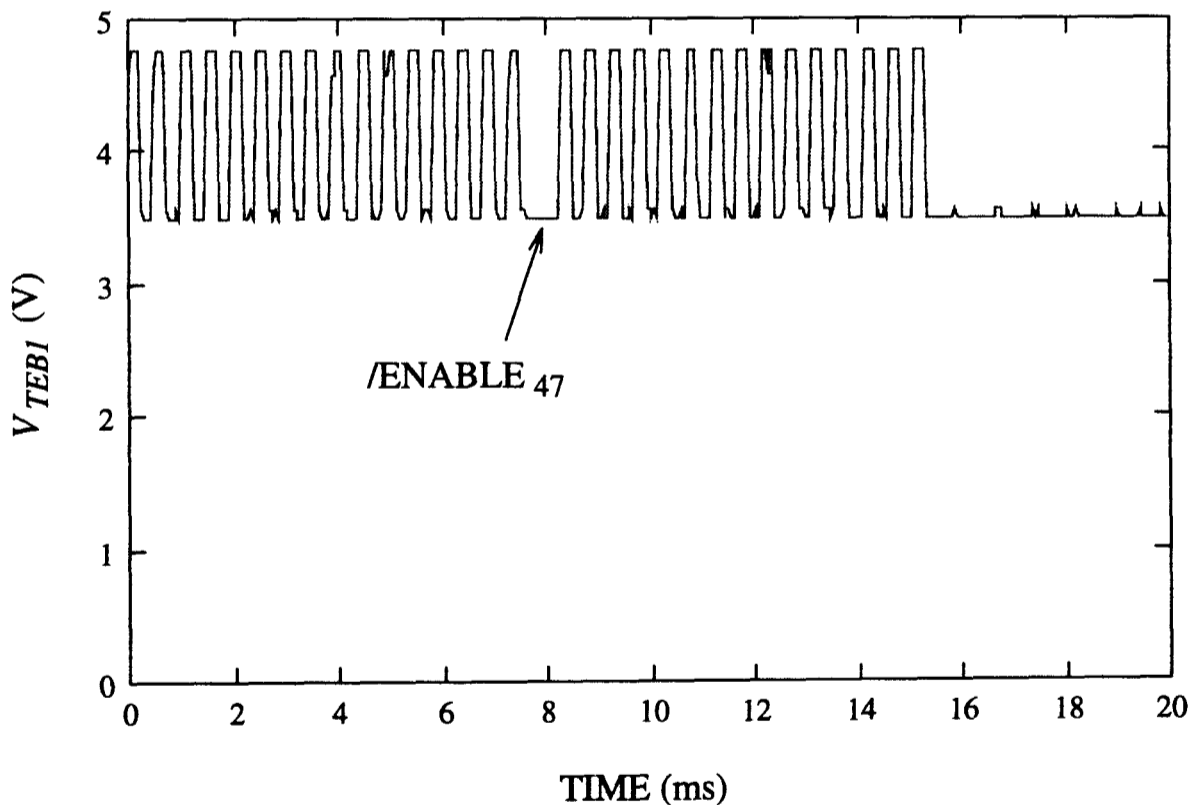
(b)

Figure 5-7: (a) A trace of V_{TDB1} on the DATA busline tester. The pulses correspond to the buslines $DATA_{223}$-$DATA_{192}$ (left to right) on die 8#01. This trace illustrates how the busline tester can be used to detect a faulty DATA busline: note that the pulse for $DATA_{202}$ is missing – this may be caused by a defect in the latch buffer circuitry, or along the busline. (b) A similar trace of V_{TEB1} on the /ENABLE busline tester.

5.2.5 Backplane yield

A custom probe card was purchased to help estimate the yield and identify yield failure mechanisms on the ten wafers. The following tests were performed in sequence on each die:

1. Measurement of quiescent current after power-up and initialisation.

2. Integrity of all eight CK busline signals when toggled.

3. Functionality of the two test point pixels.

4. DATA busline integrity test.

5. /ENABLE busline integrity test.

After power-up, the backplane is quickly initialised by scanning in a test pattern to ensure that all the pixels and signals are in a known state. If the quiescent current is more than 150 mA (900 mW when $V_{DD} = 6$ V), the power supply is immediately switched off so that the interface board is not damaged. The die is recorded as 'failed' and no further testing is performed. If the quiescent current is less than 150 mA the remainder of the tests can be performed.

On the probed wafers, the percentage of fully working backplanes is 44% . In calculating this I have ignored the results of wafer 9 because it was used (and damaged) as a test wafer for post-processing planarisation (see Section 5.3.1). A further 29% of 'partial pass' die could be constructed into working SLMs with perhaps one or two broken lines, or a section of 32×256 pixels that cannot be fully charge balanced. These yield figures are extremely encouraging considering the high packing density required to fit the electrode mirror, circuitry, and signals into a 40×40 μm^2 area. It is worth noting that wafer 4 tends to skew the yield results – without it the full pass yield is 49%, while the partial pass yield is 32%.

I will now use some of the probe tests results of wafer 1 to illustrate some of the failure mechanisms that can occur on the backplane.

Wafer number	Full pass die	Partial pass die	Failed die
1	11	6	4
2	7	10	4
3	13	4	4
4	3	1	17
5	9	6	6
6	12	8	1
7[a]	. . .	. . .	. . .
8	10	7	4
9[b]	0	2	19
10[c]	. . .	. . .	. . .

[a]Supplied to CRL Smectic Technology.

[b]Damaged by planarisation before probe testing.

[c]Not probe tested - used for initial devices.

Table 5-2: Wafer yield results from electrical probe testing. Die that have partially passed can be made into SLMs, but will have a broken line (or lines) or section.

Failure mechanisms of die on wafer 1

Of the 21 potentially working die on wafer 1 (see Fig. 5-8), 11 draw no current and pass all the functionality and busline tests, 6 have some minor faults, and 4 have serious or fatal faults. The failure mechanisms can usually be derived from the observed test signals (in some respects, the faulty die are more interesting because of this!). As an example, here are some observations of die 1#6:

- Die draws 3 mA in quiescent state,

- CK4 cannot be set to 0 V (it settles at 0.24 V),

- /ENABLE busline test signal V_{TEB1} = 3.35 V (rather than 0 V),

- /ENABLE$_{154}$ pulse is missing from busline scan test,

From these, I can infer that the /ENABLE$_{154}$ busline is shorted to the adjacent branch of the CK4 signal somewhere along the column: A similar short occurs on die 1#8, where /ENABLE$_{72}$ is shorted to CK2 (die draws 27 mA). If this die was assembled into an SLM, the 32 × 256 block of pixels associated with the CK2 signal would not be fully charge balanced as V_{CK2} cannot be set below 2.6 V. In use, the FLC over this block would gradually degrade but the rest of the device would be unaffected, illustrating the advantage of splitting CK into eight separate signals. The same would occur for die 1#6 but the degradation would probably take longer because CK4 can be set to 0.24 V.

On die 1#4, the /ENABLE$_{171-173}$ signals do not function but the adjacent clock signal (CK5) is unaffected. In this case, the fault is caused by a 'puddle' of metal on the METAL1 layer stretching across the appropriate three decoder cells. Die 1#4 also illustrates another failure mechanism – V_{CK7} cannot be pulled lower than 4.5 V without overloading the driver interface, but no /ENABLEs are affected, therefore CK7 is probably shorted to at least one electrode mirror somewhere in the appropriate block of 32 × 256 pixels. As a final example, die 1#14 draws 1 mA and, from the DATA busline testers, DATA$_{41}$ does not function
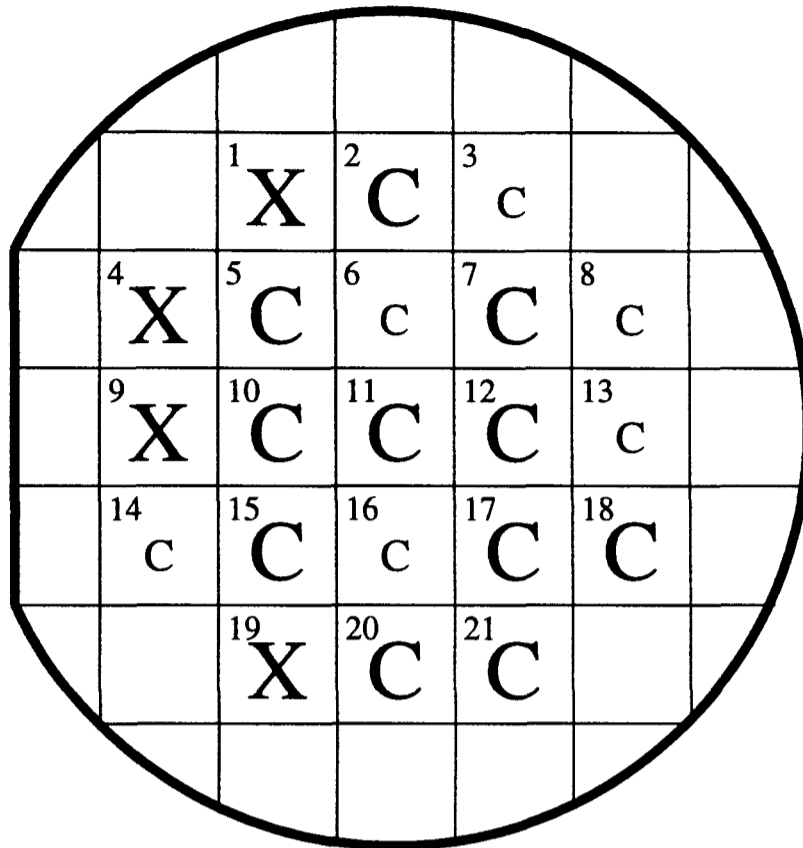
Figure 5–8: Probe test yield of wafer 1 where C = full pass (11 die), c = partial pass (6), and X = fail (4).

properly. If assembled into an SLM, this die would have one broken row near the top of the array, and so could be used for applications where a fully working SLM is not essential.

5.3 Post-processing planarisation and SLM assembly

The techniques described in this section were performed by various people in the Applied Optics Group, in the Physics department and the Silicon Research Group in the Electrical Engineering department, at the University of Edinburgh. They are acknowledged where appropriate.

5.3.1 Wafer planarisation

The fill factor and quality of the electrode mirrors of the backplanes can be substantially improved with post-processing planarisation [79]. It is not appropriate to discuss the post-processing procedures in great detail, as the work is carried out by A. O'Hara, a research fellow with the Applied Optics Group. Briefly, a thick layer (4 μm) of SiO_2 is deposited onto a wafer by an Oxford Plasma Technology AMR electron cyclotron resonance plasma-enhanced chemical vapour deposition system. The oxide layer retains the profile of the underlying circuitry, with surface variations of up to 2 μm. The wafer is then polished by a Logitech PS2000 polishing system until surface variation of the polished surface is reduced to less than 10 nm. This process is referred to as chemical-mechanical polishing (CMP). High fill-factor (84%) METAL3 electrode mirrors are then added using standard MOS procedures[3]; photolithography, oxide etch, metal deposition, and metal etch. The wafer is not sintered on completion to ensure that mirror quality is not reduced by stress-relief hillock formation inherent in the sintering process.

---

[3]These post-processing procedures are performed in the fabrication facilities of the Silicon Research Group.

Planarisation yield results

At the time of writing, post-processing planarisation is progressing rapidly but is still at the experimental stage. As part of the on-going development by the Applied Optics Group, two SBS256 wafers were processed. Unfortunately, probe testing of the first planarised wafer, wafer 9, revealed that all the die sink substantial currents (remember a die should sink negligible current in its quiescent state), with only two that could be made into partially functioning SLMs. Note that this wafer was not probe tested before planarisation because the post-processing commenced before the probe card had been purchased. It could have been a low yield wafer to start with, but the failure mechanisms on it are different from those seen on wafer 3 (the low yield wafer), so the low yield is probably a consequence of post-processing.

Planarisation of wafer 2 was more successful. Table 5-3 shows the probe test yield results before and after post-processing. Although two of the fully working die have been damaged slightly, they can still be made into partially working SLMs. However, the METAL3 quality on this wafer is very poor – the layer is covered in what look like water marks. These are a result of inadequate cleaning after the CMP stage.

	Before	After
Full pass	7	5
Partial pass	10	11
Fail	4	5

Table 5-3: Yield results for wafer 2 before, and after, post-processing planarisation.

The photomicrographs in Fig. 5-9 show the effect of planarisation on the pixel fill factor. Although this technique is in the early stages of development, it promises to improve significantly, the quality of SLM devices. Many of the

procedures must be characterised and optimised, but these initial results for the SBS256 backplane are very encouraging.

5.3.2   SLM assembly

After a wafer has been probe tested, it is cut up into individual die, and some of the candidates glued into PGA packages and bonded by A. Ruthven of the Silicon Research Group. The chips are then assembled into SLMs by J. Gourlay and A. J. Stevens of the Applied Optics Group. The full details of the assembly procedures are described in Ref. [31]; only a brief summary is given here.

The assembly procedure is performed in the Applied Optics Group LC clean room. All equipment and substrates are meticulously cleaned to minimise contamination. A 12 mm × 12 mm × 1.1 mm block of optically-flat, Indium Tin Oxide (ITO) coated glass is used as the cover glass. A thin aluminium film is deposited onto one of the thin faces and onto the edge of the ITO-coated face. The film serves as a connection between the conducting ITO film and front electrode bonding wire (see Fig. 5–10). SiO_x alignment layers are evaporated onto the cover glass. The SiO_x is obliquely evaporated at an angle of 30° to the substrate. This is commonly referred to as medium-angle deposition. This structure gives relatively strong surface anchoring of the FLC molecules. Four spots of glue, impregnated with 2 μm spacer balls, are placed at the corners of the cover glass. The cover glass is then placed carefully atop the active area of the backplane chip, and the glue is cured with an ultra-violet lamp. The front electrode bonding wire is then glued to the aluminium film on the edge of the cover glass. Finally, the device is filled with Merck SCE13 FLC material [59] under vacuum at an elevated temperature of 120 °C, and then slowly cooled back down to room temperature.
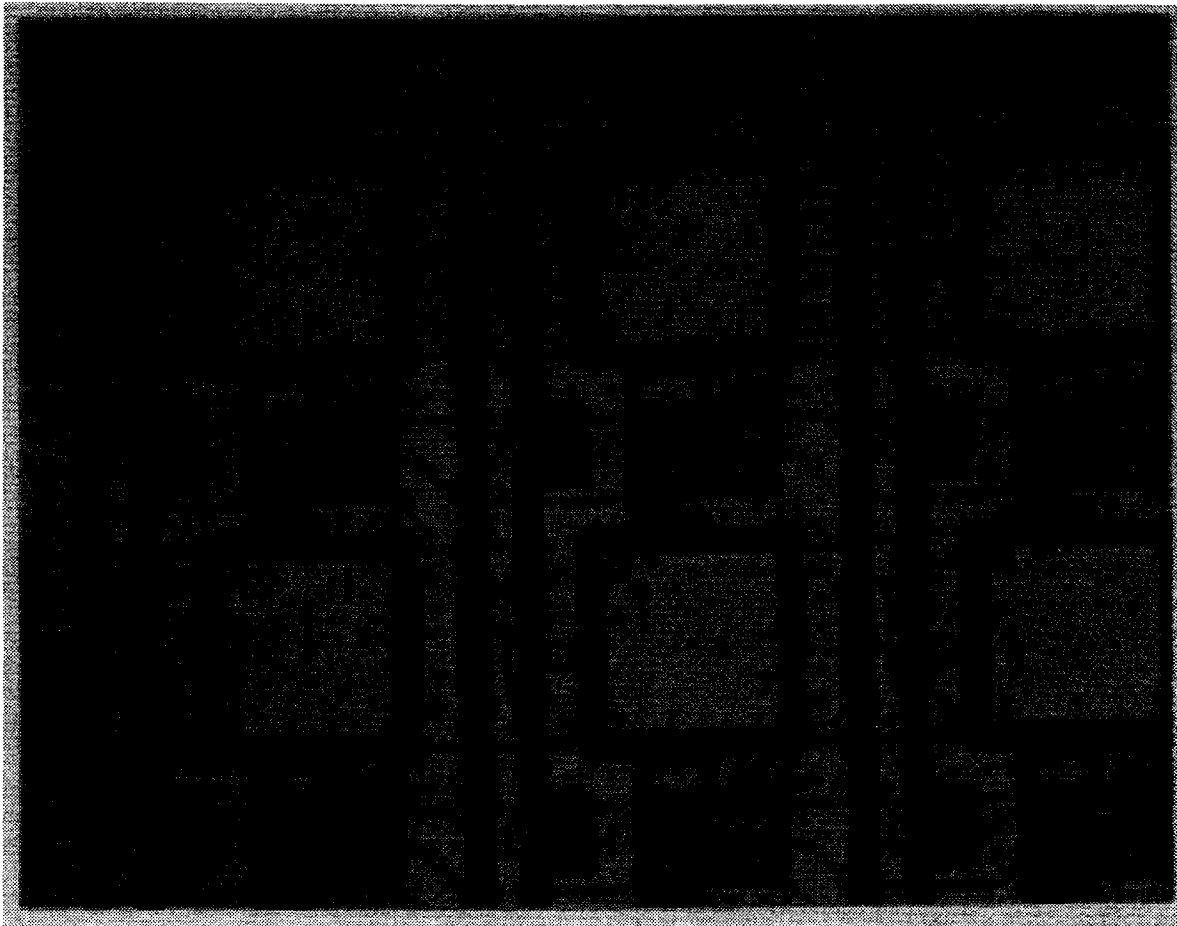
(a)
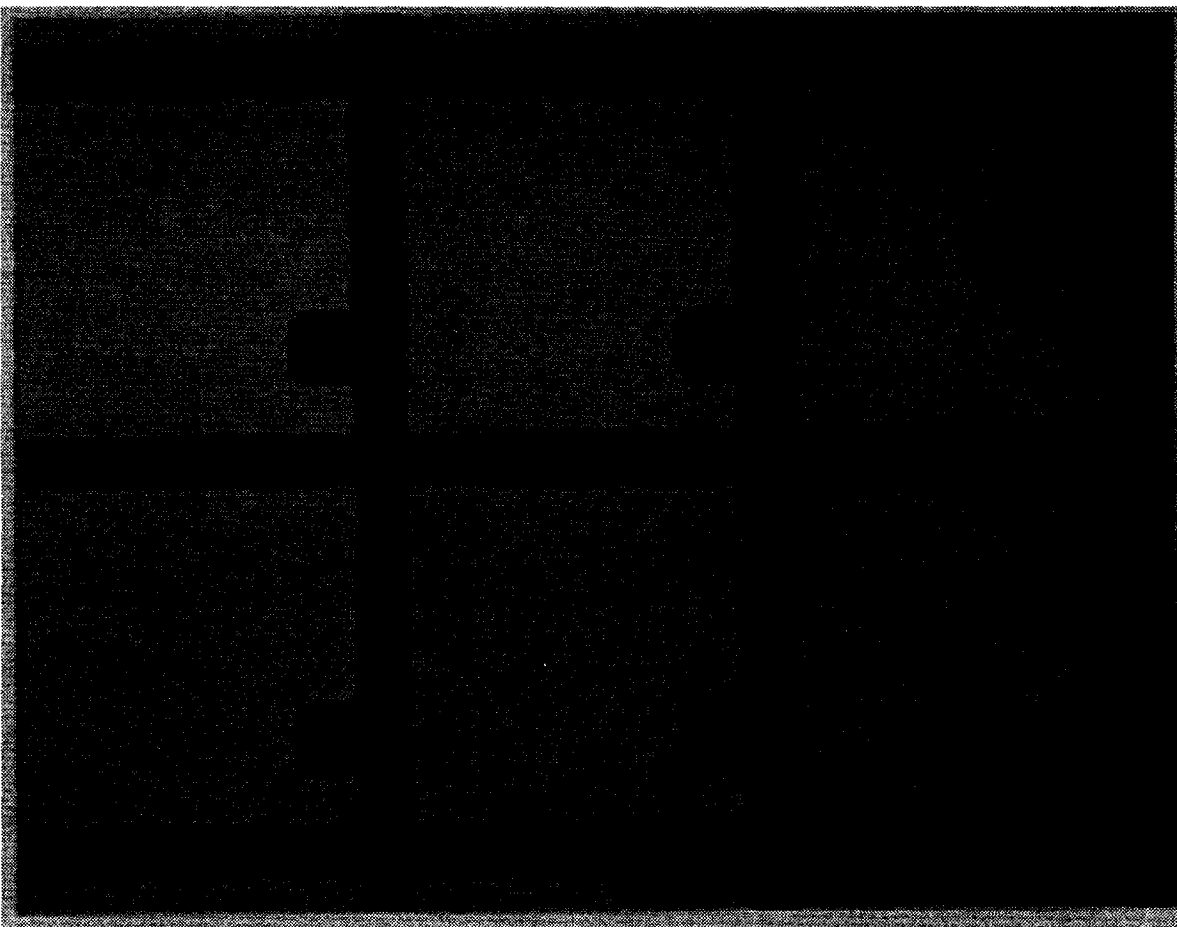
(b)

Figure 5–9: Photomicrographs of a group of six pixels (a) without, and (b) with post-processing planarisation on die 2#2. For unplanarised devices the flat fill factor is 23%, while the for planarised devices it is 84%. Courtesy of A. O'Hara.

(a)
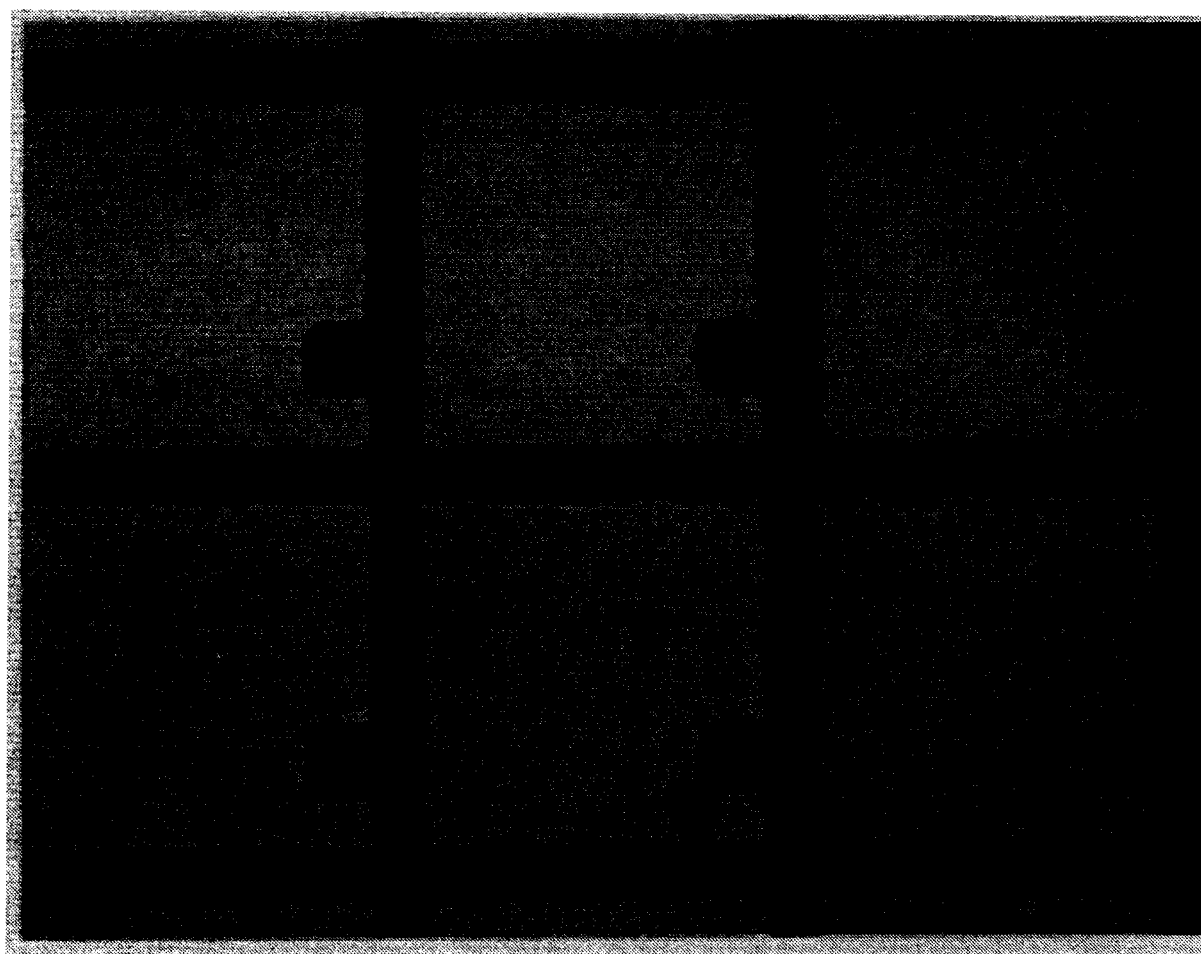
(b)

Figure 5–9: Photomicrographs of a group of six pixels (a) without, and (b) with post-processing planarisation on die 2#2. For unplanarised devices the flat fill factor is 23%, while the for planarised devices it is 84%. Courtesy of A. O'Hara.
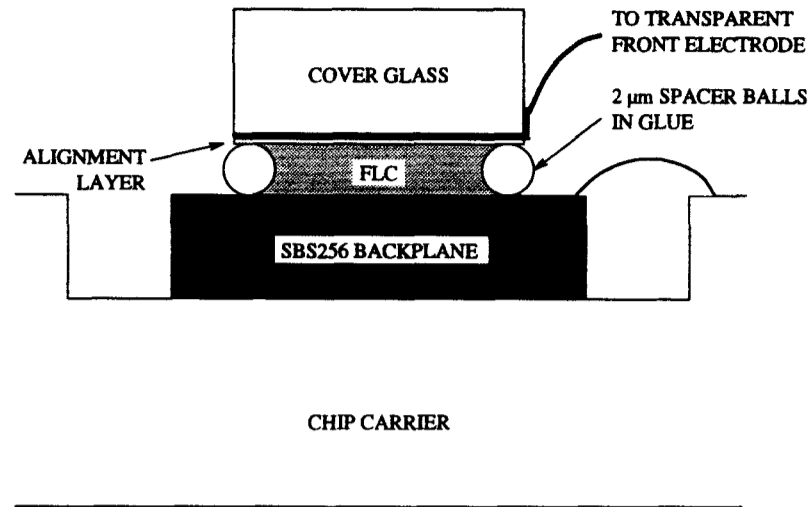
Figure 5–10: Cross-section of the SBS256 SLM.

5.4 Viewing systems

Here I describe two SLM viewing systems. The first system is similar to the configuration used in a polarising microscope. Both can use an incoherent or coherent light source; however, the second is more efficient if a coherent source is used.

5.4.1 Using a non-polarising beam-splitter and two polarisers

The first system , shown in Fig. 5–11(a), uses a non-polarising beam-splitter and two crossed polarisers (P1 and P2). The light source is partially collimated by the lens and is then linearly polarised by P1, which is oriented with its polarisation axis vertical (aligned normal to the page for plan view given in Fig. 5–11(a)). The non-polarising beam-splitter transmits 50% of the light onto the SLM, while the other 50% is lost by reflection. If the SLM and P1 are aligned correctly, the light reflected from ON pixel electrode mirrors has its polarisation rotated by 90°, while that reflected from OFF pixels is unaffected. The reflected light is split again by the beam-splitter and 'analysed' by P2 which is oriented with

its polarisation axis horizontal. Light from ON pixels is transmitted through P2 because the FLC has rotated (the light's) polarisation, while light from OFF pixels is blocked. With this configuration, the output beam is about 10% the intensity of the original beam because of losses in the polarisers but mainly from the two passes through the beam-splitter.
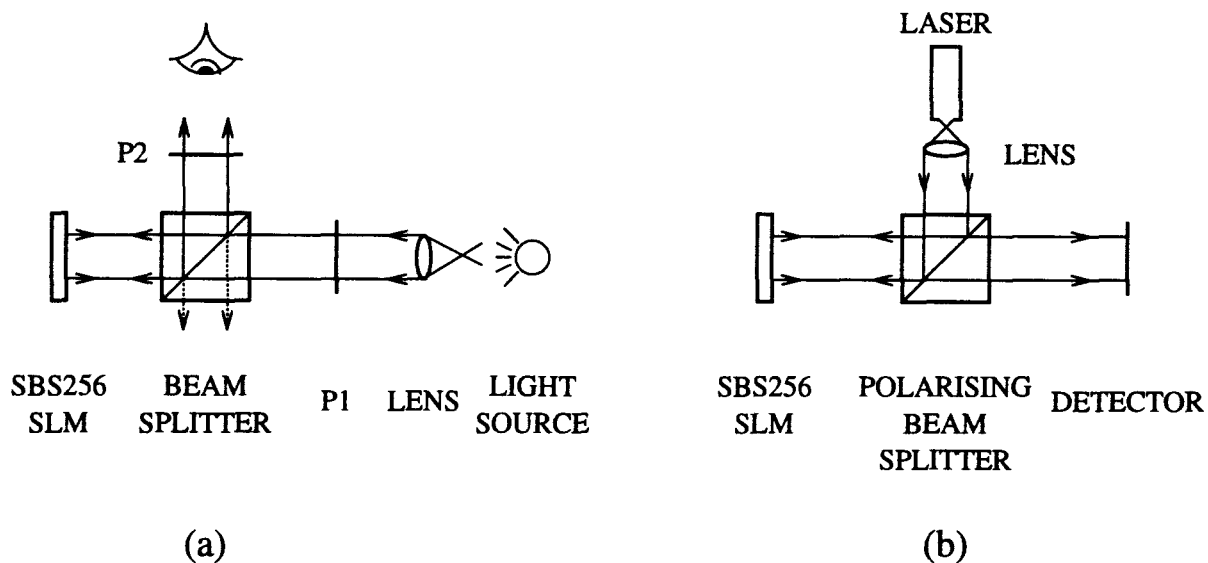
Figure 5–11: Two SLM viewing configurations. (a) Uses a non-polarising beam-splitter and two polarisers, while (b), uses a polarising beam splitter.

5.4.2   Using a polarising beam-splitter

The second configuration is illustrated in Fig. 5–11(b), Here a laser that generates linearly polarised light is used. Its polarisation is set to vertical (out of the page) so that the polarising beam-splitter (PBS), directs almost all the light (> 98% for a wavelength-specific PBS) onto the SLM. If the SLM is oriented correctly, the light reflected by ON pixels is rotated by 90°, while that reflected by OFF pixels is not rotated. The PBS then splits the light according to polarisation: horizontally polarised light (ON pixels), is transmitted straight through to the detector system[4], while the vertically polarised light (OFF pixels) is reflected back

---

[4]Not a human eye when a laser is used !

into the laser. This configuration can have a much higher throughput ($\sim$ 90% ignoring SLM losses) than the non-polarising beam-splitter system, and is used in the neural-backplane characterisation system described in Section 6.7.4.

5.5   Optical characterisation

Once assembled and re-tested to ensure that cell assembly has not damaged the backplane, the device can be viewed under a polarising microscope. The microscope polarisers are set to give a good OFF state when the device is not being addressed and then, while it is being addressed, they are adjusted to give optimum contrast. Fig. 5–12 shows a test pattern on a section of a working unplanarised device (8#2). By toggling CK very slowly ($\sim$ 1 Hz), the optical response can be seen to match the truth table given in Table 4–2. There is no visible difference between a pixel with a LOW stored in its latch ($V_{FLC} = 0$ V), and a pixel with a HIGH stored and CK set LOW (i.e. $V_{FLC} = -V_{DD}$ during the charge balancing section of the drive sequence). These observations agree with results obtained from the 16 $\times$ 16 device.

5.5.1   Data storage

The SRAM pixel provides robust data storage under high-read beam intensities. To verify this qualitively, a pattern was scanned onto an unplanarised device, and CK was toggled at 1 kHz. The device was viewed under the polarising microscope, and then placed in front of a 250 W mercury-vapour light source for just over two hours. It was positioned far away enough so that the infra-red radiation being emitted by the light source did not overheat it. The device was then viewed under the microscope to verify that the pattern was unchanged.
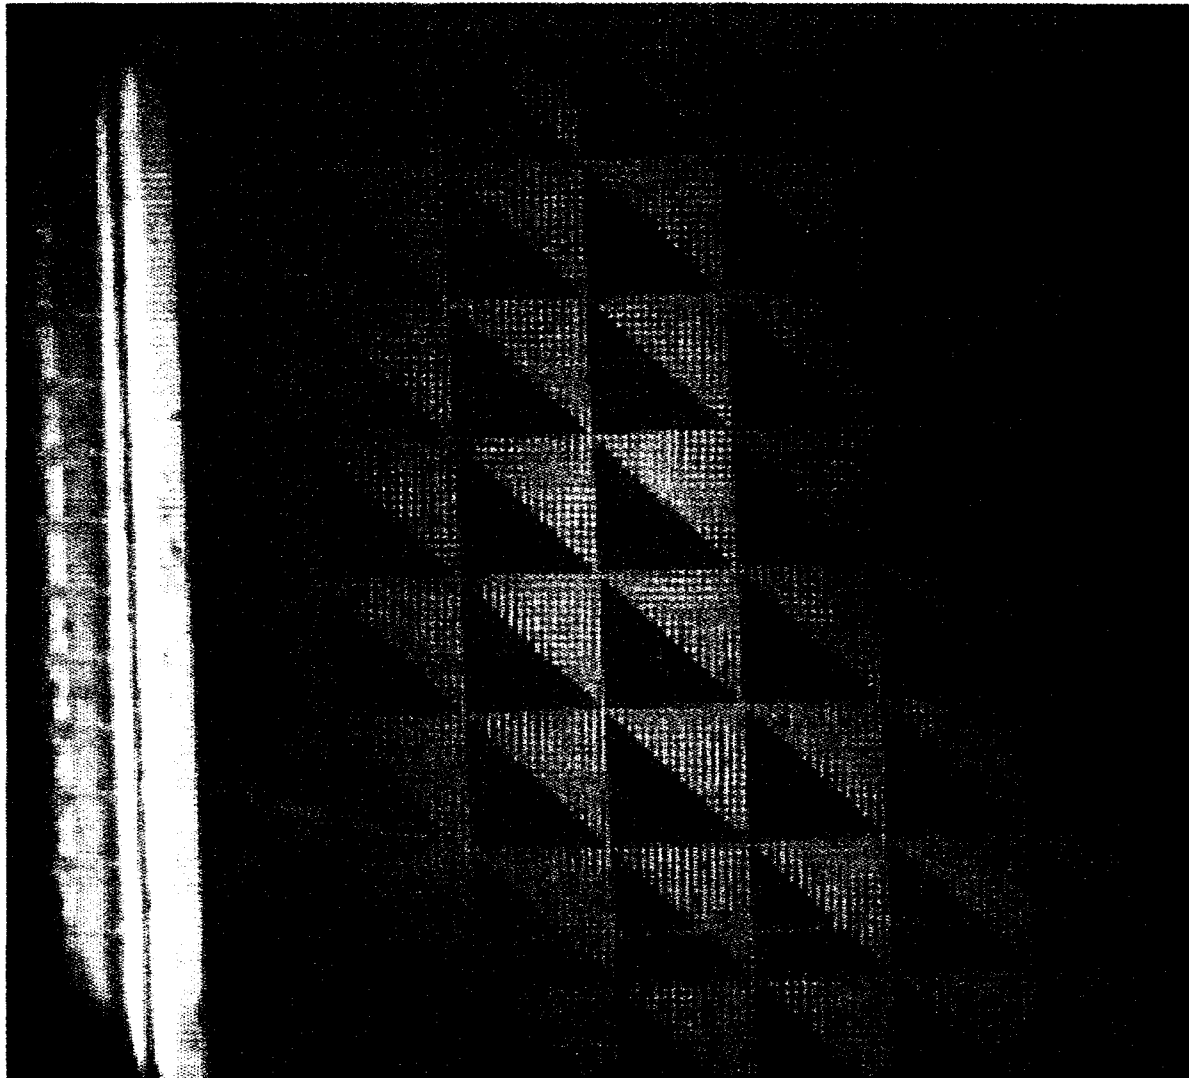
Figure 5–12: Test pattern on a section of SLM 8#2.

5.5.2 Planarisation

Fig. 5–13 shows the improved fill-factor and quality of a planarised device compared to an unplanarised one.

5.5.3 FLC switching time

The optical system described in Section 5.4.1 was used to determine the FLC switching time. A neutral density filter reduced the light falling onto the photodiode detector to ensure its integral amplifier was not saturated. An 'all ON' pattern was loaded onto the SLM, and CK was toggled at 1 kHz. The rise time $T_{FLC_{RISE}} = 150\pm5$ μs, while the fall time $T_{FLC_{FALL}} = 55\pm5$ μs (see Fig. 5–14). Mao and Johnson [62] have also reported asymmetrical switching times for their neural-backplane FLCOS SLM. The asymmetry may be due to the single-sided FLC alignment layer.

The potential frame rate of the device is limited by $T_{FLC_{RISE}}$. Assuming data is loaded onto the backplane at 16 MHz ($T_{BPS} = 128$ μs), and using Equation 4.9, the maximum frame rate is 1.8 kHz. Increasing the data shift register frequency 24 MHz gives a frame rate of 2.1 kHz. As with post-processing planarisation, device assembly techniques are still at the experimental stage, so with better alignment layers and a thinner (~ 1 μm) cell gap, the FLC switching time could be significantly reduced, which in turn, would help increase the device frame rate.

Some devices displayed an inverse of the desired pattern during the charge balancing section of the drive sequence rather than the expected 'all OFF' pattern. They also exhibited a very slow optical response ($T_{FLC} > 1000$ μs). These effects can be attributed to a faulty connection between the ITO front electrode and the evaporated aluminium along the edge of the cover glass: the front electrode voltage is undefined and tends to float to a value dependent on the pattern loaded onto the underlying electrode mirrors.
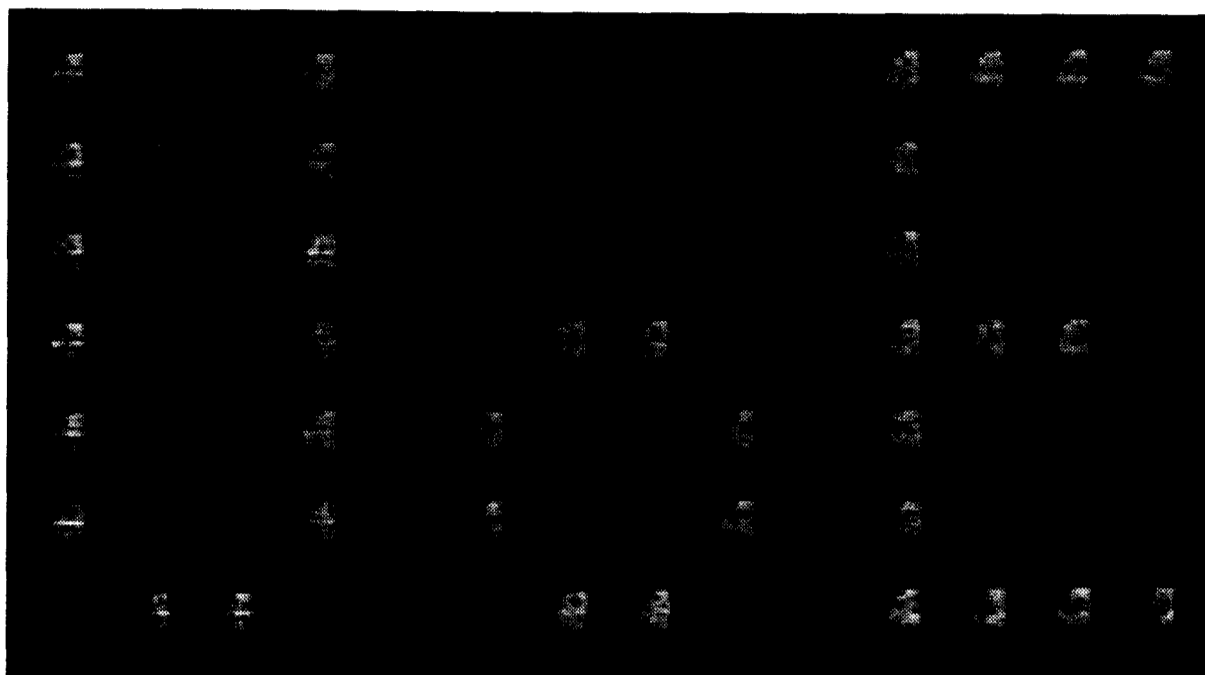
(a)
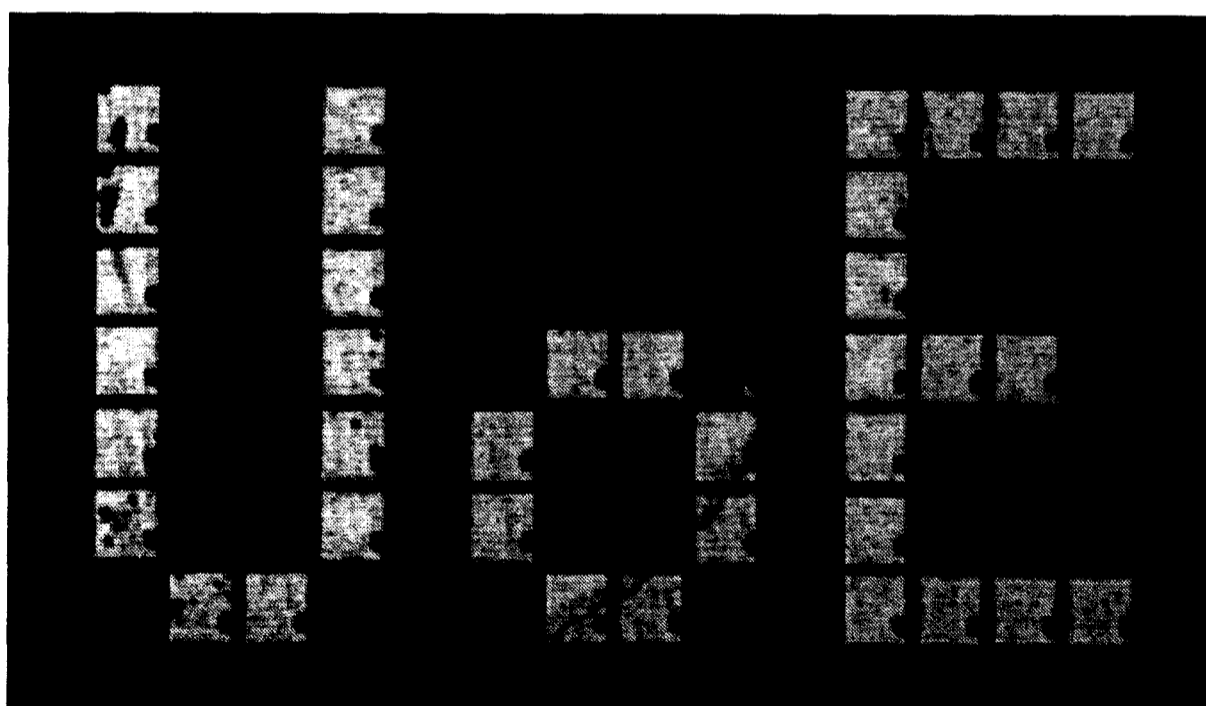
(b)

Figure 5–13: A pattern on a section of an SLM (a) without, and (b) with post-processing planarisation.
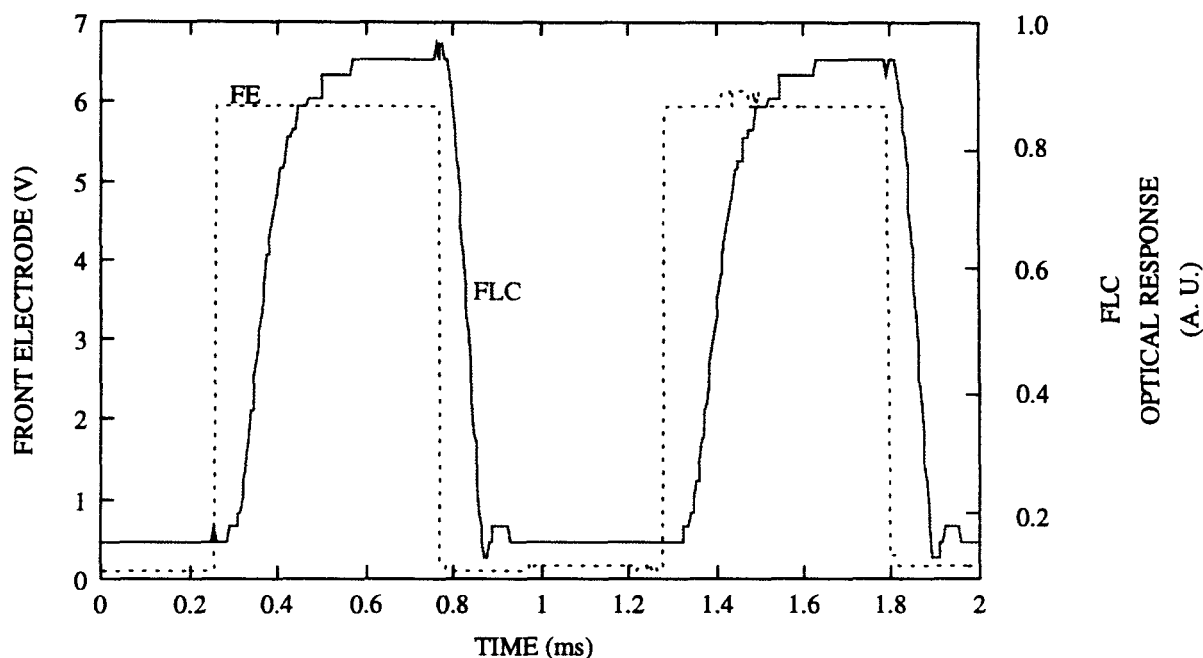
Figure 5–14: Optical response of the FLC. Here the global CK and FE signals are toggled at 1 kHz. The rise time $T_{FLC_{RISE}} = 150\pm5$ μs, while the fall time $T_{FLC_{FALL}} = 55\pm5$ μs.

5.5.4   Contrast ratio

A charge coupled device (CCD) camera connected to a frame-store was considered convenient for investigating the SLM on/off contrast ratio. The camera was attached to the polarising microscope and the microscope was set to medium magnification. The microscope light source was adjusted to ensure the reflected light did not saturate the camera. A pattern was scanned onto the SLM and CK was toggled very slowly (~0.2 Hz). The image was grabbed by the frame-store during the 'pattern-on' section of the SLM drive scheme. A section of the image was then used to estimate the contrast ratio (see Fig. 5–15). As can be seen, there are variations across each pixel. These are caused by defects in the FLC layer, so to estimate the contrast ratio, average intensity values for area of an 'on' and 'off' pixel are calculated. The intensity of each data point on the image is stored as an 8-bit number and so has a value between 0 and 255. The average value for the intensity within the black square over the 'on' pixel is 203, while that for the white square over the 'off' pixel is 24; therefore the contrast ratio is ~8:1 – a typical value for SLM backplanes. Note that here the image was

grabbed when the device was operating very slowly; when the device is operated at higher frame rates the contrast ratio is effectively reduced by 50% because of the charge-balancing section of the frame time.
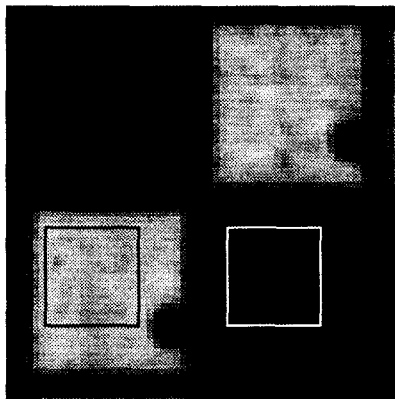
Figure 5-15: Estimating the pixel on/off contrast ratio. The superimposed white and black squares illustrate the areas over which each intensity average was taken.

5.5.5 Uniformity

SSFLC devices require a thin 1–2 μm cell gap to suppress the helical precession of the FLC molecules. The SBS256 backplane wafers supplied by AMS exhibit a significant curvature across their diameter, typically 2–3 μm per die. This curvature probably relieves the stress introduced by by the various deposition, heating, and etching steps during processing. Once a wafer is sawn up, die curvature is reduced to less than 1 μm, but this is still significant when compared to the 2 μm cell gap defined by the spacer balls we use. The die bonding procedures also add to the die curvature:

- The die are glued into chip packages which are not optically flat.

- The thermal-setting glue used is quite viscous, and is therefore difficult to spread into a uniform film.

- The die cannot be pressed flat into the package because they do not have a protective over-glaze coating covering the top-level metal.

These problems result in the assembled SLMs exhibiting significant fringing patterns (see Fig. 5–16).

(a)   (b)
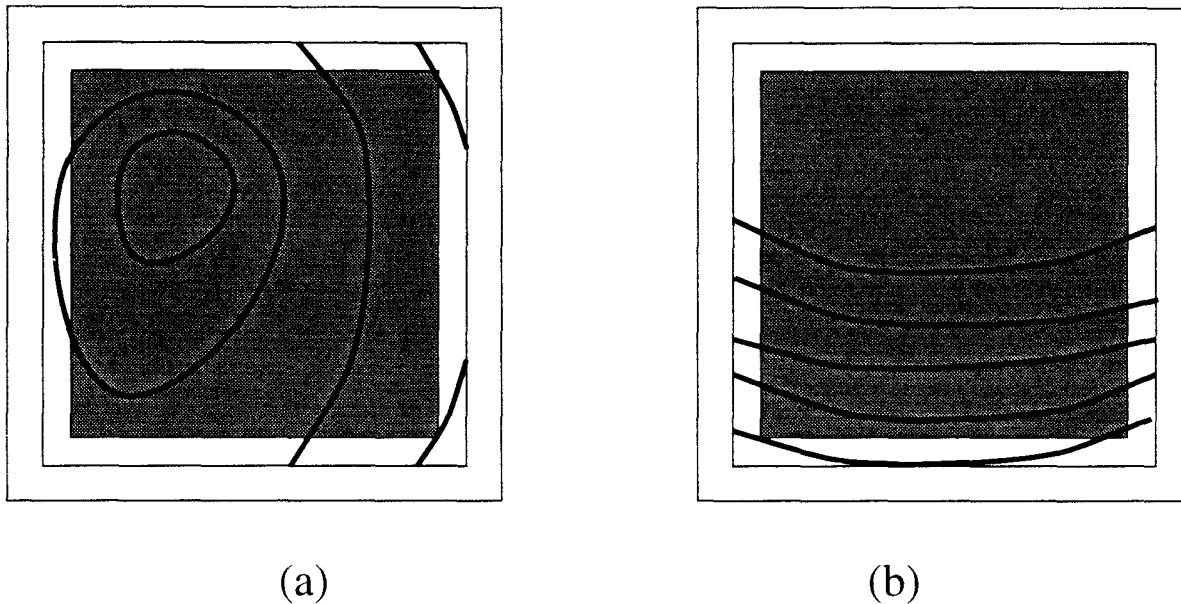
Figure 5–16: Two typical fringe patterns present on fabricated SLMs. (a) The bull's-eye fringe pattern – a result of backplane warpage, and (b) the wedge pattern – probably caused by a piece of dirt trapped at between the backplane and cover glass the edge of the array.

Solutions to die curvature problems are critical to the success of large area FLCOS devices such as the SBS256 and 512 × 512 DRAM SLMs and, at the time of writing, the problems are being investigated by A. O'Hara [78]. Possible solutions include:

- Using a less viscous glue, perhaps one available in film form.

- Coating the die with a protective layer of baked photo-resist so that force can be applied without damaging the top-level metal.

- Bonding the die to an optical flat, then bonding this to the chip package.

5.5.6 Multiplexed grey-scale

As mentioned in Section 3.3, multiplexing techniques can be used to emulate grey-scale on binary mode SLMs.

Temporal multiplexing

The simple linear encoding temporal multiplexing technique is used to generate the grey-scale test pattern shown in Fig. 5–17. For the 16 grey levels shown, 15 bit-planes were required. The device was controlled by the FastInt interface, with $f_{CK} = 1$ kHz, so the grey-scale frame rate is ~66 Hz. Even though the FLC quality on the device shown is not very good, the grey levels are fairly linear. The intensity of some of the pixels are significantly affected by domain defects in the FLC.

Fifteen bit-planes are used to generate the 'Monalisa' test pattern shown in Fig. 5–18.

Spatial multiplexing

Fig. 5–19 shows the optical response of the FLC for the grey-scale test pattern when the reflected light is imaged onto a photodiode. This graph effectively illustrates spatial multiplexing because in each successive frame, an extra column of pixels is switched on. Exactly the same effect is achieved when blocks of 4 × 4 pixels are used rather than columns. When grouped into 4 × 4 blocks, the device functions as a 64 × 64 array, capable of generating 17 grey levels.

5.6 Summary

In this chapter I have presented electrical and optical characterisation results for the SBS256 SLM (see Table 5–4). The current spikes generated by the charge balancing clock signal CK are significant but do not affect the electrical functionality of the device. From probe testing the yield was estimated at 44% for fully working die, and 32% for partially working die. Post-processing planarisation was successful and increased the flat fill factor of the pixel electrode mirrors from 23% to 84%. The SLM is capable of generating a 2 kHz fully charge balanced frame
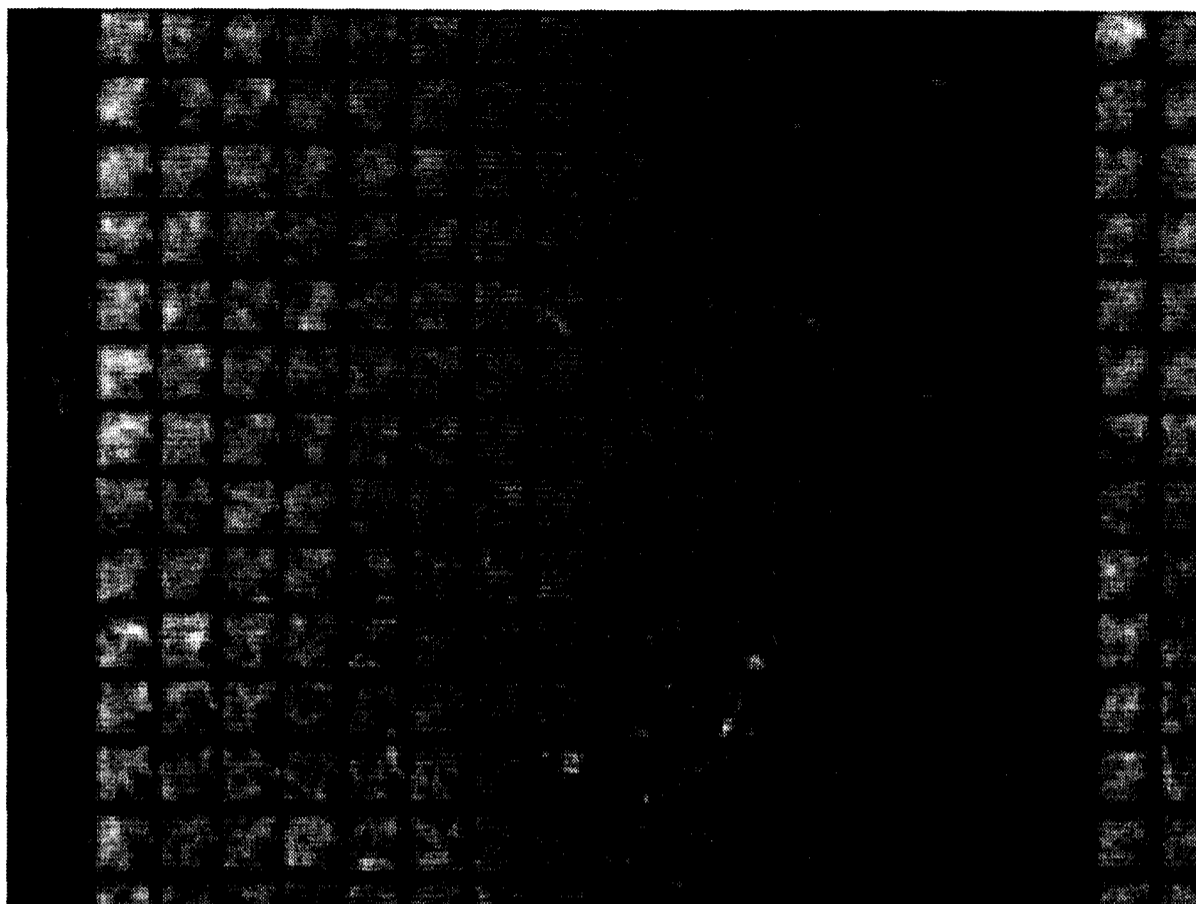
(a)
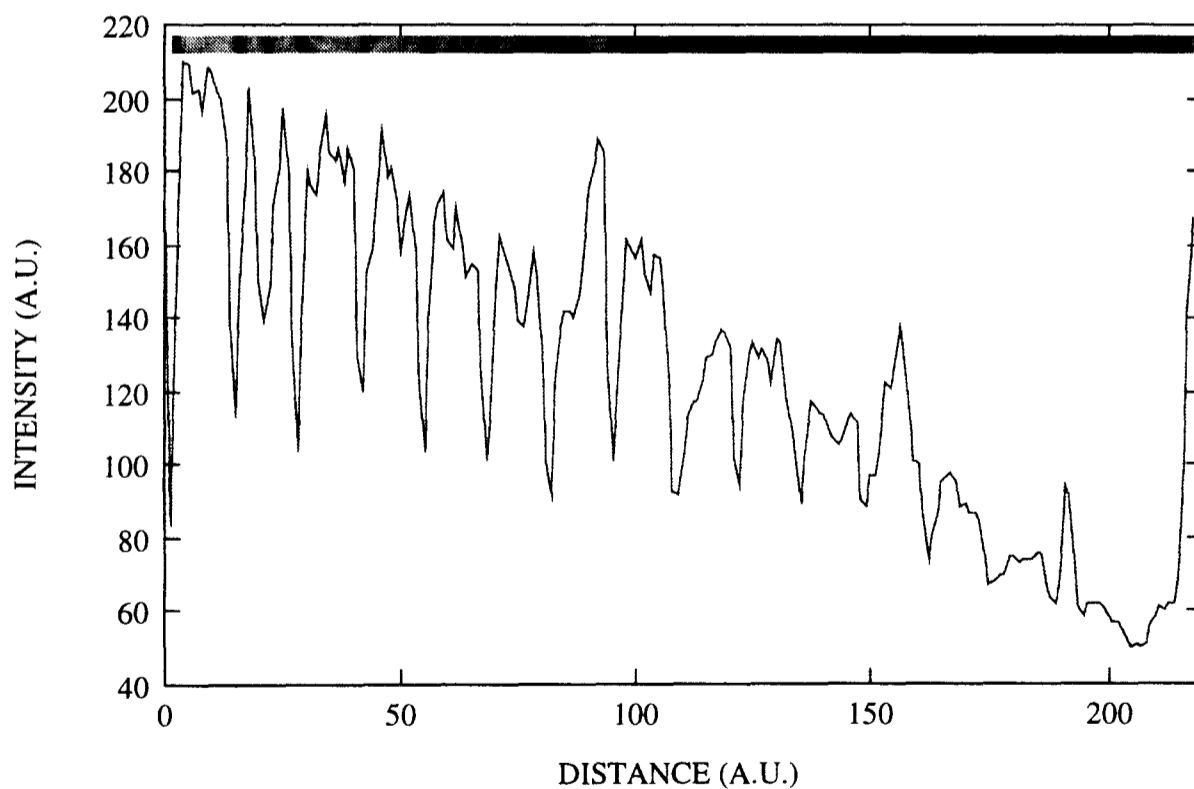
(b)

Figure 5–17: (a) Temporally multiplexed grey-scale on a section of a planarised device (2#19) and, (b) a cross-section intensity profile through a row of pixels (8^{th} row from the top).

Figure 5-18: 'Monalisa' grey-scale test pattern on a section of an unplanarised SLM (8#2). Fifteen bit-planes are used to render the temporally multiplexed grey-scale. Because of an error in the driving software, this device was not properly charge balanced - note that a previous 'triangle' pattern is still partially visible.

Backplane specifications

Number of pixels	256×256
Active pixel array area	10.24×10.24 mm^2
Die area	14.00×14.00 mm^2
Technology	1.2 μm n-well CMOS
Pixel circuit	SRAM-XOR (10 transistors)
Pixel pitch	40 μm
Electrode mirror area	19×19 μm^2 (unplanarised)
	37×37 μm^2 (planarised)
Backplane operating frequency	24 MHz
Frame scan time	84 μs
Drive voltage	5–6 V
Power dissipation	< 90 mW

Optical characteristics

Charge-balanced frame rate	2 kHz
Modulation	Binary (phase or amplitude)
Duty cycle	50% (with unpulsed light source)
Pixel contrast ratio	4–8:1
Grey scale	Temporal or spatial multiplexing

Table 5–4: SBS256 specifications.
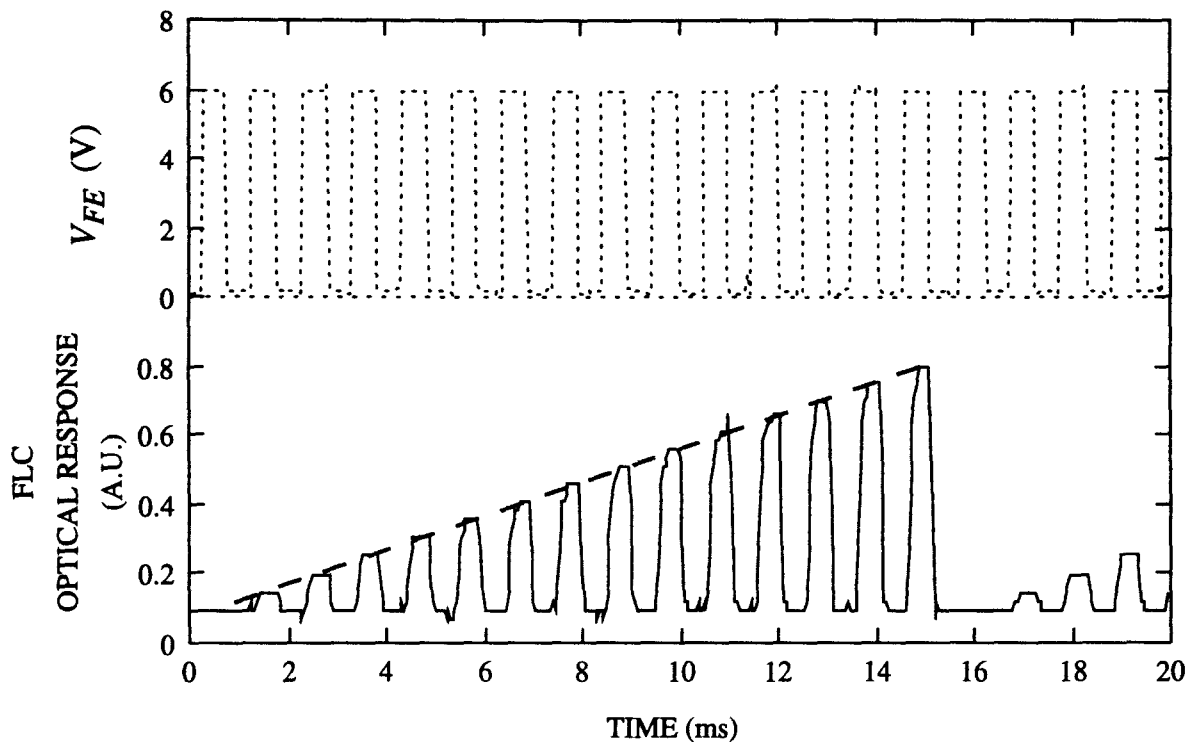
Figure 5–19: Optical response of the FLC for the grey-scale test pattern shown in Fig. 5-17.

rate with a contrast ratio of ∼8:1. Backplane warp was identified as a major performance limiting parameter. Finally, I demonstrated temporal and spatial multiplexed grey-scale techniques for generating real-time grey-scale.

Chapter 6

OANN - a prototype neural-detector backplane

In this chapter I present a prototype neural-detector backplane that can sample the temporally multiplexed grey-scale generated by the SBS256 SLM. Both devices form the basis of an optoelectronic implementation of the hidden layer in a multilayer perceptron artificial neural network architecture. I begin by discussing the requirements for such a system and then describe the simplified system used to characterise the neural-detector backplane (OANN). Photo-induced charge leakage is identified as a potential problem for the activity storage circuits on the backplane, so a technique to significantly reduce charge leakage is presented. The neuron circuit and backplane drive scheme are then described in detail. Finally, test and characterisation results are presented.

6.1 Introduction

Artificial neural networks (ANNs) are useful tools for solving many prediction, modelling, and classification problems. As classifiers, they can be used to identify patterns and trends in data, which can be in numerical, electronic, or optical form. In most ANN systems, the input layer is the location of the massive analogue parallelism, while the subsequent layers are narrower. Optoelectronic implement-

ations are particularly suitable for applications where a high fan-in architecture can be used, and the input pattern is available in the optical domain. As with other optical systems, FLCOS SLMs offer the possibility of compact, low-power, and high frame rate solutions. However, ANNs generally require adaptive, analogue interconnection weights between the neuron processing units. FLCOS SLMs are adaptive but normally operate in a binary mode because the backplanes are based on digital circuits, and the commonly used SSFLC configuration is binary in nature. Nevertheless, as shown in Section 3.3, their high frame rates permit temporal multiplexing (TM) to be used as a means of real-time grey-scale emulation. Although processing rate is sacrificed because of the multiple subframes required, the resulting grey-scale frame rates can still be significantly higher than those available with active-matrix nematic LCDs. For example, nematic LCDs can exhibit about 16 grey levels and can operate at up to 50 Hz, while the SBS256, operating at 2 kHz, can generate 16 grey levels at 500 Hz using four subframes and a binary weighted pulsed light source: a factor of ten improvement.

6.2 System requirements

I begin this section with a brief description of a system that can perform the forward pass function of the hidden layer in the multiple layer perceptron (MLP) architecture. I follow this with a discussion of the advantages and limitations of using TM grey-scale in such a system. I then describe a simplified system suitable for testing and characterising the OANN backplane.

6.2.1 An optoelectronic hidden layer ANN

Fig. 6-1 shows the basic components of a system based on the matrix-matrix crossbar architecture (see Section 3.2.1). The input SLM can be optically or electrically addressed. If optically addressed, the input scene can be converted directly from the incoherent to coherent domain. A device such as the Isophote

SLM described in Section 2.4.3, would be suitable; it would also perform a useful edge enhancement operation on the input scene. If the input SLM is electrically addressed, there will be some data transfer overheads because the image must be captured by, say, a CCD camera, and then scanned onto the SLM.

INPUT	FAN-OUT	WEIGHT	FAN-IN	NEURAL
(SLM)	OPTICS	MATRICE	OPTICS	BACKPLANE
	(HOLOGRAM)	(SLM)	(LENSLETS)	
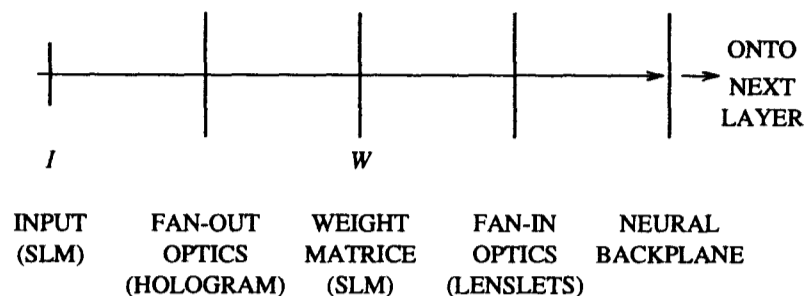
Figure 6–1: Input-to-hidden layer section of an optoelectronic ANN.

Once the image I is in the coherent domain, Fourier filtering techniques could be used to preprocess the image. A fan-out hologram replicates the (preprocessed) image onto the weight matrix SLM, W. After the multiplication, the resulting IW submatrices are fanned-in onto the neural-detector backplane using a lenslet array. To take full advantage of the weight SLM pixel count, a lenslet array with square aperture micro-lens could be used. The backplane performs the summing and thresholding operations and then generates optical or electrical outputs that can be fed onto further layers.

Note in this system, the interconnection weights would first be calculated in software using the backpropagation learning algorithm, and then downloaded to the interface controlling the weights SLM. To cope with the imperfections in the components of the optical system, the weights can then be re-adjusted by placing the system in a learning loop. This approach is used in a variety of electronic pulse-stream [34] and analogue [42] hardware implementations.

6.2.2 Using temporally multiplexed grey-scale

The TM approach has a considerable impact on the design, operating conditions, and performance of the optoelectronic ANN system.

In analogue and pulse-stream neural hardware, capacitors are normally used to store packets of charge that represent the neuron activations. However, with TM grey-scale, there is a significant mismatch between the overall frame time of the SLM and the time it takes to charge a typical capacitor on an IC backplane. For example, it takes 50 μs to increase the voltage on a 5 pF capacitor by 1 V using a 100 nA current, whereas the time to display and charge balance a single bit-plane on the SBS256 is 500 μs for a 2 kHz frame rate. The overall frame time depends on the number of greys required, but even for a single subframe, the voltage on a 5 pF capacitor would rise by 10 V if a 100 nA current was sourced onto it (ignoring the fact that the current sourcing circuit would probably be operated with a 5 V power supply). To overcome this mismatch, a sampling scheme must be used in the neuron circuit to access the activity capacitor.

The long overall frame time of the SLM can also lead to charge leakage problems. Low charge leakage is critical to the functionality of the neuron circuit, so I will discuss it separately in Section 6.3.

Below is a list of some advantages and limitations associated with using TM greys in an optoelectronic implementation.

Advantages

1. Flexible number of grey level weights. Some simple applications may only require a few grey levels, while others may require 6 to 10-bit greys. More greys can be generated simply by adding more subframes. However, there will be a corresponding decrease in the system processing rate.

2. The polarisation based systems described in Section 3.4.3, have two optical pathways, one for positive activations and one for negative activations. This architecture is less amenable for use with integrated neuron backplanes. By implementing negative weights with 2's complement temporal multiplexing (see Fig. 6–2), a single pathway system can be realised.

BITPLANE :   SIGN   2   1   0

WEIGHTING :   (-8)   (+4)   (+2)   (+1)   EQUIVALENT GREY LEVEL
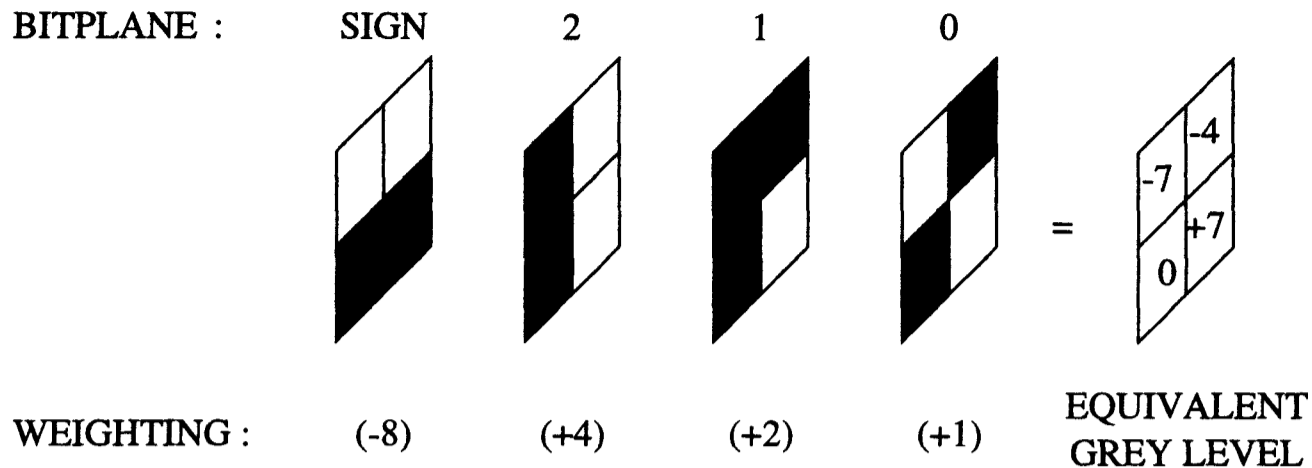
Figure 6–2: Temporally multiplexed 2's complement grey-scale scheme for generating positive and negative activations. Using four bit-planes, sixteen grey levels can be realised – here four different activations are illustrated. The neuron circuit drive scheme must sample each bit-plane with the appropriate binary weighting. Note also, that one of the bit-planes is used for the sign-bits, so the neuron sampling circuit must respond in the appropriate manner.

Limitations

1. Decreased dynamic range within the neuron circuit. The negative activations must be subtracted from the activity capacitor, and then the positive ones added (or vice versa) – they are not done simultaneously, and thus the dynamic range is lower.

2. Two temporally multiplexed SLMs cannot be efficiently cascaded because one must cycle through all its bit-planes for each of the bit-planes on the other. If the weight SLM is temporally multiplexed, then the input SLM must:

   – generate true grey-scale,

   – use spatial multiplexing for grey-scale, or

   – work in a binary mode.

6.2.3 Simplified system

With enough time and resources, the system shown in Fig. 6-1, could be built to perform useful optoelectronic neural computation. However, within the limited time-scale of my PhD project, I built a simplified system to verify that temporal multiplexing can be used for weight generation.

The simple optical system shown in Fig. 6-3, provides the prototype 4 × 4 OANN backplane with optical input signals corresponding to a 16 × 16 input pattern. To eliminate the problems of cascading two SLMs (the input and weight devices), the multiplication of the input pattern by the weights matrix is calculated in software. The resulting *IW* products for each of the hidden layer neurons are encoded into the appropriate bit-planes and down-loaded to the SBS256 FastInt RAM store. Therefore, there is no need for the input SLM or the fan-out hologram. The fan-in lenslet array is replaced by a microscope objective lens, so to match the 1:1 mark-space ratio of the input photodiodes on OANN, the combined *IW* patterns must also be placed on the SBS256 with a 1:1 mark-space ratio. Each of the sixteen 16 × 16 *IW* patterns occupy 640 × 640 μm^2 on the SBS256, while the photodiodes are 100 × 100 μm^2. The ×6.3 microscope objective scales and focuses the *IW* patterns onto the appropriate photodiodes on the neural-detector backplane. The ×6.3 objective matches the desired demagnification to within 2%.

The OANN backplane (see Fig. 6-4) detects the incoming light signals, stores the resulting activations, and then encodes the activations as electronic pulse width outputs [12]. Pulse width outputs are desirable because they can be fed directly into the EPSILON[1] pulse-stream ANN chips [34,45], so that further electronic layers can be added to perform more complex classification tasks.

---

[1]EPSILON – Edinburgh Pulse Stream Implementation of a Learning Oriented Network.
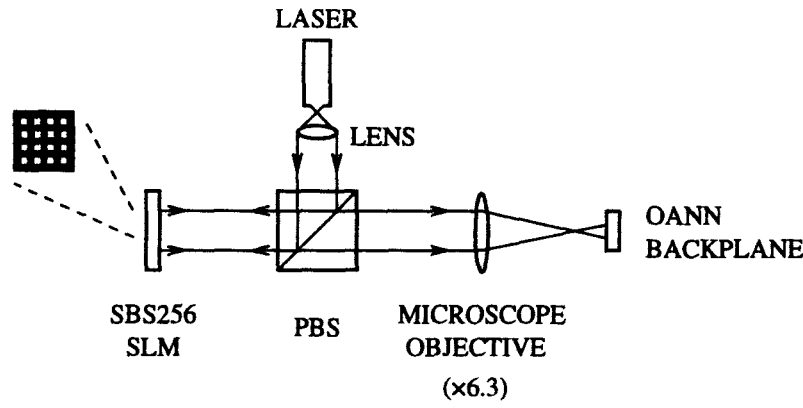
Figure 6–3: The simplified optical system suitable for characterising the OANN neural-detector backplane. The 1:1 mark-space ratio of the *IW* submatrices on a 128 × 128 section of the SBS256 are required to match the mark-space ratio of the input photodiodes on the OANN backplane.

Note that the simplifications made to the original system shown in Fig. 6–1, are purely to make the test system easier to build, align, and use, and are not limitations of the TM approach.

6.3 Activity Storage

As mentioned in the previous section, a sampling scheme must be used to access and update the activity capacitor C_{ACT} within the neuron circuit. The single-transistor DRAM circuit, as used in the DRAM-type SLMs, is a simple sample and hold circuit (see Fig. 6–5). However, if used in an optical system without any modification, it can suffer from severe photo-induced charge leakage effects. Once C_{ACT} has been charged and isolated, the charge can leak away through the transistor channel off resistance R_{OFF} (typically $10^{11\pm1}\Omega$), and through the drain-substrate reverse-biased junction D_{DB} of the pass transistor. The reverse-biased junction D_{DB} is the more significant leakage route. Any thermally or optically generated minority-carriers (electrons in the p-substrate) near it can be swept across its depletion region so discharging C_{ACT}.
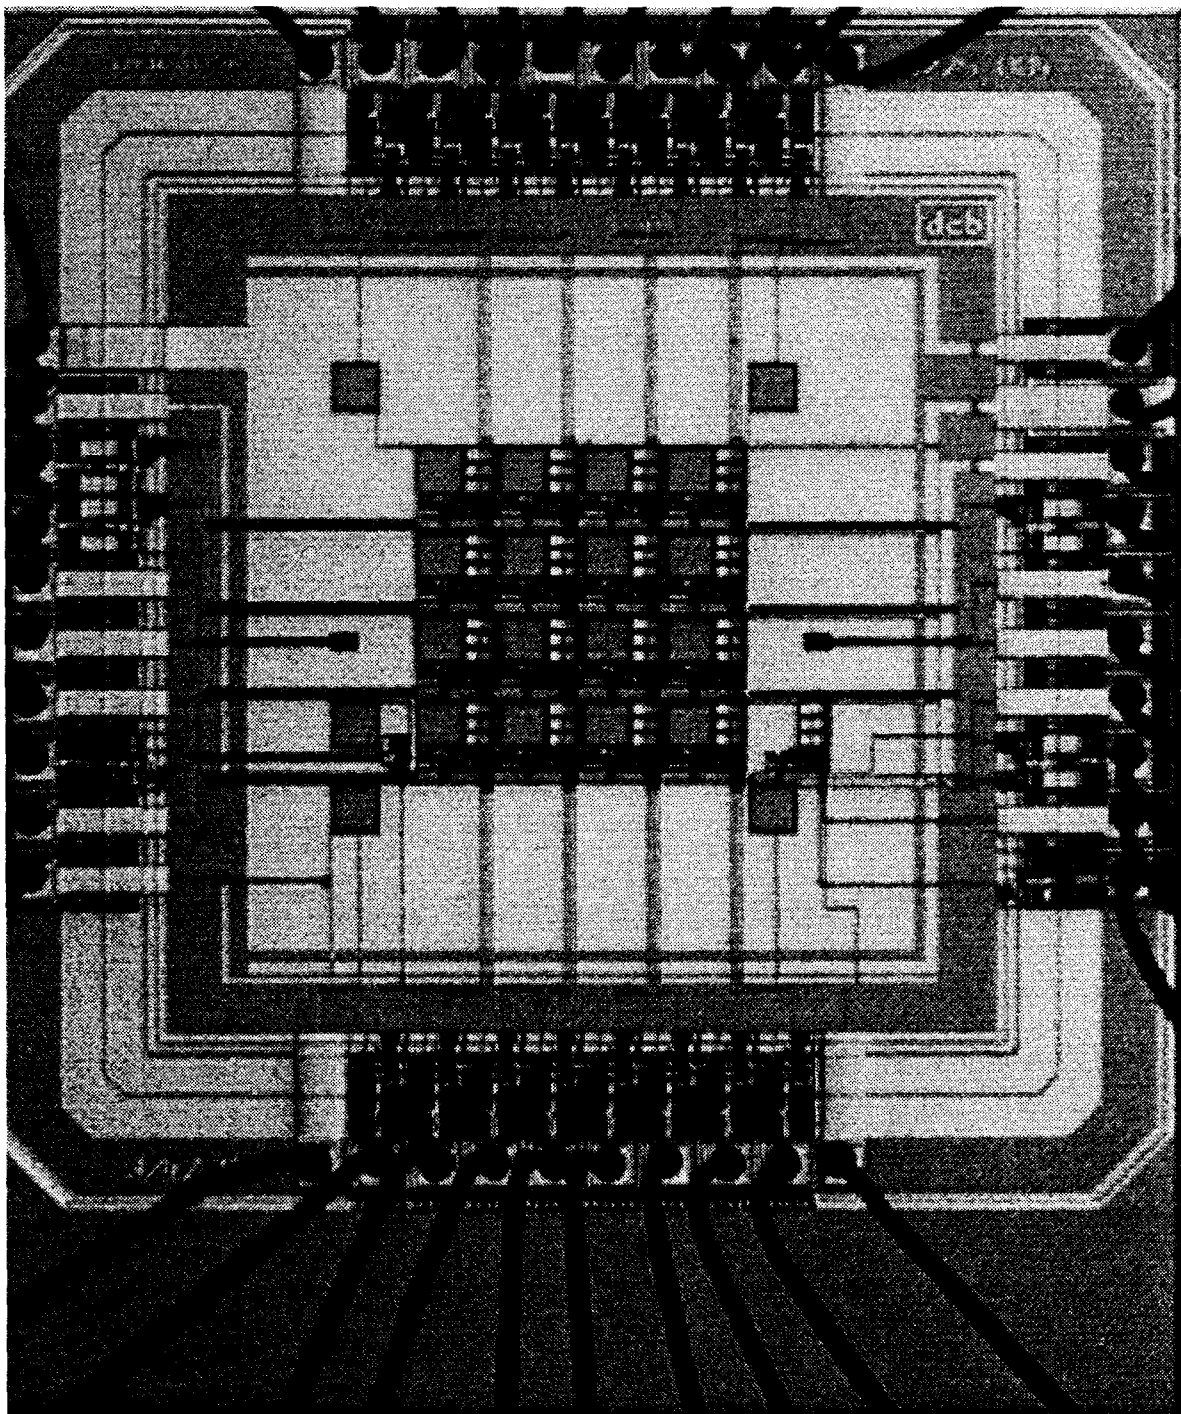
Figure 6–4: Photomicrograph of the 4 × 4 OANN backplane. The 3 × 3 mm ^{2} die was fabricated by AMS on a CAE multi-product wafer brokered by EURO-CHIP. I chose to implement it on the CAE process because I had used it for the SBS256 backplane, and also because it is an epitaxial process which can help reduce charge leakage effects (see Section 6.3). Each neuron occupies 200 × 200 μm^2, with the input photodiodes placed on a 1:1 mark-space ratio. Note that the active area is almost entirely covered in METAL2 to reduce the number of photo-induced electrons in the p-substrate.
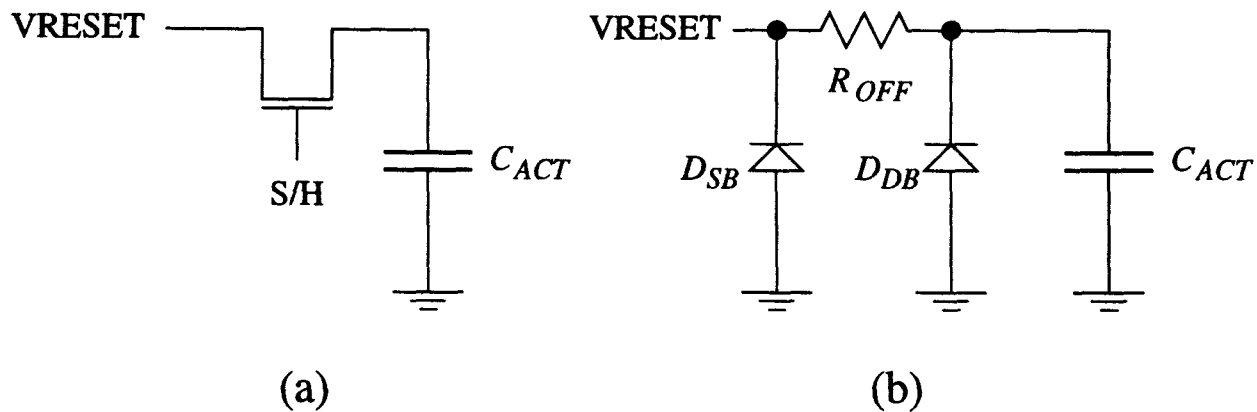
Figure 6–5: (a) Simple sample and hold circuit using one transistor and an activity capacitor, and (b) equivalent circuit when S/H = 0.

Thermally and optically generated carriers are constantly being created and then eliminated by recombination. When carriers are generated, they can either recombine locally or diffuse and drift away from the region where they were generated. The time it has to diffuse is called the carrier life-time and the distance it moves is known as the diffusion length [92]. Carrier recombination can occur in the bulk or at a surface. The surface recombination rate depends on the how the surface was prepared, and often dominates the recombination process. For Si the electron carrier lifetime τ_n is typically between 10 and 100 μs. The diffusion length L_n is given by

$$L_n = \sqrt{\frac{kT\mu_n\tau_n}{e}} \tag{6.1}$$

so when the electron mobility μ_n = 497 cm^2/Vs (for the CAE process), τ_n = 10–100 μs, and kT/e = 25 meV at room temperature, the consequent value for L_n is between 113 and 357 μm. These values are significant because they show that metal shielding layers must be extended over quite a distance beyond range to reduce the effect of the photo-induced electrons on sensitive circuitry.

Even though a metal protective layer should help to reduce charge leakage effects, 100% coverage is unlikely if only two layers of metal are available on a particular process. It is normal layout practice to route signals both horizontally and vertically on the metal layers, so there will almost certainly be gaps for light to impinge onto the silicon substrate. Therefore, photo-induced minority carriers

will probably be generated near storage capacitors if a reasonable circuit packing density is used: further protection is required.

6.3.1 Improving charge storage – the n-well ring

As most of the photo-induced carriers are generated near the surface, the access transistor to C_{ACT} can be surrounded with an n-well ring tied to the positive power rail so that minority electrons tend to be swept into it, before they have a chance to reach the access transistor drains (see Fig. 6–6). The n-well ring structure has been used by other researchers to reduce latch-up [99] and to improve the signal-to-noise ratio of the Hughes optically addressed liquid crystal light valve [111]. An epitaxial process, such as the CAE process, can also improve the efficiency of the n-well ring because the built-in electric field created by the doping gradient from the epi to substrate, attracts minority carriers towards the surface [99].
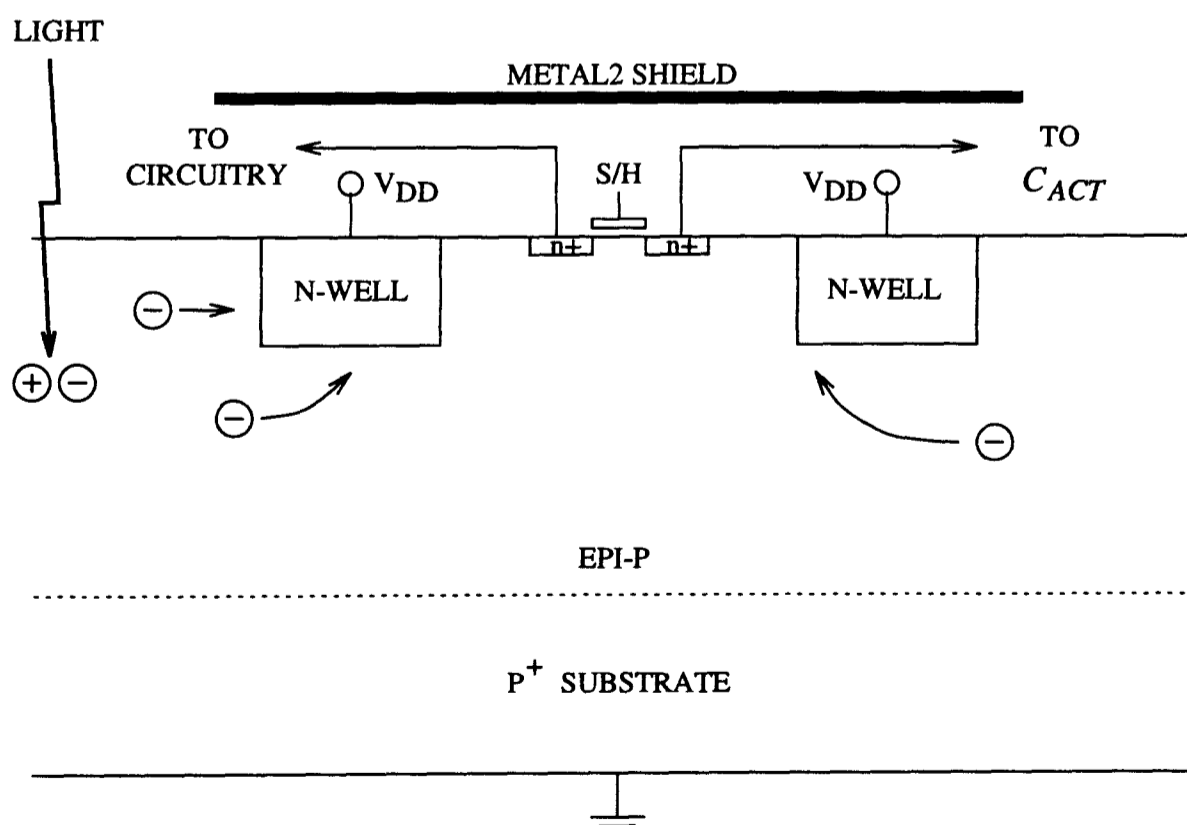
Figure 6–6: Charge storage n-well ring protection to reduce charge leakage.

I incorporated two test circuits on the SBS256 backplane to compare the effect of using, and not using, the n-well ring protection structure (see Fig. 6–7). The simple source-follower circuit used to monitor the activity voltage has linearity and range limitations but was considered adequate. Fig. 6–8 shows the results of the 1% charge decay time for a 5 pF capacitor with, and without, n-well ring protection. The measurements were taken by Matthew Aitken of the Electrical Engineering department as part of his honours degree project. This simple protection structure decreases the decay rate by a factor of over 2000 for a range of light intensities. Extrapolating to the 5.6 mW/cm^2 intensities expected in the optical system, the 1% leakage time is just less than 100 ms. Therefore a complete forward pass calculation can be performed without any appreciable charge leakage from C_{ACT}.
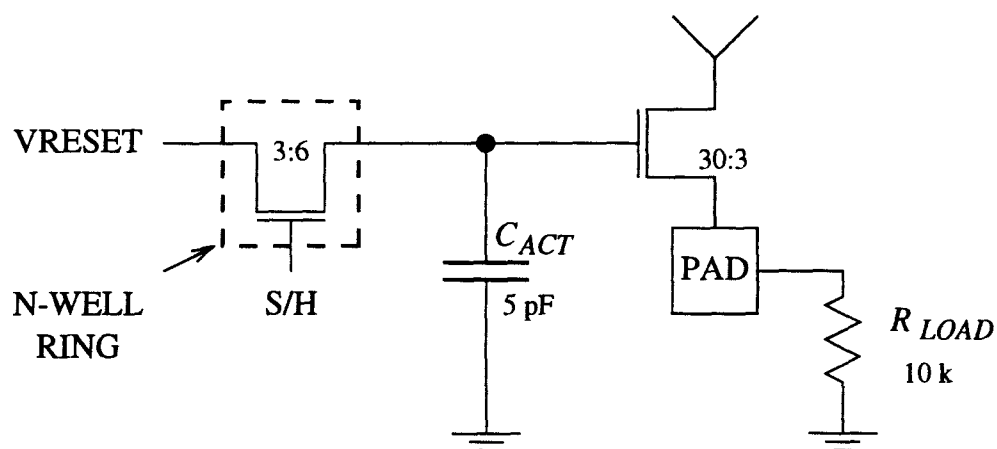
Figure 6–7: Test circuit included on the SBS256 backplane to compare charge leakage from C_{ACT} with, and without, n-well ring protection.

6.4 Neuron circuit

The neuron circuit is a simple optical-in, electronic out, processing node. Each neuron (see Fig. 6–9) consists of a photodiode, current-routing circuitry, an activity storage capacitor C_{ACT}, and a two-stage comparator to encode the neuron activity into a pulse width signal.
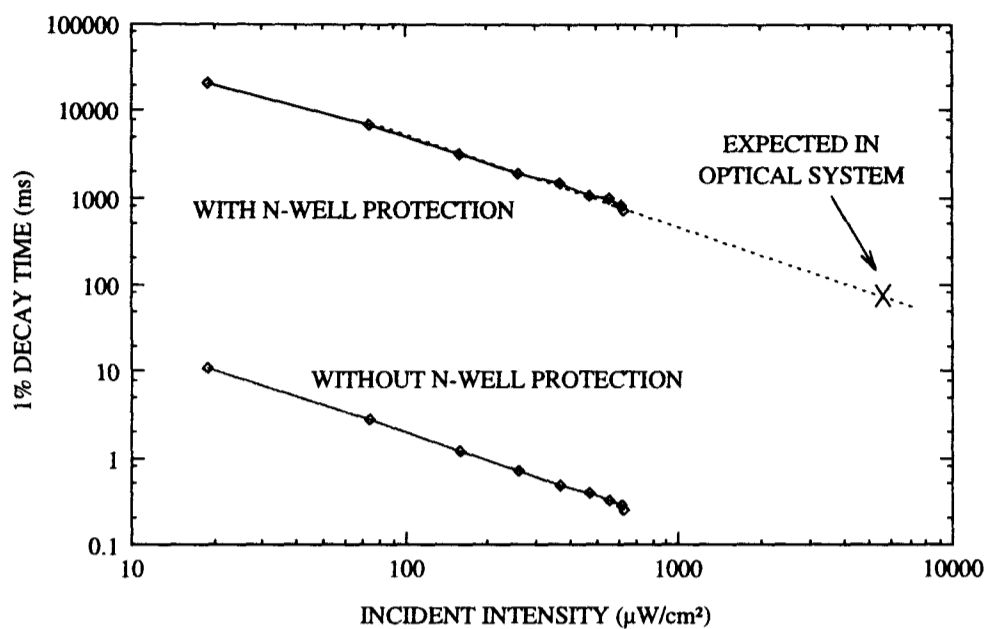
Figure 6–8: Charge leakage with, and without, n-well ring protection on the the SBS256 test circuits. Results courtesy of Matthew Aitken.
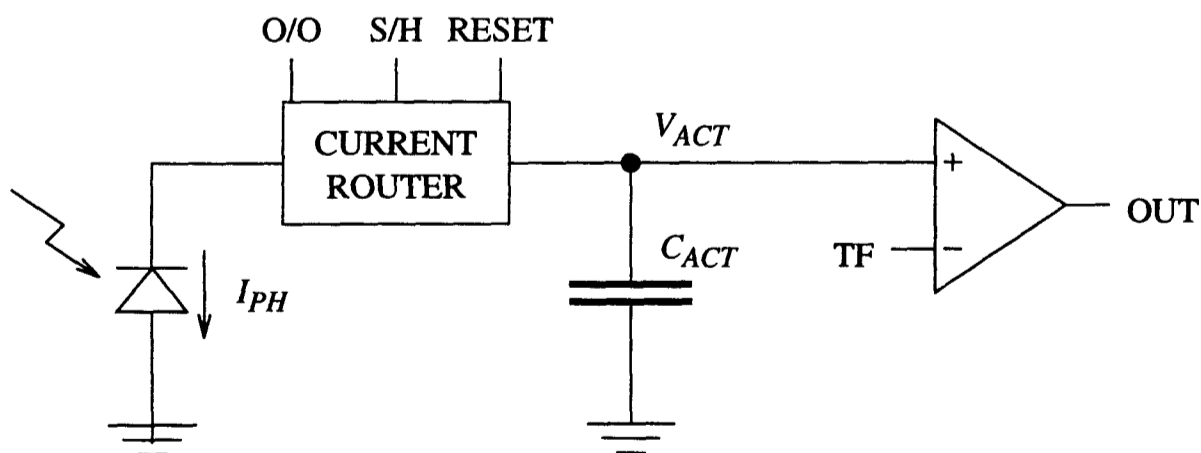
Figure 6–9: Schematic of a single neuron.

Inputs			Effect on
RESET	O/O	S/H	V_{ACT}
0	0	0	No change
0	0	1	Increases in proportion to I_{PH}
0	1	0	No change
0	1	1	Decreases in proportion to I_{PH}
1	x	x	Resets to VZERO

Table 6-1: Truth table for the current-routing control signals. When S/H =1, V_{ACT} rises or falls (depending on O/O) by an amount $I_{PH}\tau_{\mathrm{S/H}}/C_{ACT}$, where $\tau_{S/H}$ is the sample and pulse time.

6.4.1 Operation

A photodiode operating in the photoconductive mode is used to detect the incoming IW bit-planes. The resulting photocurrent is sunk from the current-routing circuitry. To implement positive and negative weights using a 2's complement drive scheme, the current-routing circuitry must be able to selectively source the photocurrent onto, or sink it from, the activity capacitor C_{ACT}. The signal O/O is used to define the direction of current flow whenever the sample and hold signal S/H is pulsed (see Table. 6-1).

Fig. 6-10 illustrates the 2's complement drive scheme used to emulate 32 grey levels - 16 positive (including 0), and 16 negative. Five subframes are required, one for the sign-bit, and four for the binary weighted bits. The sequence starts with a RESET pulse to set V_{ACT} to a zero activation state, VZERO. When S/H is high, V_{ACT} rises or falls (depending on O/O) by δV_{ACT} given by,

$$\delta V_{ACT} = \frac{\tau_{\mathrm{S/H}} I_{PH}}{C_{ACT}} \tag{6.2}$$

where $\tau_{S/H}$ is the sample and pulse time.

In subframe 5, the O/O signal is high so when S/H is pulsed, a current I_{PH} is

sunk from C_{ACT}. Then for the remaining subframes, O/O is low so when S/H is pulsed, the current I_{PH} is sourced onto C_{ACT}.
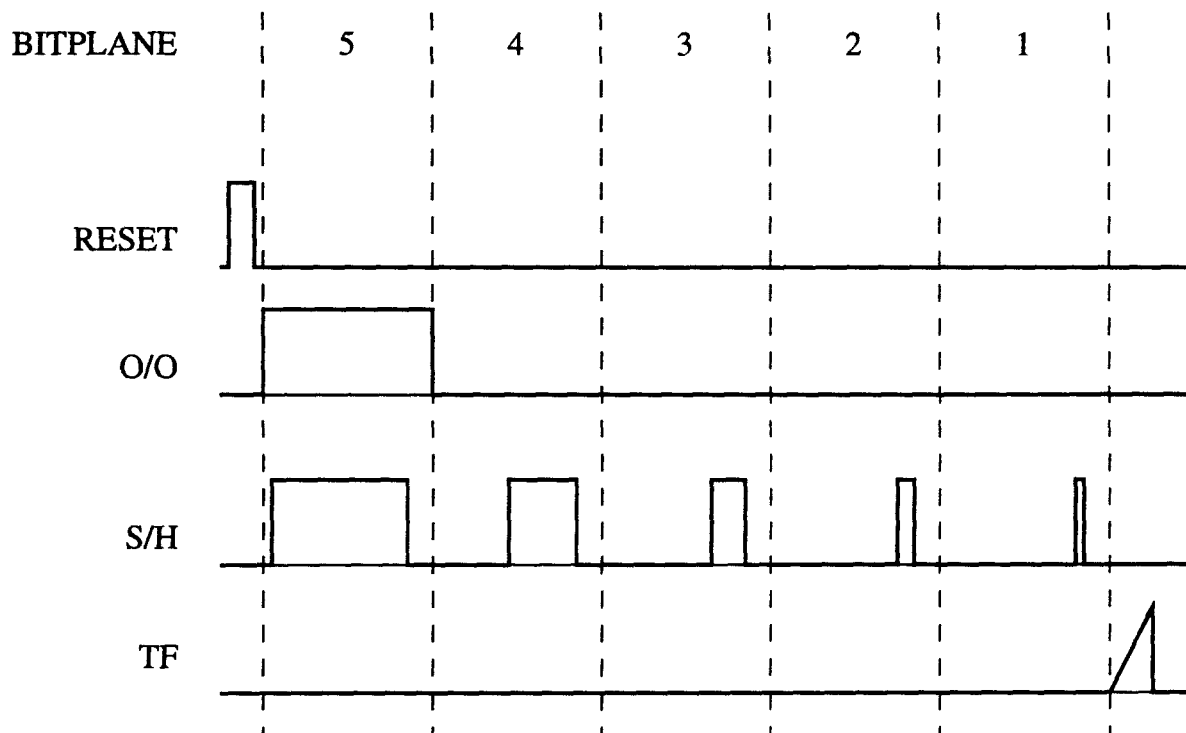
Figure 6-10: OANN control signals for the 2's complement drive scheme to implement 16 positive and 16 negative weights. The S/H sequence (not to scale) is a binary weighted pulsed signal. The pulses occur just before the falling edge of the SLM charge-balancing global clock signal CK (not shown). This gives the FLC over 'on' pixels as much time as possible to switch, so that there is not a spurious pattern on the SLM when S/H is pulsed.

The S/H signals are synchronised with the bit-planes that appear on the SLM. More specifically, they occur just before the falling edge of the SLM global clock signal CK: remember that CK is used to charge balance the FLC (see Section 4.6). This ensures that the FLC over 'on' pixels has as much time as possible to switch, so that there is not a spurious pattern on the SLM when S/H is pulsed.

Once all the bit-planes have been presented, the comparator encodes the activation voltage V_{ACT} as a pulse width encoded signal by applying the desired transfer function signal TF to the other input of the comparator.

6.4.2 Back-end circuit details

Fig. 6–11 shows a transistor-level diagram of the neuron back-end: the input photodiode, current-router, and activation capacitor.
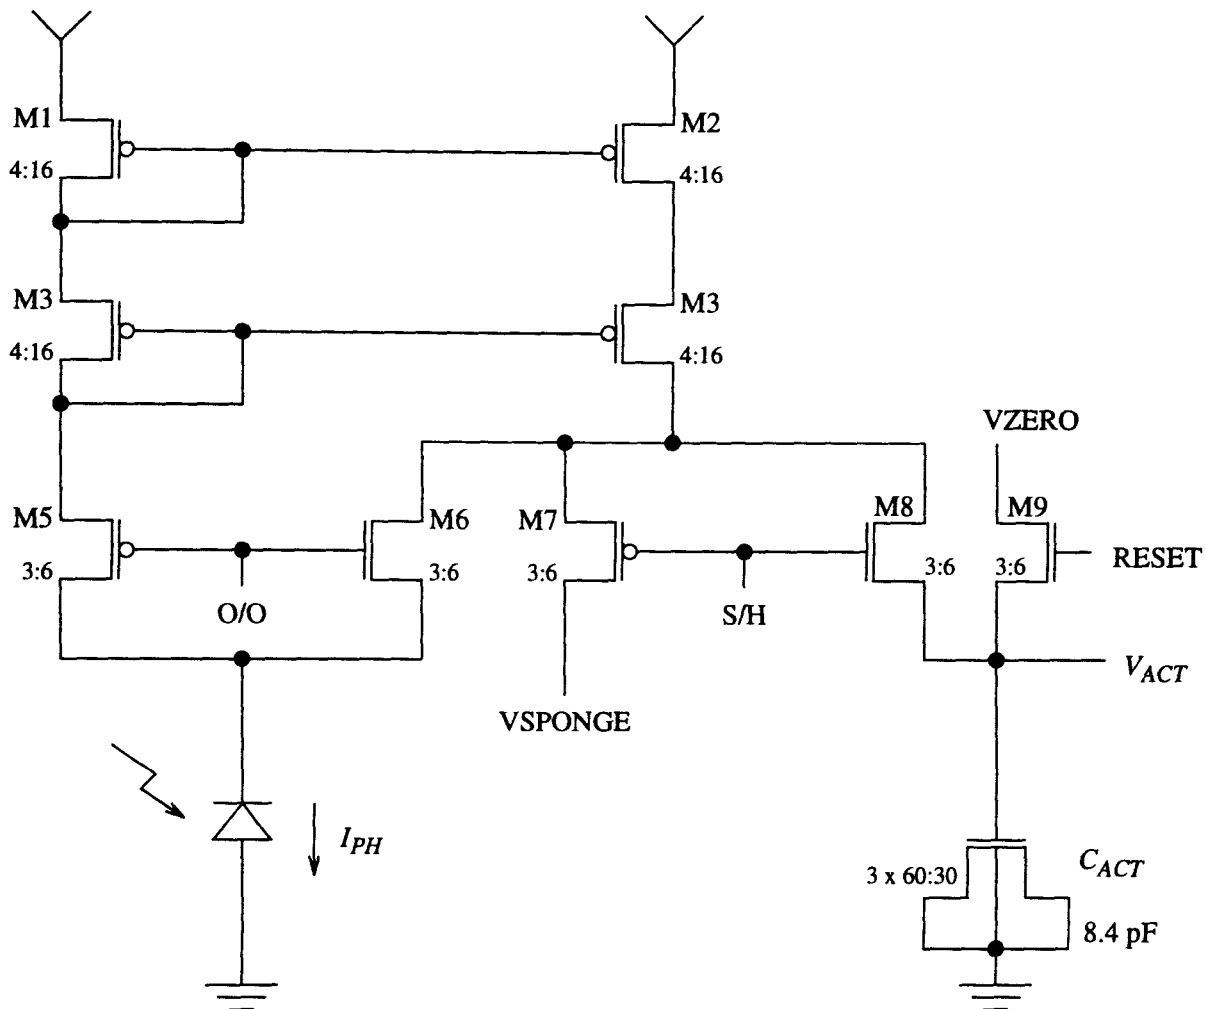
Figure 6–11: Transistor level diagram of the neuron back end.

The back-end circuit should respond linearly to the intensity of the incoming light, so a photodiode, rather than a phototransistor, is used to detect the incoming light because it has a more linear response. The dark-current of a typical pn$^+$ photodiode on the n-well CAE process is approximately 1 nA [94], so the incoming light levels should be high enough to generate currents of at least 10 nA in the photodiode.

The O/O signal is used to enable or disable the cascode current mirror. A cascode current mirror has a much more linear response than that of a simple two-transistor current mirror (see Ref. [2, pages 219–239]). The current mirror transistors are long and thin (W = 4 μm, L = 16 μm) to decrease the effects of

process variation. When S/H is high and O/O is low, the copied version of I_{PH} is sourced onto C_{ACT}, but when O/O is high, the current mirror is disabled, and I_{PH} is drawn directly from C_{ACT}. When S/H is low, current is either sourced or sunk by the reference voltage VSPONGE, depending on whether O/O is high or low. The sizes of the switching transistors are not critical to the functioning of the circuit.

The actual value of activation capacitor C_{ACT} is not critical, so it was made as large as space would permit. The gates of three large nMOS transistors with their sources and drains tied to ground (each with a gate area $= 65.2 \times 30 \ \mu\text{m}^2$), result in $C_{ACT} = 8.4 \pm 0.4$ pF. The thin gate oxide within the transistors' structure helps provide the largest capacitance per unit area (1.4 ± 0.1 fF/μm^2) on the CAE CMOS process. However, when a transistor gate is used as a capacitor, care must be taken to ensure the voltage between the gate and the substrate is kept greater than the transistor threshold voltage (typically $V_T = 0.6$–0.8 V for nMOS devices).

V_{ACT} operating range

The circuits and components limit the range over which the activation voltage V_{ACT} can operate. The upper range is limited by the minimum voltage required across the output transistors of the cascode current mirror to ensure they remain in the saturation mode. It is also limited by the transfer characteristic of the sample and hold pass transistor M8. HSPICE simulations (SS worst-case) showed that V_{ACT} could rise to just above 4 V when $V_{DD} = 6$ V. As mentioned above, the lower limit is constrained by the requirement to keep the voltage on a transistor-capacitor above V_T. Therefore, the drive scheme should ensure that the combination of the I_{PH} photocurrent levels, and the total duration of the S/H pulses, do not result in V_{ACT} rising above 4 V, or below 1 V.

Charge-sharing event

A significant charge-sharing event can occur because of the way the current-router is implemented. When O/O is low, and the S/H pulse-train is applied, the circuit behaves properly by dumping the appropriate currents onto C_{ACT}, because the photodiode junction capacitance is isolated from C_{ACT} by the current mirror. However, when the O/O signal is high, and S/H is pulsed, the photodiode and C_{ACT} are effectively connected in parallel by the transistor M8, so a charge-sharing event occurs because the photodiode junction capacitance (voltage dependent, but typically 2 pF for the voltages present in circuit) is similar to that of C_{ACT}.

This problem was identified early in the backplane design cycle – an extra current mirror isolating-stage could have been added but would have increased the complexity of the circuit. Instead, I decided to ensure the drive scheme sequence minimised the charge-sharing event (see Fig. 6–10). This is achieved by presenting the sign-bit subframe at the start of the sequence, and ensuring $V_{VSPONGE} = V_{VZERO}$. Therefore, after the RESET pulse at the start of the sequence, the voltage across the photodiode and V_{ACT} should be equal, so negligible charge-sharing occurs when S/H is pulsed. Note that without an extra isolating current mirror, this is the only way negative activations can be satisfactorily presented to the neuron. A scheme where, say, four positive and then four negative subframes are presented to this current-router circuit, would not work, because there would be four charge-sharing events associated with the negative activation subframes. However, this scheme is less attractive anyway, because it requires almost twice as many subframes as the 2's complement scheme, to represent the same number of grey levels.

AC analysis

The back-end circuit exhibits a relatively low frequency response. HSPICE AC analysis simulations showed that when $I_{PH} = 100$ nA and a small 10 nA AC signal

is superimposed on the input, the circuit has a 25 kHz 3 dB bandwidth response (see Fig. 6–12). This limited response is a combination of the capacitance of the large (100 × 100 μm^2) photodiode, and the relatively low currents present. Even though the circuit operates in a pulsed mode, and so a standard HSPICE AC simulation is not strictly appropriate, it shows that the circuit can cope with the 1–2 kHz frame rates available from the SBS256.
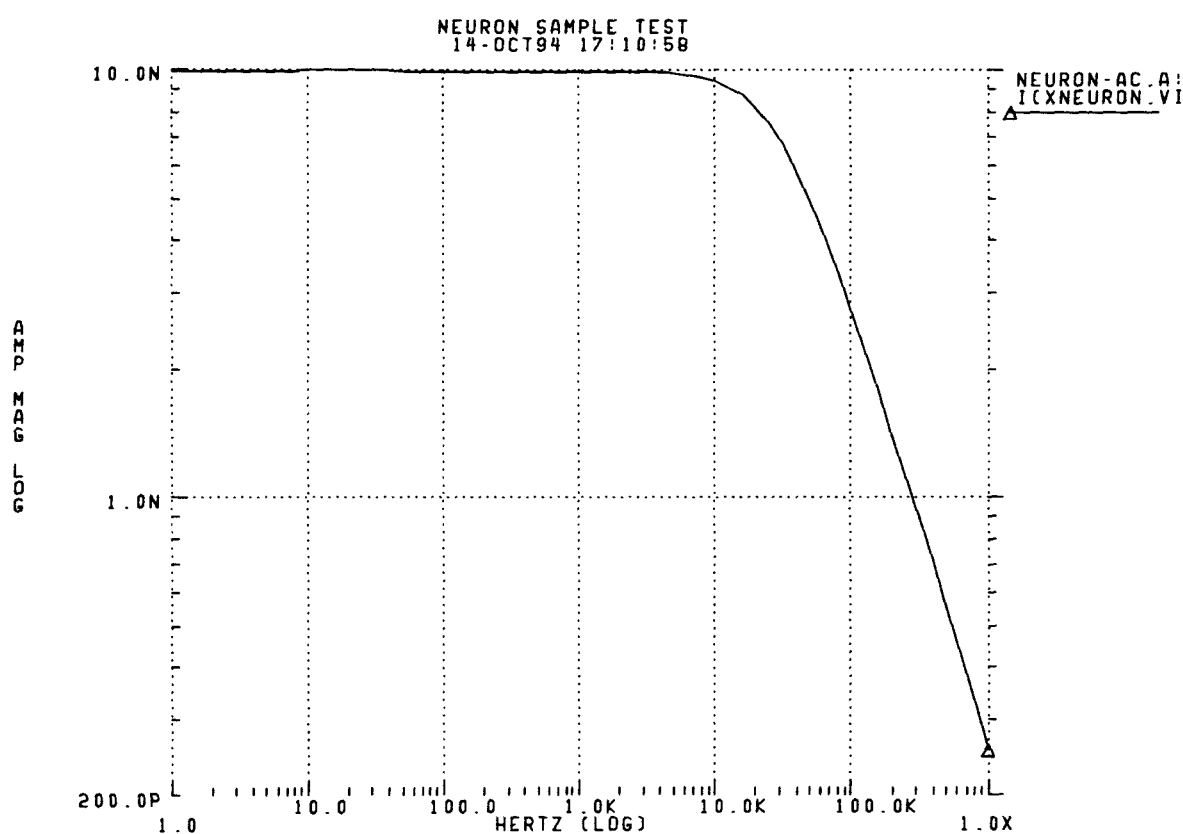
Figure 6–12: HSPICE AC analysis simulation of neuron back-end. For I_{PH} = 100 nA and a small 10 nA AC signal superimposed on the input, the circuit has a 25 kHz 3 dB bandwidth response.

Noise considerations

HSPICE noise analysis simulations showed that the thermal, shot, and flicker noise [100] generated within the circuit components was equivalent to a few millivolts of noise on the activity capacitor, and so would probably be overwhelmed by system noise present on the power supply rails and from capacitive coupling of the electrical input and output signals.

6.4.3 Front-end circuit details

After all the bit-planes have been presented and sampled onto C_{ACT}, the 2-stage comparator (see Fig. 6-13) encodes V_{ACT} as a pulse width signal by applying the desired transfer function signal TF to the other input of the comparator. The comparator is based on the standard design given in Allen and Holberg (Ref. [2, pages 333-349]). It does not need to have high performance characteristics so has been designed to be quite small, with the only design constraints being a common mode input range of 1-5 V and a tail current of 2 μA. The output of the comparator is then passed to an inverter busline-driver. It can drive the busline that routes the neuron output to an output pad, with a rise/fall-time of $\sim$ 100 ns. If the comparator drove the busline directly, the fall-time would be almost 3 μs because the charge on the busline capacitance ($\sim$ 0.5 pF) would have to be discharged through M16, which sinks a constant current of 2 μA.
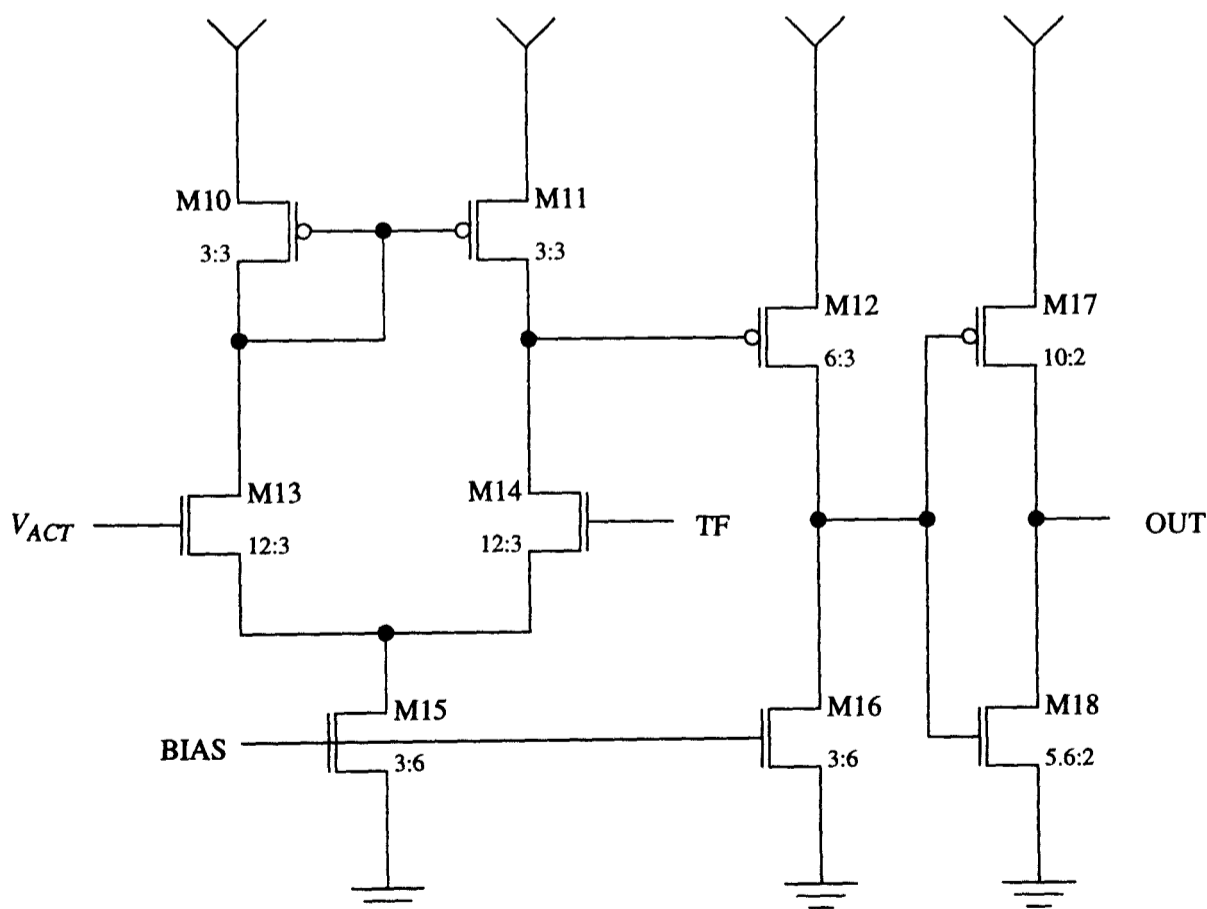
Figure 6-13: Transistor level diagram of the neuron front end.

6.4.4 Neuron layout

For this prototype backplane, high packing density was not critical, therefore the neuron (see Fig. 6–14) occupies a relatively large 200 × 200 μm^2 area. With a 200 μm pitch, the 100 × 100 μm^2 photodiodes are on a 1:1 mark-space ratio, and there is plenty of space for the transistor circuitry and activity capacitor. The photodiode is clearly visible in the top-left quarter of the circuit, while the triple transistor activity capacitor is placed in the top-right quarter. Note the n-well guard ring surrounding the access transistors in the bottom-centre section. Most of the control signals are routed horizontally in METAL1. However, the power rail VDD, BIAS, and the neuron output buslines are routed in METAL2. VDD also acts as an optical shield covering as much of the circuitry as possible.

6.5 OANN backplane

The 3 × 3 mm^2 OANN backplane (see Fig. 6–4) was fabricated by AMS on a CAE multi-product wafer brokered by EUROCHIP. As there are only 16 neurons, the output signals (N00–15) are routed individually off the array to output buffer pads, so no addressing circuitry is required. For larger arrays (> 8 × 8), multiplexing decoders would probably be necessary. The four photodiodes at the corners of the array can be used to help align the optical system. A few other test circuits are also included but will not be discussed here.

6.6 OannInt interface

I built the OannInt interface to test and characterise the OANN backplane (see Fig. 6–15). The interface is computer controlled and uses the same control signals as the SBS256 FastInt interface so that they can both be operated in parallel by a single PC. The interface control sequence is defined by a simple state machine

PHOTODIODE

ACTIVITY
CAPACITOR

200 μm

CURRENT ROUTER
& N-WELL RING

2-STAGE
COMPARATOR

LAYER KEY:

N-WELL   CONTACT

DIFFUSION   METAL1

POLY   VIA

BORON   METAL2
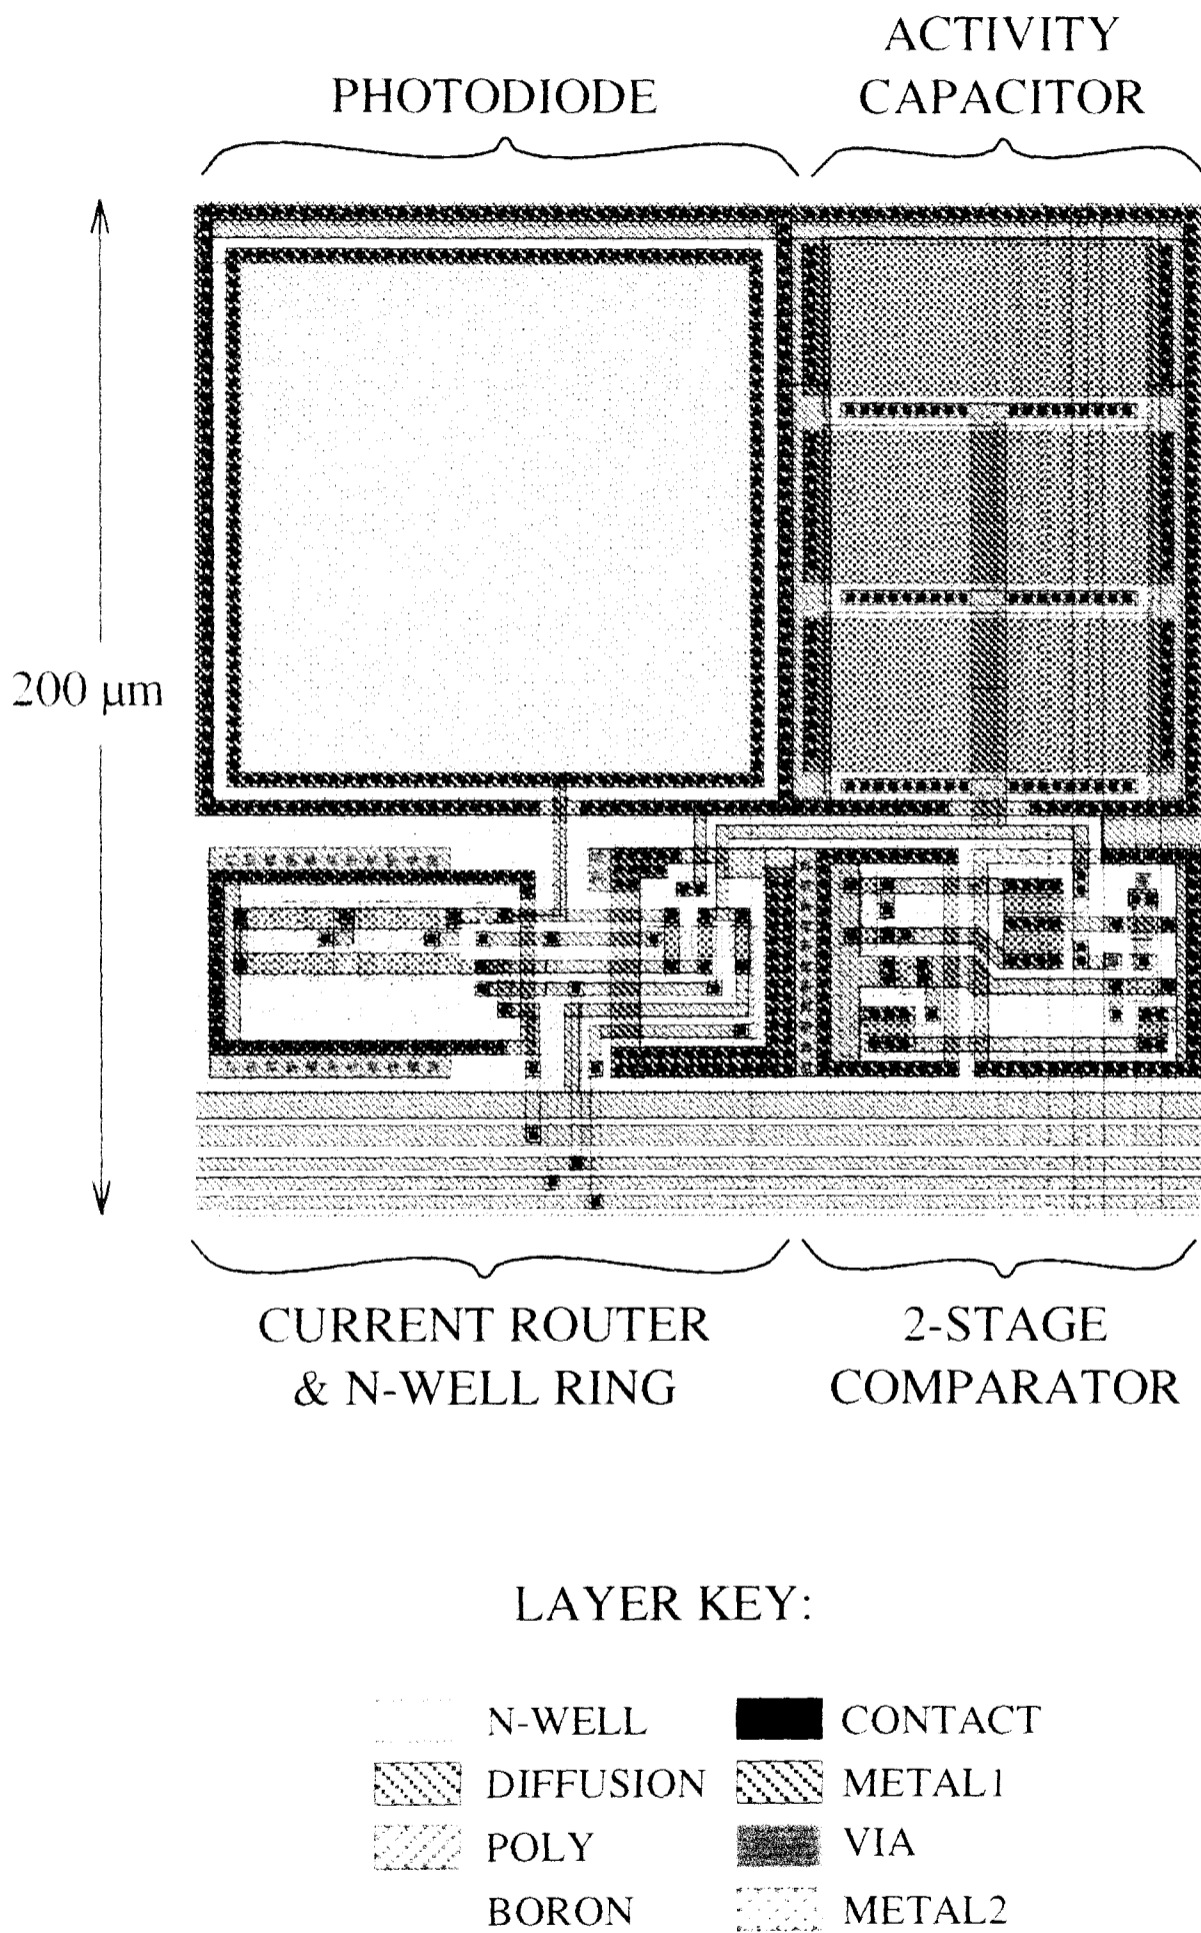
Figure 6-14: Layout of a single neuron.

clocked by a signal derived from the crystal driving the FastInt state machine. This ensures that both interfaces can operate synchronously when configured correctly.
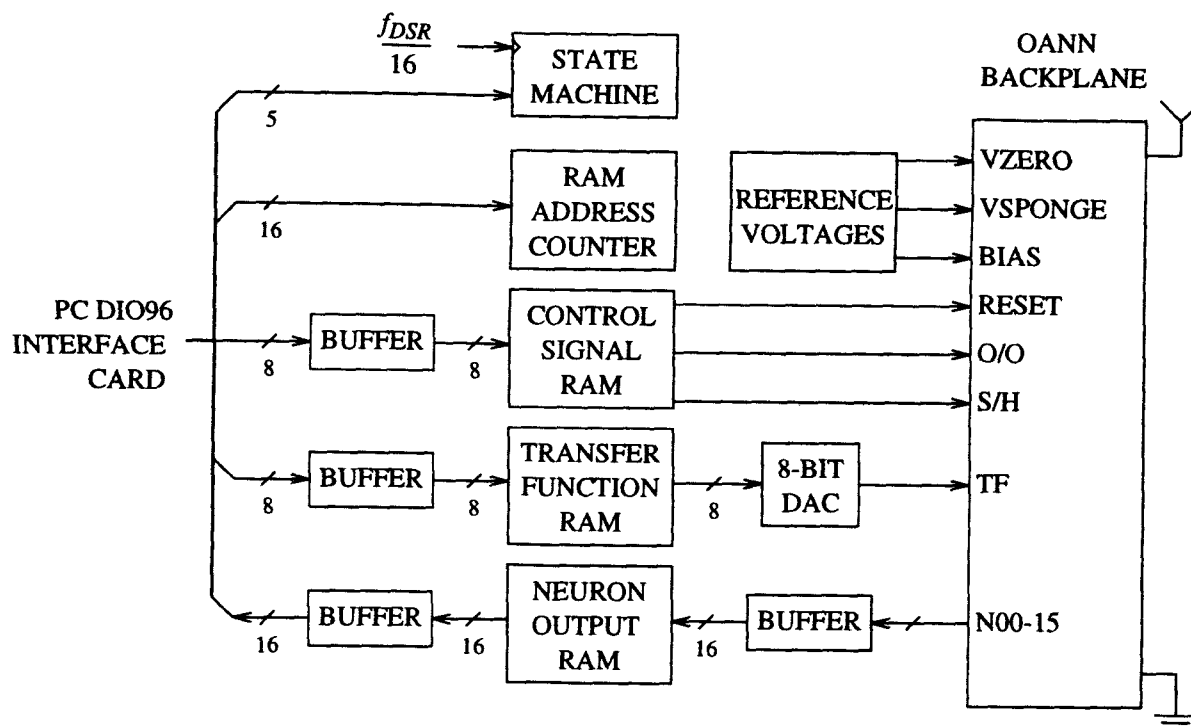
Figure 6–15: Block diagram of the OANN backplane interface (OannInt). For clarity, the various control and RAM addressing signals have not been drawn.

I designed the interface to be as flexible as possible. The RESET, O/O, and S/H control signals are all stored as a RAM look-up table so that the control sequence can be redefined easily. Similarly, the transfer function RAM supplies the 8-bit DAC with a sequence of bytes corresponding to the desired transfer function TF. Simple triangular ramp or variable-gain sigmoidal functions can be produced simply by loading the RAM with the appropriate bit-pattern. To simplify backplane characterisation, TF is always a simple single-sided ramp.

The DAC that converts the byte sequence in the transfer function RAM into TF can only provide a full scale deflection of 0 to 4 V, so the full range of the 2-stage comparator cannot be verified. However, as the neuron current-routing circuitry can only operate up to $V_{ACT} = 4$ V, the extra 1 V (1–5 V) range of the comparator is actually redundant anyway.

After the FastInt has been loaded with the appropriate bit-planes and delay counts, and the OannInt has been configured to generate the control signals at the appropriate times, both interface state machines are triggered to cycle once through their programmed sequences.

As with the counters in the FastInt, the RAM address counter is configured as a programmable down counter. When it reaches 255, both the transfer function RAM and the neuron output RAM are enabled. While the TF signal is ramped, the neuron states are read into the neuron output RAM until the address counter reaches zero. The state machines then enter their idle states until the next trigger pulse from the PC. The sequence in the RAM corresponding to the neuron states can then be read by the PC.

The voltage references VZERO, VSPONGE, and BIAS can be adjusted using potentiometers. The reference VSPONGE must be able to sink as well as source current: a LM324 op-amp configured as a buffer connected to the VSPONGE potentiometer provides this capability.

For all the experiments described in the next section, the OannInt state machine and RAM counters are clocked at a frequency of 1 MHz, or $f_{DSR}/16$ (where the clocking frequency of the FastInt state machine and SBS256 data shift registers, $f_{DSR} = 16$ MHz). This helps provide a good matching between the programmable counters on both the FastInt and OannInt, and permits a wide range of frame rates and S/H pulses widths to be selected. Because of this, the transfer function TF is defined over a 256 μs range, and thus the neuron pulse width outputs will lie between 0 and 256 μs. However, this means that the neuron outputs are not fully compatible with the EPSILON chip set, which require 0–20 μs input signals. This is not a limitation of the OANN backplane – the OannInt could be modified to generate the required EPSILON pulse widths, but I have not done so, purely for the convenience of having the two interfaces well matched.

6.7 Testing and characterisation

Ten die, already bonded into 40-pin chip packages, were supplied by EUROCHIP. Of these, three had been damaged during bonding. Initial visual inspection of the candidate chips also revealed that there were significant residue deposits on many of the neuron photodiodes (see Fig. 6–16). This may have been caused by an incomplete cleaning step during the fabrication sequence. Although initial electrical testing showed that the residue did not short the photodiodes to the neighbouring buslines, its opacity would obviously affect the intensity of the light actually hitting the active area. Chip #1 was least affected by this problem, and so was used in all the *optical* characterisation experiments described below.
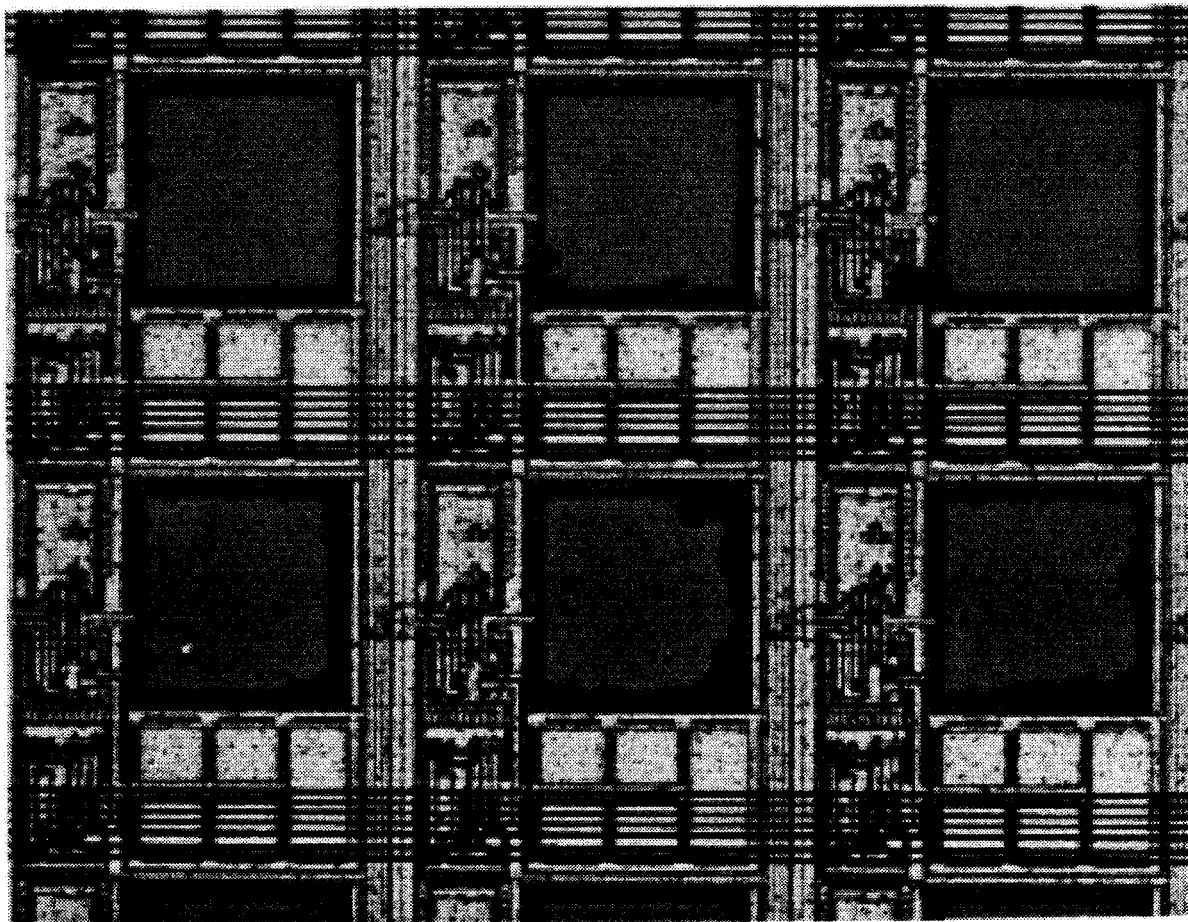
Figure 6–16: Photomicrograph showing the residue on the photodiodes on chip #4. This may have been caused by an incomplete cleaning step during the the fabrication sequence at AMS.

6.7.1   Pulse width output versus V_{ACT}

The two-stage comparator of the neuron front-end encodes the voltage V_{ACT} on the activity capacitor into a pulse width signal. By pulsing the RESET signal, V_{ACT} can be set to the voltage reference VZERO. Fig. 6–17 shows the pulse width output with respect to V_{ACT}, where each data point represents the averaged value of all 16 neurons on chip #2. The standard deviation for each average is less than 1 μs so all the neurons generate essentially the same output for a given V_{ACT}. The response is also approximately linear over the range tested.
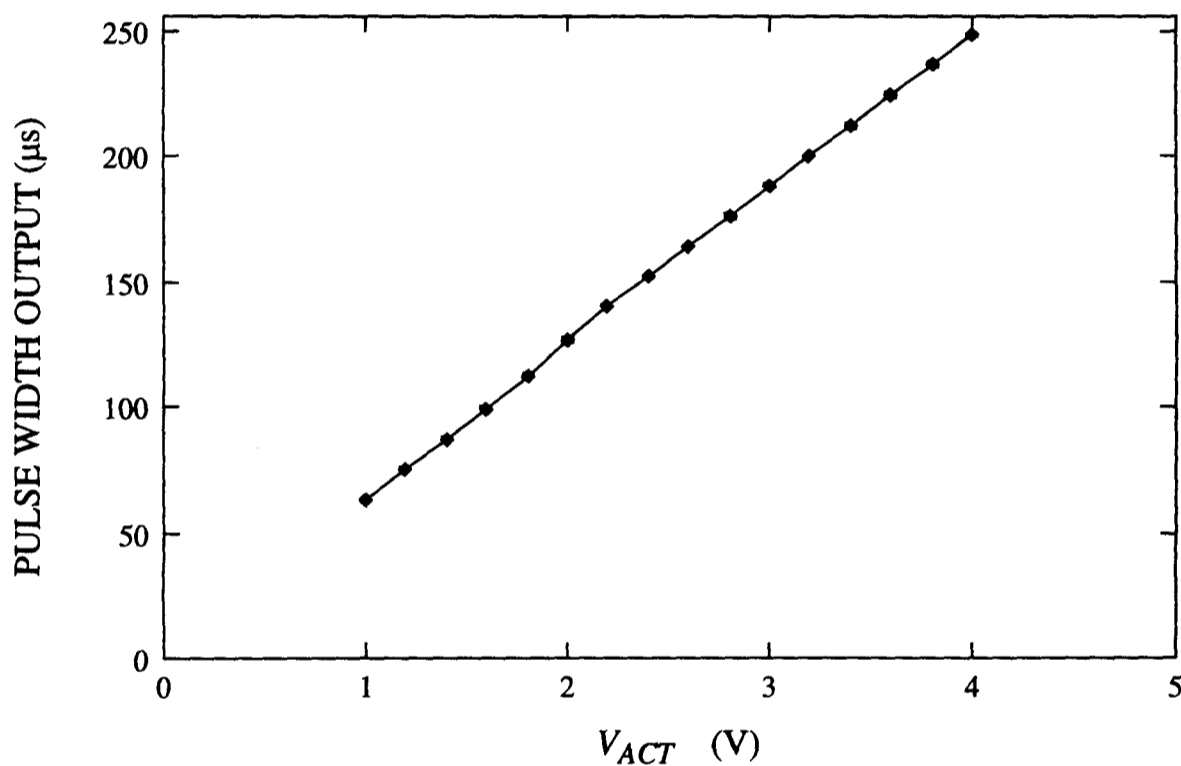
Figure 6–17: Pulse width output with respect to neuron activation voltage V_{ACT}. Here each data point represents the averaged value of all 16 neurons on chip #2. The standard deviation for each average is less than 1 μs.

6.7.2   Sample and hold characteristics

For a given intensity of light, the duration of the S/H signal determines how much current is sourced onto, or sunk from, from the activity capacitor. The backplane was placed in a simple optical system (see Fig. 6–18) to investigate the effect of varying the S/H pulse width.

LIGHT SOURCE   P1   P2   OANN BACKPLANE
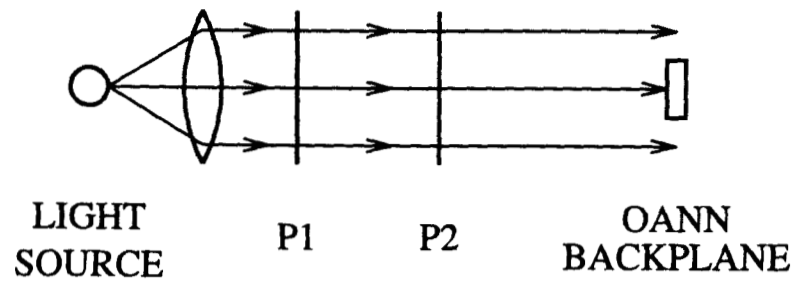
Figure 6–18: Simple optical system to investigate the current-router functionality. The polariser P2 can be rotated with respect to P1 to attenuate the incoming light according to Malus' Law.

The white light source is partially collimated by the lens and then passes through the two polarisers P1 and P2. By rotating P2 with respect to P1 (or vice versa), the transmitted light can be attenuated according to Malus' Law [91]. The top-left alignment photodiode APN_0 at the edge of the neuron array was used to help set P2 to give linearly increasing light intensities for each S/H sweep. The values at the end of each trace correspond to the voltage across a 500 kΩ load resistor connected to APN_0.

Fig. 6–19 shows the pulse width output of neuron N10 on chip #1 with respect to the S/H pulse width for five linearly increasing light intensities, with O/O set low and high. The results show that the basic functionality of the back-end circuit is correct. There is a slight mismatch between sourcing current onto, and sinking current from, C_{ACT} for the same light intensity. This is probably due to the imperfect mirroring capability of the cascode current mirror.

6.7.3 Effects of process variation

The process variation inherent in all CMOS processes was expected to affect the uniformity of response of neurons across the array. Fig. 6–20 illustrates the variation across chip #1 for uniform illumination. Here, the neurons were reset to give a 143 μs pulse width output, and then the S/H signal was pulsed for 128 μs. After this, the average neuron output was 228 μs, while the minimum
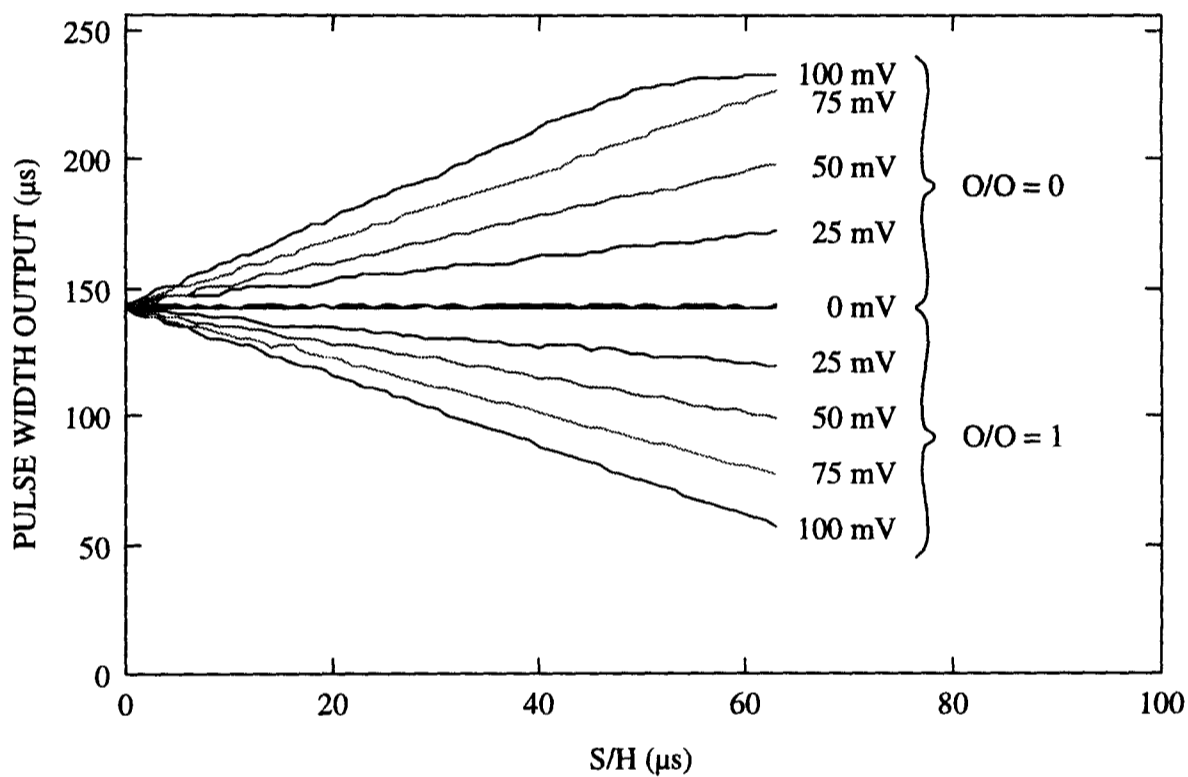
Figure 6–19: Pulse width output of neuron N10 versus S/H for several different intensities, with O/O set low and high. The values at the edge of each trace correspond to the voltage across a 500 kΩ load resistor connected to the alignment photodiode APN_0.

was 220 μs and the maximum was 234 μs. This corresponds to a ±9% variation on the average if 143 μs is taken as the zero point. Variations in the input photodiodes, the cascode current mirror transistors, and the transistor-capacitors, all contribute to this variation.
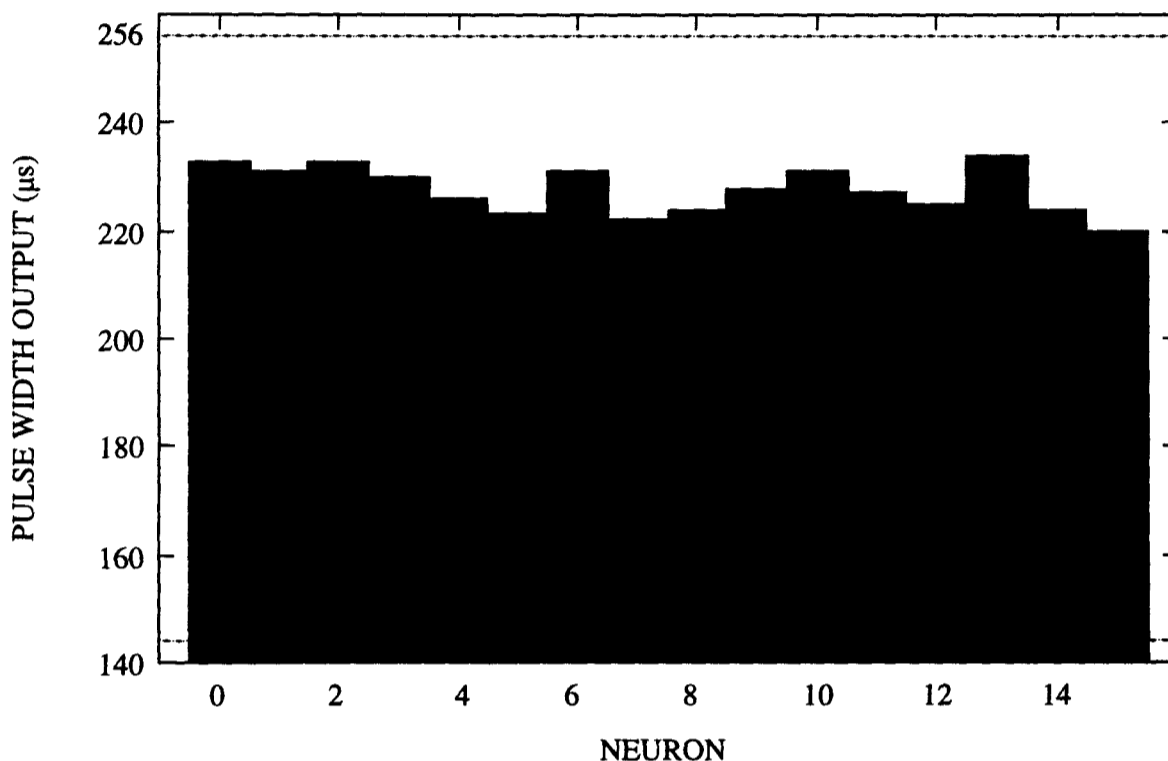
Figure 6-20: Neuron pulse width output variation on chip #1 while under uniform illumination. Here the neurons were reset to give a 143 μs output, and then S/H was pulsed for 128 μs.

6.7.4 Cascading the SBS256 and OANN

Fig. 6-21 shows the experimental set-up used to investigate how the SBS256 SLM and OANN backplane perform in the optical system described in Section 6.2.3.

The He:Ne laser (wavelength = 632.8 nm) is rated at 10 mW, but is almost 10 years old, and actually generates just over 1 mW of optical radiation. A planarised SBS256 SLM (device 2#11) is used to ensure that enough light passes through the system to fall onto the neuron photodiodes. While an 'All ON' pattern was displayed on the SLM, S/H was pulsed for 64 μs (with VZERO set to give an initial 142 μs neuron pulse width signal). The resulting pulse widths are
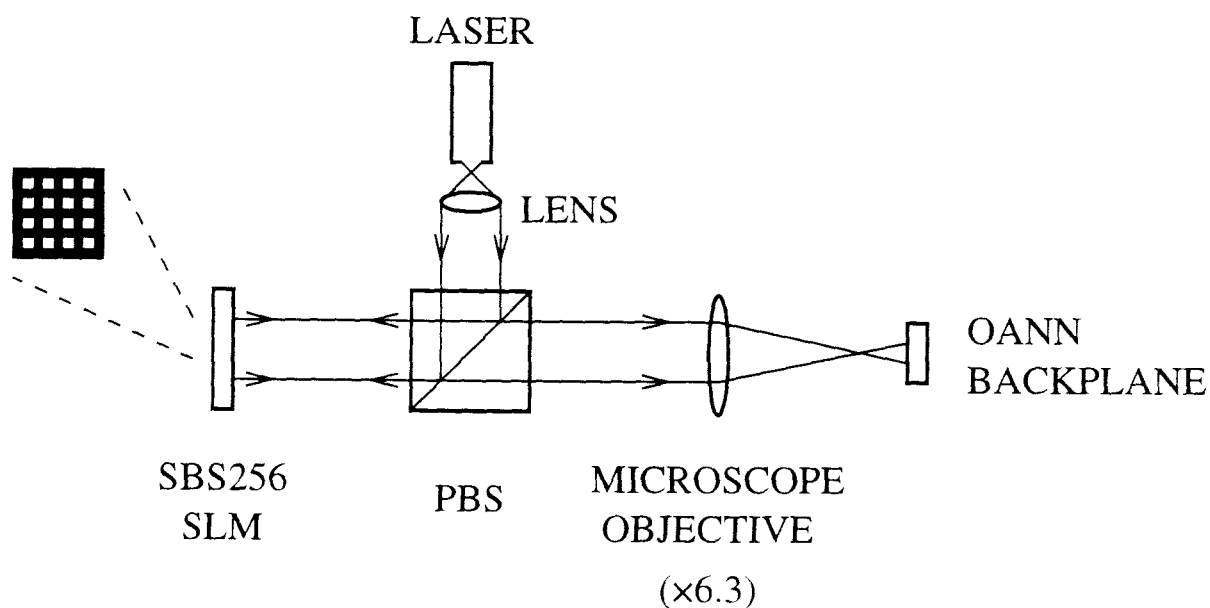
(a)
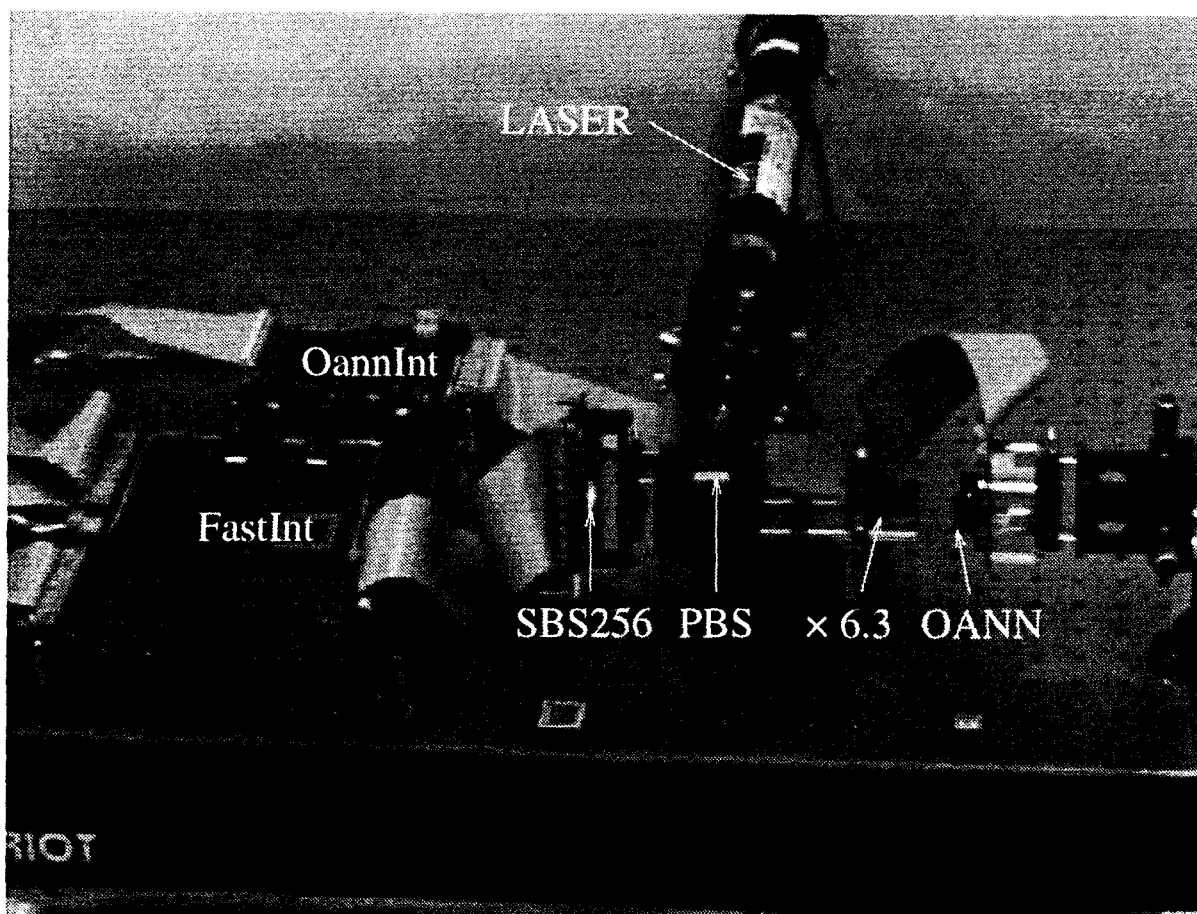
(b)

Figure 6-21: (a) The simplified optical system described in Section 6.2.3, and (b) a photograph of the actual system and interfaces.

superimposed onto the SLM fringe pattern shown in Fig. 6-22. As can be seen, the fringes present on the SLM severely affect the amount of light each neuron photodiode receives. and so affect the pulse width outputs. Note the pulse width outputs match quite well with the fringe patterns.
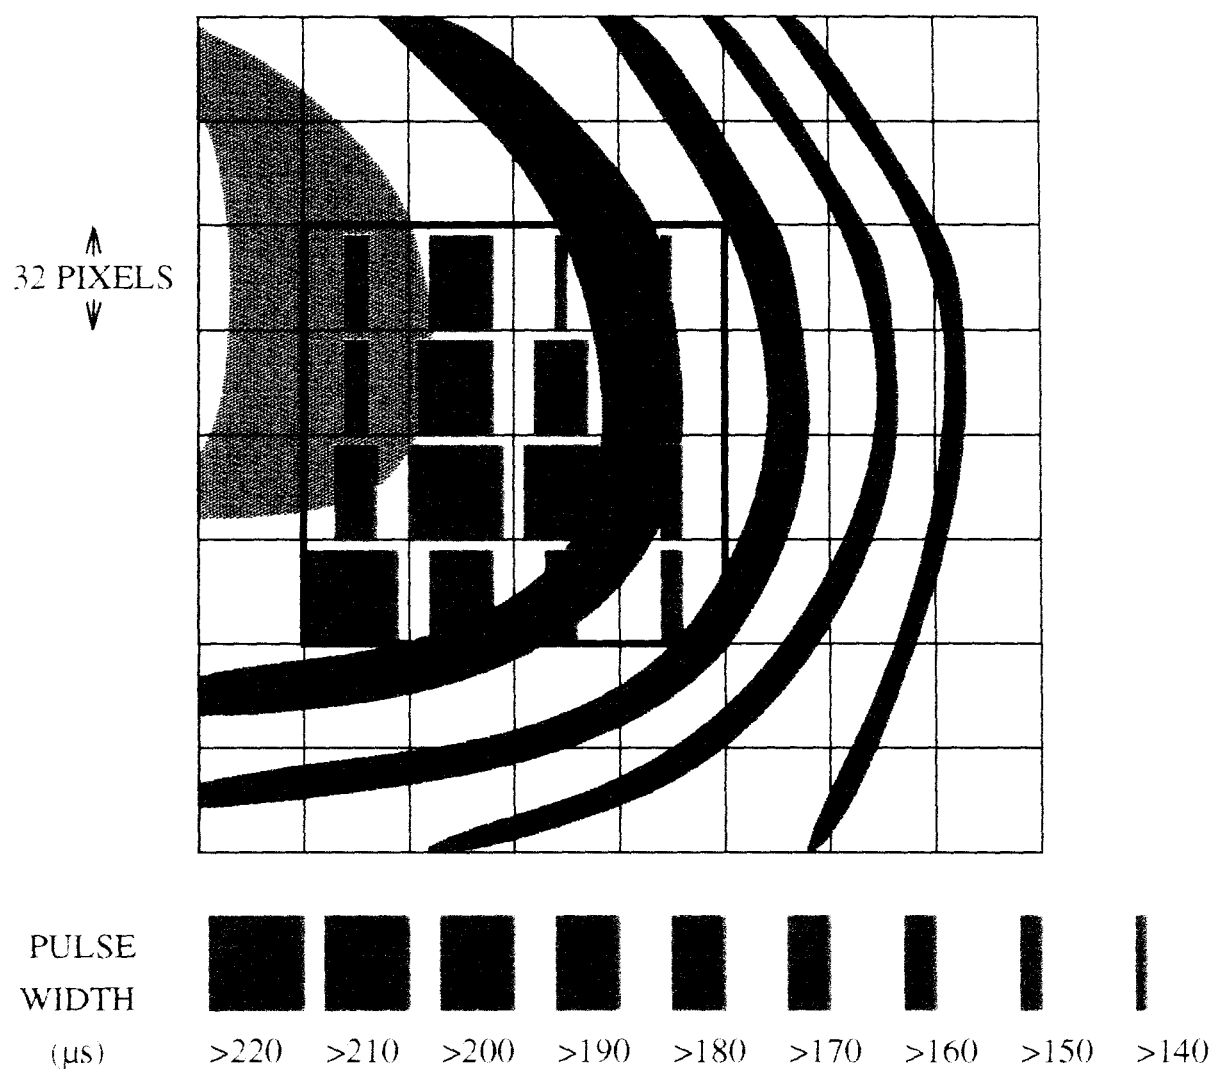
Figure 6–22: Effects of the SLM fringes (device 2#11) on the neuron pulse width outputs (superimposed). Note that the fringes drawn are only an approximate representation.

6.7.5   Neuron characterisation

The results above show that the neuron responses are affected by the non-uniformities on the SLM device. Nevertheless, by selecting a neuron that is not affected by a fringe. its response to the 2's complement drive scheme can be investigated.

Using a 6-bit 2's complement drive scheme, 64 grey levels can be realised (-32 to +31). Fig. 6-23 shows the transfer characteristics for five binary input test patterns with the weights swept from -32 to +31. The test patterns are 8 × 8 rather than 16 × 16 because the software controller program on the PC had a limited amount of memory available for data storage. For each value on a weight sweep, every pixel in the input pattern is multiplied by the weight value. These *IW* multiplications are calculated in software as explained in Section 6.2.3. To reduce alignment requirements in the optical system, the resulting 8 × 8 *IW* pattern is tiled to make a 32 × 32 pattern, and then encoded into the six 2's complement bit-planes These are then down-loaded to the FastInt RAM store, and a forward pass calculation is performed.
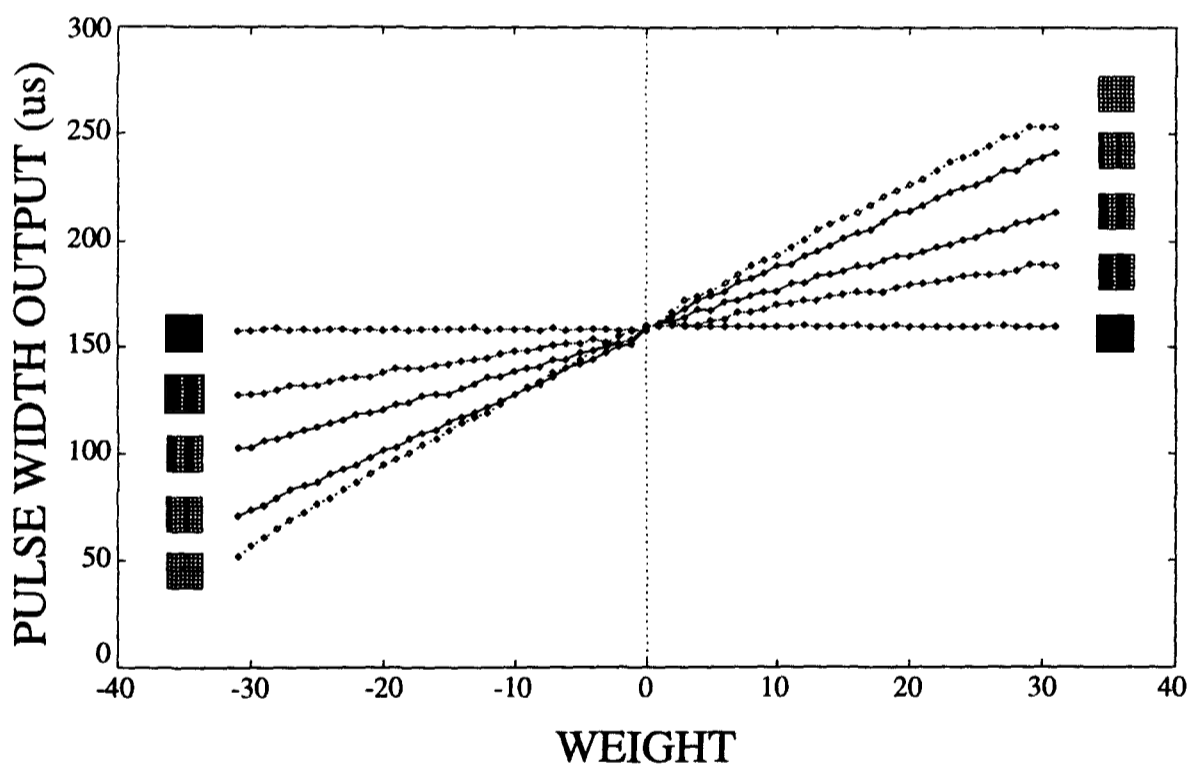
Figure 6–23: Transfer characteristics for neuron N00 on chip #1. For each of the five input test patterns, the trace corresponds to a sweep of all the possible weights using a 6-bit 2's complement drive scheme.

Fig. 6-23 shows the transfer responses of neuron N00 on chip #1 for the test patterns. Here the S/H signal is pulsed for 4 μs in the least significant bit-plane, 8 μs in the second least significant bit-plane, and so on. Using six bit-planes, the most positive weight corresponds to a total of 64 μs. The resulting pulse

width output for this S/H pulse is slightly different (252 μs as opposed to 225 μs) to that obtained in the previous experiment because VZERO has been adjusted slightly to give a different starting point.

The response is fairly linear with respect to the number of neurons on in the input pattern. Defects in the FLC layer on the SLM are likely to be the main cause of any nonlinearity. As with the sample and hold experiment described above in Section 6.7.2, there is a slight mismatch between sourcing current onto, and sinking current from, the activity capacitor. Again, this is a result of mismatching between the current generated in the photodiode, and the copied current in the current mirror. This could be compensated for by adjusting the S/H pulse in the sign subframe so that it is slightly longer, a refinement that has not been carried out here.

6.8 Summary

In this chapter I have presented the OANN prototype neural-detector backplane. The device has a 4 $\times$ 4 array of neurons that can sample the temporally multiplexed grey-scale generated by the SBS256 SLM. At the start of the chapter, I discussed requirements for an optoelectronic system incorporating the SBS256 SLM and the OANN backplane. After presenting a technique to reduce significantly photo-induced charge leakage from storage capacitors, I described the neuron circuit in detail. Test and characterisations results were then presented. I demonstrated a 2's complement temporally multiplexed SBS256 drive scheme to generate positive and negative activations that the neuron array could sample. However, the non-uniformity of the SLM device affected the uniformity of response across the neuron array.

Chapter 7

Discussion and conclusions

In this chapter I evaluate the performance of the two IC designs presented in the preceding chapters. I also discuss possible improvements to the devices, and give pointers to possible areas of future work.

7.1 SBS256 SLM

In Chapters 4 and 5, I presented the SBS256 SLM.

7.1.1 Backplane evaluation

The backplane circuits function as designed: the pixel exhibits the expected storage and XOR-gate functionality, while the peripheral addressing circuits load the image bit-planes onto the array as expected.

The current spikes generated by toggling the global clock charge-balancing signal CK, are approximately 40–50% of the value estimated by the simple superposition of current spike predicted by HSPICE simulations for a single pixel. Although lower than estimated, they illustrated the need to ensure that transistor sizing and addressing architecture must be taken into account when designing a large array SLM backplane.

The bit-plane scan time of the array is limited by the rise/fall times of the data shift register's non-overlapping clock buslines. HSPICE simulations predicted that the data shift register would operate at up to 25 MHz: operation at 24 MHz was verified.

On the probed wafers, the percentage of fully working die is 44%, while another 29% will function as partially working SLMs. These yield figures are extremely encouraging, considering the high packing density of the transistor circuits and buslines running through the array. The busline yield circuits suggested by I. Underwood, proved to be a simple but effective means of determining the yield failure modes of the backplane.

The backplane was also used to help develop the post-processing chemical-mechanical planarisation technique being investigated by A. O'Hara. Preliminary results show that this technique can be successfully used to increase significantly the optical quality of SLM devices. It is particularly suitable for improving the fill-factor of devices such as the SBS256 that have several transistors and buslines in each pixel. Once the technique has been fully characterised, it is likely to become an integral part of SLM fabrication.

7.1.2 Optical evaluation

The device functions as an electrically-written, optically-read, memory device and exhibits the expected FLC optical response. Qualitative experiments show that data storage is robust under high read-beam intensities. The 4–8:1 pixel on/off contrast ratio is similar to that observed on other SLMs. This is likely to improve as better FLC alignment techniques are developed.

The asymmetric FLC switching times are curious but have been observed in other SLM devices and test cells. The relatively slow 155 ± 5 μs switch-on time limits the device's maximum frame rate to $\sim$2 kHz. Again, better alignment techniques and thinner cells (by using 1 μm spacer balls rather than the 2 μm balls) should improve the switching times, and therefore frame rate. For example,

if both the FLC switching times could be reduced to 40 μs, and data is scanned in at 24 MHz, the device could operate at a fully-charge balanced frame rate of 4 kHz (with the pattern appearing on a 50% duty cycle). The frame rate of present devices is sufficient to demonstrate temporally multiplexed grey-scale using the simple linear encoding scheme. The grey-scale rendition is approximately linear but can be affected by defects in the FLC.

The main limitation of the assembled devices used in this project is the non-uniformity of the FLC layer. This is caused by a combination of the warp introduced to the backplane by (1) the CMOS processing steps, and (2) the gluing procedure when the die are inserted into chip packages. When viewed under white light illumination, a series of approximately circular 'rainbow' fringe patterns is observed over the pixel array. These significantly reduce the optical quality of the device. The problem is also observed on the 16 × 16 and 176 × 176 devices but is much more noticeable on the large active area of the SBS256. At the time of writing, die-flattening experiments are being carried out by A. O'Hara.

7.1.3 Improvements

The Smectic Technology group at Thorn EMI's Central Research Laboratories (CRL) have recently procured a 25 wafer batch of SBS256 backplanes for commercial exploitation (see Appendix B). Fabrication at AMS should commence soon. To help improve the backplane yield, I have redesigned the METAL2 mask. Because the transient current spikes were lower than estimated, I have made the VDD and CK buslines slightly narrower. This increases the spacing between them and the adjacent /ENABLE lines to at least 2.6 μm (the minimum design rule spacing is 1.8 μm), so the number of busline shorts ought to decrease, thus increasing the yield.

As mentioned above, improvements in the device assembly techniques are required. Nevertheless, the results obtained so far with limited resources, equipment, and manpower are very encouraging. Possibly the most dramatic improvement in optical quality will come if and when the die-flattening experiments are

successful. The cell thickness can then be reduced to 1 μm, and so improvements in FLC switching times and alignment quality can be expected. These improvements, along with planarisation, should help the SBS256 become a powerful, high quality SLM, suitable for a variety of applications. For example, researchers on the SCIOS [90] collaborative project plan to use the device as a reconfigurable holographic router in an optical computing architecture, while researchers on the HICOPOS [41] project plan to use two SBS256 devices in a matched-filter correlator system.

7.1.4   A more compact SRAM-type design

Now that planarisation is available, some of the constraints on the pixel layout can be relaxed. The most obvious is that the electrode mirror can be reduced so that it can just accommodate the various contacts and vias between the circuit transistors and the top level metal. The single pixel design, where every pixel is oriented identically, can be replaced by the 'Quad' pixel design illustrated in Fig. 7-1. Here the VDD, CK and GND buslines can be shared between adjacent pixels, thus saving on area. All the pMOS transistors can be placed in the same n-well, again saving area by eliminating some of the spaces required between n-wells and nMOS transistors. These modifications should reduce a single SRAM-XOR type pixel to about 30 × 30 μm^2 using the same 1.2 μm CMOS process (that is, the pixel pitch is about 25 times the minimum transistor size). This could either increase the operating frequency of the backplane because of shorter data shift register clocking buslines, or increase the pixel count if the the same active area is used.

7.1.5   SRAM versus DRAM

For a given CMOS process, the SRAM-XOR pixel will occupy more area than the single-transistor DRAM-type. This results in either a larger active area for a given number of pixels, or a smaller number of pixels for a given active area.
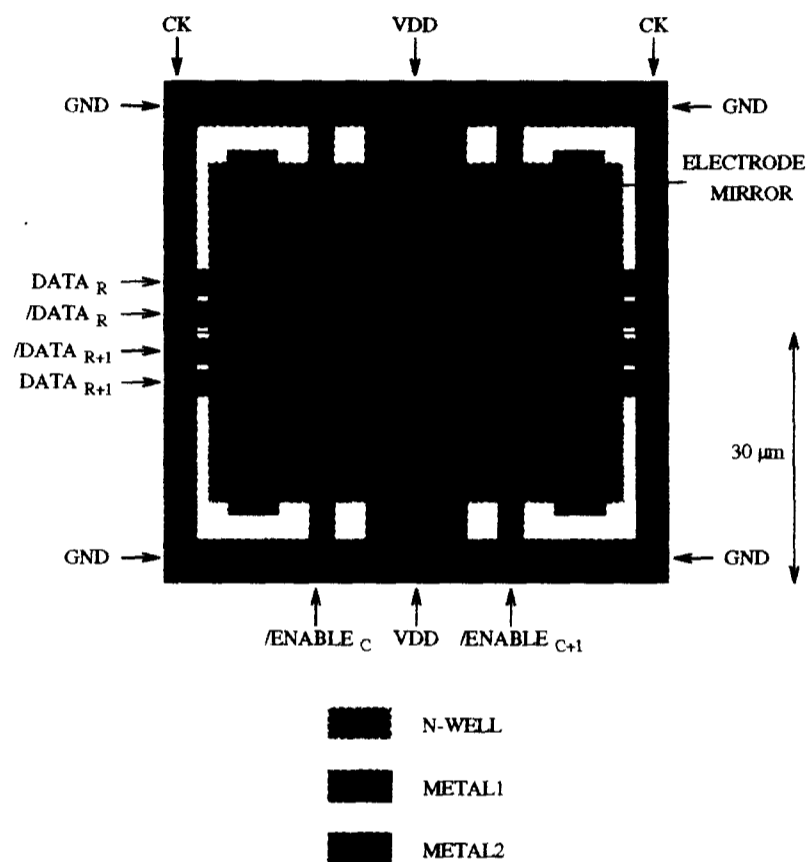
Figure 7-1: The 'Quad' pixel layout. Here the pixels are no longer oriented identically. This permits (1) the VDD, CK, and GND lines to be shared between adjacent pixels, and (2) all the pMOS transistors for four pixels to be placed in one n-well. With planarisation the electrode mirrors can also be much smaller. For a 1.2 μm CMOS technology, the 'Quad' pixel should fit into an area of about 60 × 60 μm^2, so the pixel pitch is about 25 times the minimum transistor size.

For example, with a 1.2 μm CMOS process similar to that used for the SBS256, McKnight *et al* [66] implemented a 256 $\times$ 256 array using a 20.6 $\times$ 20.6 μm^2 single-transistor pixel. Using this size of pixel, a 512 $\times$ 512 device would be realisable on the same active area as the SBS256. However, the SRAM-XOR pixel has improved data storage and FLC drive capabilities. Including the SRAM latch overcomes the photo-induced charge-leakage problems associated with the DRAM-type – an important factor for a device that operates in an optical system ! With an XOR-gate in each pixel, a charge balancing scheme can be used that does not result in an inverse pattern appearing on the device. Therefore the SBS256 does not require a pulsed light source to view it.

The FLC switching time is inversely proportional the FLC's spontaneous polarisation P_S (in nC/cm^2), so faster FLC materials generally have higher values. In a DRAM-type device, sufficient charge must be dumped onto the data storage capacitors to switch the overlying FLC. The active-addressing technique permits the array to be addressed very rapidly, ideally within one FLC switching time. However, for FLC mixtures with P_S greater than 40 nC/cm^2, the charge dumped onto the capacitor may not be enough to switch the FLC so the access transistors must remain on until the FLC switches. Therefore, the array takes N FLC switching times to address, where N is the number of columns, so the potential frame rate is severely affected. In the SRAM-XOR pixel, the electrode mirror is actively driven via circuit transistors by the power rails, so there is an effectively unlimited amount of charge available to switch the FLC. The pixel design is therefore more suitable for driving faster FLC mixtures.

As process geometries shrink, backplane operating voltages will tend to decrease to limit the effects of hot-carrier degradation [25]. This will also tend to favour the SRAM-XOR type over the DRAM-type, again because the electrode mirror is actively driven.

Both types of device are suitable for a variety applications, but the DRAM-type device is favoured where high resolution is the most important system requirement. The SRAM-XOR type device is particularly suitable for systems

where the read-beam is intense, say at the input stage of a cascaded SLM system, or in an application where a pulsed light source cannot be used. The synchronous FLC switching available with the pixel design may also be useful in applications where a whole image must appear on the device simultaneously, with no addressing skew. The static nature of the SBS256's data storage permits sections of the array to be addressed while other sections retain a previous pattern. This permits architectures such as the joint transform correlator (see Section 3.2.2), to be efficiently implemented using the device. It should be noted that both types of device show the desired pattern on a 50% duty cycle. Some applications may require a 100% duty cycle, so a device based on the three-transistor pixel [94] may be more suitable.

Summing up, SRAM-XOR type devices offer improved functionality over the DRAM-type and are still competitive in terms of pixel count compared with current DRAM-type devices. However, using the 1.2 μm process with a 10.24 $\times$ 10.24 mm^2 active array, the DRAM-type has the potential to increase to about 750 $\times$ 750 if a 13 μm pixel pitch can be used, while the SRAM-type could probably increase to 350 $\times$ 350, if a 30 μm pitch pixel based on the 'Quad' design is used. Obviously with a smaller geometry process, such as a 0.8 or 0.6 μm one, a higher pixel density could be achieved for both types but, at present, these processes are not available on large area die.

7.2 OANN neural-detector backplane

In chapter 6, I presented OANN – a prototype neural-detector backplane.

7.2.1 OANN evaluation

The OANN backplane consists of a 4 $\times$ 4 array of optical-in, electronic-out, processing units, or neurons. It was designed to sample the grey-scale generated

by the SBS256 SLM. By using a 2's complement grey-scale encoding scheme, the neurons can respond to positive and negative activations.

Early in the backplane design cycle, I identified photo-induced charge leakage as a potentially severe problem for the activity storage capacitors. Therefore I included a simple protection technique that had been verified using a test circuit on the SBS256 (see Section 6.3). This n-well protection circuit reduced charge leakage by a factor of about 2000, making charge leakage almost insignificant during the grey-scale frame rate of the SLM (typically 2–8 ms). The technique is also suitable for improving charge storage in standard electronic circuit designs where optical radiation is not present, that is, the structure also reduces the number of thermally generated minority carriers close to the access transistors. The n-well ring does result in a small area increase for the storage circuit; however, this could probably be offset by using a smaller storage capacitor.

The neuron circuits function as designed. The input photodiode and current mirror sub-circuit can respond linearly to low light intensities, with currents as low as 50 nA being generated and copied. The current routing circuits can source current onto, or sink current from, the activity capacitor; however, a charge-sharing event can occur if an inappropriate drive scheme is used. The 2-stage comparator generates a pulse width output signal dependent on the voltage stored on the activity capacitor.

A 2's complement grey-scale drive scheme was successfully demonstrated using six bit-planes to generate the 64 levels of activation – 32 positive, and 32 negative. The neuron circuit responded in the appropriate manner. The drive scheme can be optimised to take account of the number of grey levels weights required for a particular application, and the light intensities present in the optical system.

The rather low power of the laser used in the experiments resulted in quite long S/H integration times (50–100 μs), but these match satisfactorily with the bit-plane frame time available from the SLM devices. A more powerful laser would be required for a cascaded SLM system where the multiplication is performed in the optical domain.

The variation across the neuron array is ± 9%. This is due to process variation and is inherent in analogue integrated circuit implementations. When the backplane is used in an optoelectronic ANN system, this variation could be taken into account by including the system in the weight-adjusting learning loop. However, the gross variation resulting from the fringes on the SLM used to characterise the backplane could not be overcome by 'system in the loop' learning: an optically-flat SLM is required.

7.2.2 Improvements and future work

The most obvious requirement for satisfactory operation of a simple ANN system, is to have an optically-flat SBS256 device. As mentioned earlier, flat devices should be available soon.

In the longer term, the system proposed in Section 6.2.1 is feasible. A SBS256 SLM could be used for the input plane, but a device such as the Isophote is preferable because it can provide a coherent image directly from the 'real-world'. Some Fourier Optics preprocessing could also be included. With a computer generated fan-out hologram and a fan-in lenslet array, a large fraction of the SBS256 pixel count could be utilised in the weight plane. This represents up to 65,536 interconnection weights. Note that the maximum size of the input-to-hidden layer is limited by the number of pixels on the weight SLM. Therefore, there must be a trade-off between the numbers of neurons in the input and hidden layers. For example, if a 64 × 64 input image is required, the hidden layer can be a 4 × 4 array, while for a 32 × 32 image, the hidden layer can be 8 × 8. From these figures, a higher pixel count weight SLM device may be required for some classification tasks.

The neuron circuit could be redesigned to increase its frequency response, and the neuron array could be increased to an 8 × 8 array without any need to add decoder addressing circuitry. Sixty-four output neurons would map well to the EPSILON pulse-stream neural chips, which could implement the hidden to output layer of a multi-layer perceptron ANN. A system such as this with, say

32 × 32 input, 64 hidden, and 64 output neurons, could perform useful image classification tasks at frames rates of 500 Hz (assuming the SBS256 can operate at 4 kHz, and eight bit-planes are required).

Summing up, I demonstrated a prototype optical-in, electronic-out neuron device. It is suitable for optoelectronic ANN networks architectures where high fan-in is required, and further layers can implemented using electronic pulse-stream neural chips. Of course, the neuron circuit design could be extended to generate an optical output for further optical processing, by redesigning it to be incorporated into a FLCOS SLM (see Appendix A for a description of a test pixel that I included on the SBS256 – it converts an analogue voltage into an optical pulse width signal).

7.3 Conclusions

The technology of ferroelectric liquid crystal over silicon has much to offer the fields of optical image processing and miniature displays systems. It has advanced rapidly over the past eight years or so, and will continue doing so for the foreseeable future. Experience has given researchers a good 'feel' for the technology's strengths and limitations, thereby allowing systems researchers to begin designing systems that take advantage of the strengths, while the device researchers address the performance limitations of present devices. The technology is also amenable to high-volume production, and so costs should gradually decrease if sufficient quantities are manufactured and sold. Devices such as the SBS256 presented in this thesis, and the 176 × 176 DRAM are now commercially available. This allows systems researchers to incorporate them into their systems as 'off-the'shelf' components.

In the field of artificial neural networks, the optoelectronic approach should feature prominently in image recognition applications, where the input exists in the optical domain. The increasing electronic and optical bandwidths of new

FLCOS SLMs devices should permit the implementation of high-speed opto-electronic ANNs, with many hundreds of thousands of adaptive interconnection weights and hundreds, possibly thousands, of neurons. These systems should be capable of performing sophisticated image processing tasks.

Bibliography

[1] G. A. Allen, 1992. A private communication.

[2] P. E. Allen and D. R. Holberg. *CMOS analogue circuit design.* Holt, Rinehart, and Winston Inc, 1987.

[3] P. M. Alt and P. Pleshko. Scanning limitations of liquid crystal displays. *IEEE Transactions on Electron Devices,* ED-21(2):146-155, 1974.

[4] G. Anderson, I. Dahl, P. Keller W. Kuczynski, S. T. Lagerwell, K. Sharp, and B. Stebler. Submicrosecond electrooptic switching in the liquid crystal smectic *a* phase. *Applied Physics Letters,* 51(9):640-642, 1987.

[5] D. Armitage and D. K. Kinell. Liquid crystal integrated silicon spatial light modulator. *Applied Optics,* 31(20):3945-3949, 1992.

[6] S. Bian, J. Li, and K. Xu. Optical implementation of Hopfield neural network with programmable interconnect matrix. In *Proceedings of the SPIE,* volume 1709, pages 682-687, 1992.

[7] J. R. Black. Electromigration - a brief survey and some recent results. *IEEE Transactions on Electron Devices,* ED-16(4):338-347, 1969.

[8] M. F. Bone, D. Coates, W. A. Crossland, P. Gunn, and P. W. Ross. Ferro-electric liquid crystal display capable of video line address times. *Displays - Technology and Applications,* 8(3):115-118, 1987.

[9] R. M. Boysel. A 128 × 128 frame-addressed deformable mirror spatial light modulator. *Optical Engineering,* 30(9):1422-1427, 1991.

[10] D. C. Burns, I. Underwood, D. G. Vass, and A. O'Hara. Electronically addressed ferroelectric liquid crystal over silicon spatial light modulators for optical computing. In *The International Conference on Optical Computing Technical Digest,* pages 27-28, Edinburgh, UK, 22-25 August 1994.

[11] L. M. F. Chirovsky, M. W. Focht, J. M. Freund, G. D. Guth, R. E. Leiben-guth, G. J. Przybylek, L. E. Smith, L. A. D'Asaro, A. L. Lentine, R. A. Novonty, and D. B. Bucholtz. Large arrays of symmetric self-electro-optic effect devices. In H. S. Hinton and J. W. Goodman, editors, *OSA Proceedings of the Topical Meeting on Photonic Switching,* volume 8, pages 56-59, Washington DC, USA, 1991. OSA.

[12] S. Churcher. *VLSI neural networks for computer vision*. PhD thesis, University of Edinburgh, 1993.

[13] N. A. Clark and S. T. Lagerwell. Submicrosecond bistable electro-optic switching in liquid crystals. *Applied Physics Letters*, 36(11):899–901, 1980.

[14] N. A. Clark and S. T. Lagerwell. Surface stabilised ferroelectric liquid crystal electrooptics: new multistate structures and devices. *Ferroelectrics*, 59:25–67, 1984.

[15] N. Collings, W. A. Crossland, P. J. Ayliffe, D. G. Vass, and I. Underwood. Evolutionary development of advanced liquid crystal spatial light modulators. *Applied Optics*, 28(22):4740–4747, 1989.

[16] A. B. Davey and W. A. Crossland. Potential and limitations of the electroclinic effect in device applications. *Ferroelectrics*, 144:101–112, 1991.

[17] A. R. Dias, R. F. Kalman, L. M. Woody, and A. A. Sawchuk. Fiber-optic crossbar switch with broadcast capability. *Optical Engineering*, 27(11):955–960, 1988.

[18] U. Efron. Spatial light modulators and applications for optical information processing. In *Proceedings of the Society of Photo-optical and Instrumentation Engineers*, volume 960, pages 180–203, 1988.

[19] S. Esener. Recent developments on Si/PLZT smart pixel modulators. In *IEEE/LEOS Summer Topical Meeting on Smart Pixels*, pages 85–88, August 1992.

[20] S. C. Esener, J. H. Wang, T. J. Drabik, M. A. Title, and S. H. Lee. One-dimensional silicon/PLZT spatial light modulators. *Optical Engineering*, 26(5):406–413, 1987.

[21] L. A. Beresenev *et al.* Electro-optic response of a thin layer of a ferroelectric liquid crystal with a small helicoid pitch and a high spontaneous polarisation. *Soviet Technical Physics Letters*, 14:116–118, 1988.

[22] N. H. Farhat, D. Psaltis, A. Prata, and E. Peak. Optical implementation of the Hopfield model. *Applied Optics*, 24(10):1469–1475, 1985.

[23] A. D. Fisher and J. N. Lee. The current status of two-dimensional spatial light modulator technology. In H. H. Szu, editor, *Optical and Hybrid Computation*, volume 634, pages 352–370, Bellingham, WA, 1986. SPIE.

[24] S. M. Fluxman. Integrated active matrix liquid crystal displays. *GEC Journal of Research*, 11(1):2–11, 1993.

[25] D. P. Foty and E. J. Nowak. MOSFET technology for low-voltage/low-power applications. *IEEE Micro*, 14(3):68–77, 1994.

[26] C. L. Giles and T. Maxwell. Learning, invariance and generalisation in higher-order neural networks. *Applied Optics*, 26(23):4972, 1987.

[27] L. A. Glasser and D. W. Dobberpuhl. *The design and analysis of VLSI circuits*. Addison-Wesley, Massachusetts, 1985.

[28] Austria Mikro Systeme International GmbH. Schoß Premstätten, 8141 Unterpremstätten, Austria.

[29] J. W. Goodman, A. R. Dias, and L. M. Woody. Fully parallel high-speed incoherent optical method for performing discrete fourier transforms. *Optics Letters*, 2(1):1–3, 1978.

[30] D. J. Goodwill, F. A. P. Tooley, A. C. Walker, M. R. Taghizadeh, M. McElhinney, F. Pottier, C. R. Stanley, D. G. Vass, I. Underwood, M. W. G. Snook, M. H. Dunn, J. Hong, and B. D. Sinclair. InGaAs S-SEED and silicon CMOS smart pixels for 1047–1064 nm operation. In *IEEE/LEOS Summer Topical Meeting Digest on Smart Pixels*, pages 20–21, Lake Tahoe, USA, July 1994.

[31] J. Gourlay. *Ferroelectric liquid crystal spatial light modulators: devices and applications*. PhD thesis, The University of Edinburgh, 1994.

[32] J. Grinberg, A. Jacobsen, W. Bleha, L. Miller, L. Fraas, D. Boswell, and G. Myer. A new real time non-coherent to coherent light image converter: the hybrid field effect liquid crystal light valve. *Optical Engineering*, 14:217–225, 1975.

[33] R. D. Guenther. *Modern Optics*. John Wiley & sons, 1990.

[34] A. Hamilton, S. Churcher, and A. F. Murray. Pulse stream VLSI circuits and systems: The EPSILON neural network chipset. *International Journal of Neural Systems*, 4(4):395–406, 1993.

[35] M. A. Handschy, H. Chase, L. K. Cotter, J. D. Cunningham, A. M. Pattee, T. J. Drabik, and S. D. Gaalema. One-transistor DRAM ferroelectric liquid crystal spatial light modulator. In *Proceedings of the Conference on Optical Pattern Recognition V*, volume 2237, pages 432–438. Society of Photo-Optical Instrumentation Engineers, April 1994.

[36] M. A. Handschy, L. K. Cotter, J. D. Cunningham, T. J. Drabik, and S. D. Gaalema. One-transistor DRAM FLC/VLSI SLM. In *Spatial Light Modulators and Applications Technical Digest*, volume 6, pages 14–17. Optical Society of America, Washington, DC, 1993.

[37] W. J. A. M. Hartmann. Ferroelectric liquid-crystal video display. *IEEE Transactions on Electron Devices*, 36(9):1895–1899, 1989.

[38] W. J. A. M. Hartmann. Ferroelectric liquid crystal displays for television application. *Ferroelectrics*, 122:1–26, 1991.

[39] S. Heddle and R. M. Sillitto. Reduction and removal of replication in an optical processor by randomisation of pixels in the fourier plane filter. *Journal of Modern Optics*, 40(2):299–313, 1993.

[40] J. Hertz, A. Krogh, and R. G. Palmer. *Introduction to the theory of neural computation*. Addison-Wesley, Massachusetts, 1991.

[41] HICOPOS Highly Integrated and Compact Optical Processor for On-board Systems.

[42] M. Holler, S. Tam, H. Castro, and R. Benson. An electrically trainable artificial neural network ETANN with 10240 "floating gate" synapses. In *Proceedings of the International Joint Conference on Neural Networks*, pages 191–196, 1989.

[43] J. J. Hopfield. Neural network and physical systems with emergent collective computational abilities. In *Proceedings of the National Academy of Science, USA*, volume 79, pages 2554–2558, 1982.

[44] W. E. Howard. Active matrix addressing techniques for displays. In *Proceeding of the Society for Information Display*, volume 27, pages 313–326, 1986.

[45] G. B. Jackson, A. Hamilton, and A. F. Murray. The EPSILON processor card: A framework for analogue neural computation. In *Proceedings of the Fourth International Conference on Microelectronics for Neural Networks and Fuzzy Systems*, pages 280–286, Torino, Italy, 26–28 September 1994.

[46] L. K. Cotter, T. J. Drabik, R. J. Dillon, and M. A. Handschy. Ferroelectric liquid crystal/silicon integrated circuit spatial light modulator. *Optics Letters*, 15(5):291–293, 1990.

[47] B. Javidi. Nonlinear joint power spectrum based optical correlation. *Applied Optics*, 28(12):2358–2367, 1989.

[48] M. S. Jin, J. H. Wang, V. Ozguz, and S. H. Lee. Monolithic integration of a silicon driver circuit onto a lead lanthanum zirconate titanate substrate for smart spatial light modulation. *Applied Optics*, 33(14):2842–2848, 1994.

[49] K. M. Johnson, D. J. McKnight, and I. Underwood. Smart spatial light modulators using liquid crystals on silicon. *Journal of Quantum Electronics*, 29(2):699–714, 1993.

[50] K. M. Johnson and G. Moddel. Motivations for using ferroelectric liquid crystal spatial light modulators in neurocomputing. *Applied Optics*, 28(22):4888–4899, 1989. And references therein.

[51] M. Killinger, J. L. De Bougrenet de la Tocnaye, and P. Cambon. Controlling the grey level capacity of a bistable FLC spatial light modulator. *Ferroelectrics*, 122:89–99, 1991.

[52] A. G. Kirk, W. A. Crossland, and T. J. Hall. A compact and scalable free-space optical crossbar. In *Proceedings of the Third International Conference on Holographic Systems, Components & Applications*, volume 1574, pages 137–141. IEE, London, UK, 16–18 September 1991.

[53] A. G. Kirk, S. Jamieson, H. Imam, and T. J. Hall. Experimental implementation of an optoelectronic matrix-matrix multiplier which incorporates holographic multiple imaging. In *Optical Computing & Processing*, 2(4), pages 293–304, 1992.

[54] S. T. Lagerwell, N. A. Clark, J. Dijon, and J. F. Clerc. Ferroelectric liquid crystals: the development of devices. *Ferroelectrics*, 94:3–62, 1989.

[55] J. Lazzaro, S. Ryckebusch, M. A. Mahowald, and C. A. Mead. Winner-take-all networks of $O(N)$ complexity. In D. Touretzky, editor, *Advances in Neural Information Processing Systems I*, pages 703–711. Kaufmann, Los Altos, California, 1989.

[56] A. L. Lentine, L. M. F. Chirovsky, L. A. D'Asaro, E. J. Laskowski, S. S. Pei, M. W. Focht, J. M. Freund, G. D. Guth, R. E Leibenguth, L. E. Smith, and T. K. Woodward. Field-effect-transistor self electro-optic effect device (FET-SEED) electrically addressed differential modulator array. *Applied Optics*, 33(14):2849–2855, May 1994.

[57] R. P. Lippmann. An introduction to computing with neural nets. *IEEE ASSP Magazine*, pages 4–22, April 1987.

[58] R. P. Lippmann. Pattern recognition using neural networks. *IEEE Communications Magazine*, pages 47–64, November 1989.

[59] Merck Ltd. Merck House, Poole, Dorset, England.

[60] A. Vander Lugt. Signal detection by complex spatial filtering. *IEEE Transactions on Information Theory*, IT–10:139–145, 1964.

[61] A. Vander Lugt. Coherent optical processing. *Proceedings of the IEEE*, 62(10):1300–1319, 1974.

[62] C. C. Mao and K. M. Johnson. Optoelectronic array that computes error and weight modification for a bipolar optical neural network. *Applied Optics*, 32(8):1290–1296, 1993.

[63] Y. Matsueda, Y. Ozawa, T. Nakamura, S. Takei, H. Kamakura, and N. Okamoto. 3.7in. HDTV poly-Si TFT-LCD light valve with fully integrated peripheral drivers. In *Proceedings of the 12th International Display Research Conference*, pages 561–564, Japan, 1992.

[64] J. L. McClelland and D. E. Rumelhart. *Parallel distributed processing*. MIT Press, Massachusetts, 1986.

[65] D. J. McKnight. *An electronically addressed spatial light modulator*. PhD thesis, University of Edinburgh, 1989.

[66] D. J. McKnight, K. M. Johnson, and R. A. Serati. Electrically addressed 256 by 256 liquid crystal on silicon spatial light modulator. *Optics Letters*, 18(24):2159–2161, 1993.

[67] D. J. McKnight, K. M. Johnson, and R. A. Serati. 256 × 256 liquid crystal on silicon spatial light modulator. *Applied Optics*, 33(14):2775–2783, 1994.

[68] D. J. McKnight, K. M. Johnson, G. D. Sharp, and R. A. Serati. Liquid crystal over silicon spatial light modulators. In *The International Conference on Optical Computing Technical Digest*, pages 25–26, Edinburgh, UK, 22–25 August 1994.

[69] M. A. Mignardi. Digital micromirror array for projection TV. *Solid State Technology*, 37(7):63–68, 1994.

[70] D. A. B. Miller. Quantum-well self electro-optic effect devices. *Optical and Quantum Electronics*, 22(5):S61–S98, 1990.

[71] G. Moddel, K. M. Johnson, W. Li, R. A. Rice, L. A. Pagano-Stauffer, and M. A. Handschy. High-speed binary optically addressed spatial light modulator. *Applied Physics Letters*, 55(6):537–539, 1989.

[72] D. M. Monro, J. A. Dallas, J. A. Nicholls, M. D. Cripps, and W. A. Crossland. An optically connected parallel machine. In *Optical Computing Topical Meeting Technical Digest*, volume 6, pages 280–283. Optical Society of America, Washington, DC, 1993.

[73] S. G. Morton. Electronic hardware implmentations. In *Handbook of neural computing applications*, chapter 16, pages 251–269. Academic Press Limited, London, UK, 1990.

[74] J. A. Neff, R. A. Athale, and S. H. Lee. Two-dimensional spatial light modulators: a tutorial. *Proceedings of the IEEE*, 78(5):826–855, 1990.

[75] D. C. O'Brien, R. J. Mears, T. D. Wilkinson, and W. A. Crossland. Dynamic holographic interconnects that use ferroelectric liquid crystal spatial light modulators. *Applied Optics*, 33(14):2795–2803, 1994.

[76] D. C. O'Brien, T. D. Wilkinson, and R. J. Mears. Programmable computer generated holograms with large space bandwidth product. In *Fourth International Conference on Holographic Systems, Components & Applications*, volume 379, pages 216–221, 13–15 September 1993.

[77] A. Ochoa and W. Davies. Latch-up control in CMOS integrated circuits. *IEEE Transactions on Nuclear Science*, NS-26(6):5065–5068, 1979.

[78] A. O'Hara, 1994. A private communication.

[79] A. O'Hara, I. Underwood, and D. G. Vass. Improved optical performance of silicon backplane spatial light modulators using chemical-mechanical polishing. In *Technical Digest of the Optical Fabrication and Testing Workshop*, volume 13, pages 86–89. (Optical Society of America, Washington, DC), 1994.

[80] J. Ohwada, M. Takabatake, Y. A. Ono, and Y. Nagae. Peripheral circuit integrated poly-Si TFT LCD with gray scale representation. In *Proceedings of the 8th International Display Research Conference*, pages 215–219, 1988.

[81] M. Oita, M. Takahashi, S. Tai, and K. Kyuma. Character recognition using a dynamic optoelectronic neural network with unipolar weights. *Optics Letters*, 15(21):1227–1229, 1990.

[82] C. O'Malley. The new microprocessors - power chips. *Popular Science*, 243(1):58–63, 1993.

[83] *The International Conference on Optical Computing Technical Digest*, Edinburgh, UK, 22–25 August 1994.

[84] D. R. Pape and L. J. Hornbeck. Characteristics of the deformable mirror device for optical information processing. *Optical Engineering*, 22(6):675–681, 1983.

[85] T. P. Rieker, N. A. Clark, G. S. Smith, D. S. Parmer, E. B. Sirota, and C. R. Safina. "Chevron" local layer structure in surface stabilised ferroelectric liquid crystal Smectic-*C* cells. *Physical Review Letters*, 59(23):2658–2661, 1987.

[86] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. *Learning internal representations by error backpropagation*. MIT Press, Massachusetts, 1986.

[87] T. Sakuri. Approximation of wiring delay in MOSFET LSI. *IEEE Journal of Solid-State Circuits*, SC-18(4):418–426, 1983.

[88] E. Sanchez-Sinencio and C. Lau. *Artificial neural networks: paradigms, applications and hardware implementations*. IEEE press, New York, 1991.

[89] M. Schadt and W. Helfrich. Voltage-dependent optical activity of a twisted liquid crystal. *Applied Physics Letters*, 18(4):127–128, 1971.

[90] SCIOS Scottish Collaborative Initiative in Optical Science.

[91] F. W. Sears, M. W. Zemansky, and H. D. Young. *University Physics*. Addison-Wesley, seventh edition, 1987.

[92] J. Seymour. *Electronic devices and components*. Longman Scientific and Technical, second edition, 1988.

[93] T. M. Slagle and K. Wagner. Winner-take-all spatial light modulator. *Optics Letters*, 17:1164–1166, 1992.

[94] M. W. G. Snook, 1994. A private communication.

[95] M. W. G. Snook, P. W. McOwan, W. J. Hossack, and D. G. Vass. Smart liquid crystal spatial light modulator for use in real-time image-intensity contouring. In *OSA Annual Meeting*. Optical Society of America, 3–8 October 1993.

[96] R. G. Stearns. Optically programmed neural network capable of stand-alone operation. *Applied Optics*, 32(26):5141–5152, 1993.

[97] Y. Takafuji, T. Yamashita, Y. Akebi, T. Toichi, T. Shimada, and K. Awane. A poly-Si TFT monolithic LC data driver with redundancy. In *IEEE International Solid State Circuits Conference Technical Digest*, pages 118–119, 1992.

[98] CRL Smectic Technology. 320 × 320 ferroelectric liquid crystal spatial light modulator, 1994. Hayes, Middlesex, UK.

[99] R. R. Troutmann. Epitaxial layer enhancement of n-well guard rings for CMOS circuits. *IEEE Electron Device Letters*, EDL-4(12):438–440, 1983.

[100] Y. P. Tsividis. *Operation and modelling of the MOS transistor*. McGraw-Hill, 1988.

[101] R. M. Turner, D. A. Jared, G. D. Sharp, and K. M. Johnson. Optical correlator using very-large-scale integrated circuit/ferroelectric-liquid-crystal electrically addressed spatial light modulators. *Applied Optics*, 32(17):3094–3101, 1993.

[102] R. M. Turner and K. M. Johnson. CMOS photodetector array for sequential peak detection. In *IEEE/LEOS Summer Topical Meeting Digest on Smart Pixels*, pages 54–55, Lake Tahoe, USA, July 1994.

[103] I. Underwood. *An nMOS addressed liquid crystal spatial light modulator*. PhD thesis, University of Edinburgh, 1987.

[104] I. Underwood, 1993. A private communication.

[105] I. Underwood, D. G. Vass, A. O'Hara, D. C. Burns, P. W. McOwan, and J. Gourlay. Improving the performance of liquid crystal over silicon spatial light modulators: issues and achievements. *Applied Optics*, 33(14):2768–2774, 1994.

[106] I Underwood, D. G. Vass, R. M. Sillitto, G. Bradford, N. E. Fancey, A. O. Al-Chalabi, M. J. Birch, W. A. Crossland, A. P. Sparks, and S. G. Lantham. A high performance spatial light modulator. In D. M. Gookin, editor, *Proceedings of the SPIE*, volume 1562, pages 107–115, 1991.

[107] H. J. M. Veendrick. Short-circuit dissipation of static CMOS circuitry and its impact on the design of buffer circuits. *IEEE Journal of Solid-State Circuits*, SC-19(4):468–473, 1984.

[108] Meta-Software's version of the Simulation Program for Integrated Circuit Evaluation.

[109] T. Wada and M. Koden. Liquid crystal flat panel displays. *Optoelectronics – Devices and Technologies*, 7(2):211–219, 1992.

[110] K. Wagner and T. M. Slagle. Optical competitive learning with VLSI/liquid-crystal winner-take-all modulators. *Applied Optics*, 32(8):1408–1433, 1993.

[111] M. S. Welkowsky, U. Efron, W. Byles, and N. W. Goodwin. Status of the Hughes charge coupled device addressed liquid crystal light valve. *Optical Engineering*, 26(5):414–417, 1987.

[112] R. Wilson. Mirror, mirror on the wall. *Electronics Weekly*, page 16, November 1993.

[113] M. R. Worboys, G. White, K. Mitchell, and A. Mosley. Miniature display technologies for integrated helmet systems. *GEC Journal of Research*, 10(2):111–118, 1993.

[114] T. Yamazaki, Y. Kawahara, S. Motte, H. Kamamori, and J. Nakamura. A liquid-crystal TV display panel incorporating drivers. In *Proceedings of the Society for Information Display*, volume 23, pages 223–226, 1982.

[115] Y.Hirai. Hardware implementations of neural networks in Japan. In *Proceedings of the International Conference on Neural Networks*, pages 435–446. Munich, Germany, 1991.

[116] J. M. Younse. Mirrors on a chip. *IEEE Spectrum*, 30(11):27–31, 1993.

[117] L. Zhang, M. G. Robinson, and K. M. Johnson. A fast, optoelectronic neurocomputer for character recognition. In *Proceedings of the SPIE*, volume 1469, pages 225–229, 1991.

[118] L. Zhang, M. G. Robinson, and K. M. Johnson. Optical implementation of a second-order neural network. *Optics Letters*, 16(1):45–47, 1991.

Appendix A

Pulse width modulation test pixel

The simple temporal multiplexing scheme described in Section 3.3.2 can be thought of as a discrete pulse width modulation scheme, where the pulse width is built up over a number of subframes. Here I describe an analogue pulse width modulation circuit that I included as a test structure on the SBS256 FLCOS SLM. I will give a brief description of how the circuit operates and show some preliminary results from the test structure.

A.1 Circuit principles and operation

As with the SBS256, the FLC is operated in a binary mode, but the width of the voltage pulse applied across it is controlled by an analogue value stored at the pixel. The circuit is based on the neuron pulse width output circuit described in Section 6.4.3. It consists of a 2-stage comparator, an XOR gate, and an electrode mirror (see Fig. A-1(a)). The comparator provides a pulse width output dependent on its input signals, V_P and V_N. V_P is the analogue voltage representing the desired activation of the pixel, and V_N is a triangular ramp voltage. By ramping V_N from 5 to 1V, pulse widths ranging from 0 to 100% duty cycle can be achieved. The XOR gate is included to ensure the FLC can be charge balanced, just as in the SBS256 pixel. Because of the charge balancing requirement, the pixel can be on for between 0 to 50% of the frame time. Fig. A-1(b) illustrates

how increasing V_P, increases the width of the pulse applied across the FLC, which increases the time the FLC is switched to its ON state.

The comparator does not need to be high performance and so was designed to be quite small. The only design constraints used were a common mode input range of 1–5 V and a 2 μA tail current. As this is a proof of principle circuit, both inputs of the two-stage comparator are connected to external pads. If implemented as a pixel array, the input voltage V_P would be held on a storage capacitor, and V_N would be a global ramp signal. As V_P would be stored dynamically, photo-induced charge leakage would occur. However, this could be significantly reduced by incorporating the n-well protection ring presented in Section 6.3.

A.2   Results

The electrode mirror signal is routed to a test pad to facilitate electrical testing. By keeping CK low, the 2-stage comparator exhibits a similar linear pulse width output response to that of the OANN neuron output circuit.

The optical response is shown in Fig. A–2. Here the CK signal and V_N are operated synchronously at 1 kHz and V_P is decreased from 5 to 0.5 V, in 0.5 V steps. A gradation is clearly visible; however the FLC contrast is not very high so there is not an obvious difference between adjacent grey levels. The response is fairly linear but defects in the FLC layer caused by hillocks on the electrode mirror make quantitative measurements very difficult.
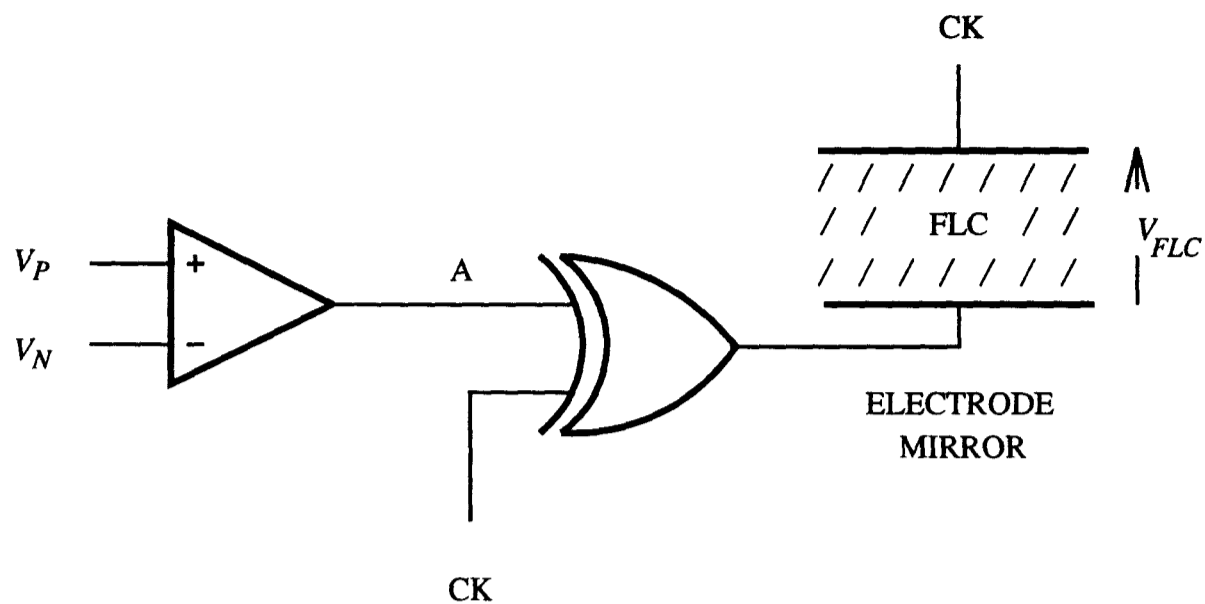
(a)
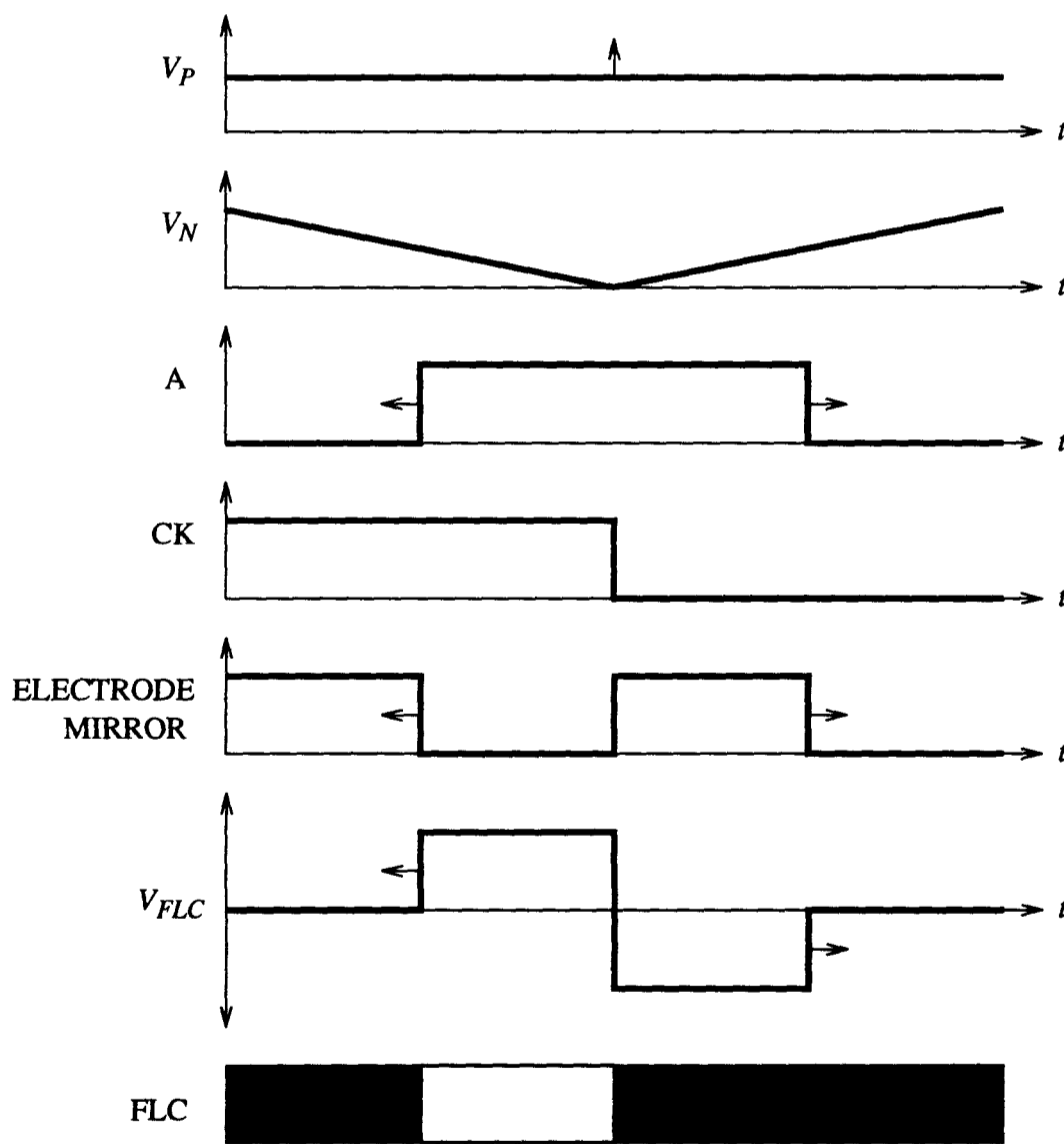
(b)

Figure A–1: (a) Schematic and (b) timing diagram for the pulse width modulation test pixel.

Figure A–2: Optical response of the pixel. Here V_P is decreased from 5 to 0.5 V, in 0.5 V steps.

A.3 Future work

The test structure shows that the pulse width modulation scheme works in principle. With fast switching ($< 50 \ \mu s$), high contrast ($\sim 100{:}1$) FLCs, the technique could be used to generate essentially analogue grey-scale (in the time-domain) at kilohertz frame rates. However, it should be noted that the grey-scale exists in the time-domain, so detector systems must integrate over the whole frame time to ensure the grey-scale is sampled correctly.

The pixel circuit could probably be implemented in a $80 \times 80 \ \mu m^2$ area using the 1.2 μm process, so a 128×128 array would be feasible. The most obvious application would be as a medium resolution raster-scanned grey-scale display device. The device could also perform thresholding and contrast enhancement functions on an input image, simply by altering the externally generated ramp function V_N. This type of pixel would also be suitable for incorporating into an optical-output version of the neuron used on the OANN backplane, thus permitting further layers of neurons to be implemented optically.

Appendix B

SBS256 press cuttings

Below are some of the press cuttings and advertisements that have appeared in various specialist electronics and optelectronic magazines, regarding CRL Smectic Technology's commercialisation of the SBS256 SLM presented in this thesis. A copy of the CRL's data sheet is also included.

Applications abound for CRL's spatial light modulator

By Peter Fletcher
Middlesex, England

Applications for spatial light modulator (SLM) devices, developed by Central Research Laboratories (CRL) of Thorn EMI plc in Middlesex, range from retail to astronomy and all markets between, said research worker John Brocklehurst.

The devices will find use in image correlation, where a stored image is compared optically in real-time with a real scene. Such a system would determine whether an object represented by a stored image is present in that scene, then trigger appropriate actions.

In optical telecommunications systems, the device will find use as a switch, directing a signal from one filament in a bundle of optical fibers to another. It also could be used in backplane wiring for high-speed electronic systems.

At a more esoteric level, Brocklehurst said, CRL has demonstrated the ability to generate and control dynamic real-time holographic displays. The SLM will modulate the phase of coherent laser lightwaves. The resultant holographic images could then be used to steer lightwaves, correct telescopic images or replicate three-dimensional images.

In optical computing, the spatial modulators will find a place as key elements in neural networks, Brocklehurst predicted. ∎

CRL'S SLM ADVANTAGES

John Brocklehurst of CRL said the lab's spatial light modulator has two distinct advantages because it operates at ambient illumination levels with ordinary white light:

First, it allows the modulator to run at very high speeds. He claims the 256-by-256 array of pixels can be changed at a rate of up to 3,000 frames per second.

Second, and important for system developers, the ability to use ambient white light helps reduce the cost of optical systems.

Early samples of the SLM 256, though, will be relatively expensive. CRL plans to make it commercially available to other researchers almost immediately for a price of US$13,500 each. "That price will fall rapidly when production volumes ramp up," Brocklehurst said. —PF

Figure B–1: Scanned from *Electronics*, 23 May 1994.

SRAM-based SLM has a faster switching speed of 3kHz, useful for colour projection screens.

Breakthrough leads to fast image recognition

by Nick Flaherty

Faster, more compact, image recognition, optical switching and display systems will be possible with the first spatial light modulator (SLM) based on SRAM technology.

It was developed by the Central Research Labs of Thorn-EMI and Edinburgh University, and is used as a reflective optical processing element in image recognition, miniature displays and optical switching.

Using an SRAM 'backplane' as the controlling element for a layer of ferroelectric liquid crystals means that an optical system half the size of previous systems can be made, with switching speed increased to 3kHz. This increase in speed

means the SLM could be used in a colour projection screen.

Some of the control electronics is integrated into the backplane with an exclusive OR cell and a latch. More control circuitry could be integrated on to the periphery of the backplane to further reduce the size of the system.

The pixels are 19 × 19µm on a 40µm pitch. CRL plans to increase the pixel size to 37 × 37µm on the same pitch to make the image appear smoother.

Other SLMs, such as the DX230 launched by CRL in February, are controlled by rows of electrodes for transmissive sys-

tems. Some SLM systems use DRAM cells, which are easier to make but need to be constantly refreshed. The DRAM SLM also needs strobe light illumination.

The SRAM base controlling the switching of the crystal was designed by Edinburgh University as part of a Scottish cross university project, and made by asic house Austria Mikro Systeme on a 1.2µm process.

CRL says the system will be able to recognise road signs in realtime as part of the European Prometheus programme. It claims the SRAM-based SLM could cut the system down from something filling a car boot to the size of a shoebox.

Figure B–2: Scanned from *Electronic Times*, 26 May 1994.

Static Memory
May Shrink
Imaging Systems

HAYES, Middlesex, England — Faster, more compact image-recognition, optical-switching and display systems will be possible with the first spatial light modulator based on static random-access memory (SRAM) technology.

The device was developed by the Central Research Labs (CRL) of Thorn EMI and Edinburgh University and is used as a reflective optical processing element.

Using an SRAM backplane as the controlling element for a layer of ferroelectric liquid crystals can produce an optical system half the size of previous systems, with switching speeds increased to 3 kHz, according to CRL. This increase in speed means the modulator could be used in a color projection screen.

Some of the control electronics are integrated into the backplane with an exclusive XOR logic cell and a latch. More control circuitry can be integrated onto the periphery of the backplane to further shrink the system.

The pixels are 19 × 19 µm on a 40-µm pitch. CRL plans to increase the pixel size to 37 × 37 µm on the same pitch to make the image appear smoother.

The SRAM base was designed by Edinburgh University as part of a Scottish cross-university project and manufactured by ASIC house Austria MikroSysteme on a 1.2-µm process.

Other spatial light modulators, such as the DX230 launched by CRL in February, are controlled by rows of electrodes for transmissive systems.

Some use dynamic random-access memory (DRAM) cells, which are easier to make but which must be constantly refreshed. Modulators based on DRAMs also need strobe-light illumination.

CRL gives as an example of the advantages of the new spatial light modulator a system to recognize road signs in real-time as part of the European Prometheus program. It claims that the new modulator could shrink the system from something filling a car trunk with the DX230 to one the size of a shoe box. □

John Walko

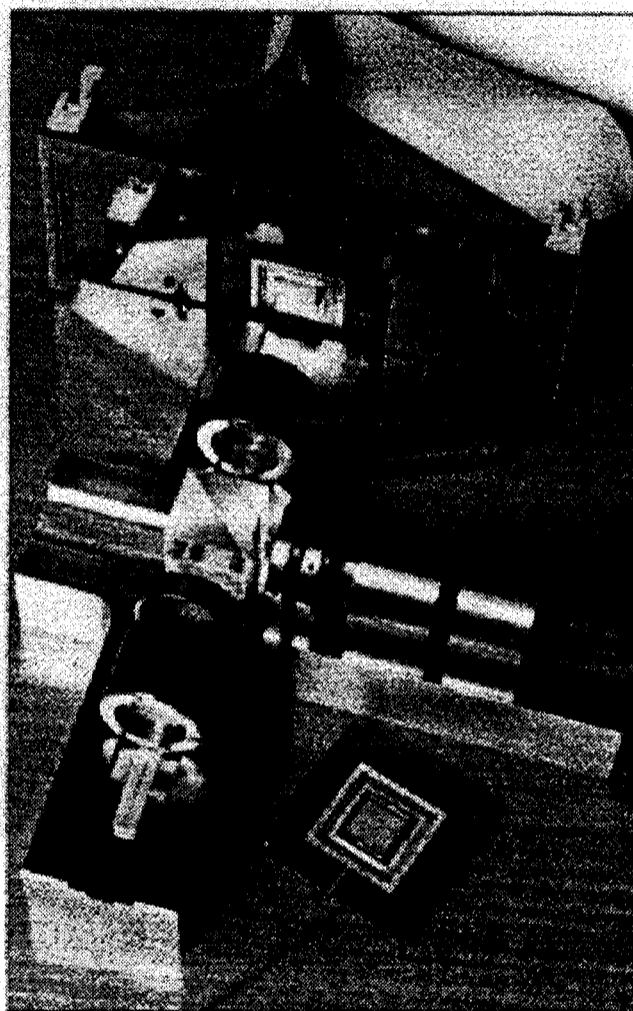Figure B-3: Scanned from *Photonics Spectra*, July 1994.

Compact SLM suits continuous illumination

The latest in a series of silicon-backplane ferroelectric liquid crystal spatial light modulators (SLMs), the SBS256, has been launched by CRL, a subsidiary of Thorn EMI. The compact PC-addressable device offers 256 × 256 resolution and high-speed operation.

Specification Standard CMOS VLSI technology to form circuitry at each pixel allows the fast switching times of the ferroelectric liquid crystal to be exploited. IC structure is insensitive to incident light so it is suited to use with high levels of continuous illumination.

The device can be operated at a frame rate of up to 3 kHz – at this speed the binary nature of the liquid crystal effect can be used to generate significant pseudo grey levels for visible observation.

Comment Applications include optical correlators for industrial inspection, compact imaging systems (helmet-mounted displays and virtual reality), and optical implementations of neural networks.

Availability Single SB256 SLM has a list price of £9000. Delivery time six to eight weeks.

Insert 334 on OLE reply card

Figure B–4: Scanned from *Opto & Laser Europe*, September 1994.

Silicon Backplane (SRAM) Spatial Light Modulator (SBS256)

Device Features

- **SLM uses single crystal silicon integrated circuit and ferroelectric smectic liquid crystal technology**

    The use of standard CMOS VLSI technology to form circuitry at each pixel allows the fast switching times of the ferroelectric liquid crystal to be utilised, and offers the potential of low prices for quantity supplies.

- **SRAM latch technology allows high incident light levels**

    This IC structure is not sensitive to incident light so it is suited to use with high levels of continuous illumination.

- **Very fast frame scan times**

    The SLM can be operated at a frame rate of up to 3kHz. As an example, at this operating speed the binary nature of the liquid crystal effect can be used to generate significant pseudo grey levels for observation by eye.

- **Compact device**

    The SLM device is physically compact, so the associated optics which are required for systems can also be small.

- **PC addressable**

    The standard device is offered with a simple PC interface. This allows straight forward operation of the SLM in a laboratory, and application specific interfaces can be supplied.

- **Small pixel pitch and phase or amplitude switching**

    The 40μm pixel pitch and the binary phase or amplitude switching of the SLM make it suitable for optical processing applications such as compact correlators.

- **Faster interface, optics and mounting hardware**

    We are able to offer standard or custom optical base plates and suitable optical components for use with the SLM. A faster interface using memory mounted close to the SLM and a video interface will be available. (These will allow more use of the fast frame scan rates available with this SLM)

SBS256 with Simple Interface System

Number of pixels	256 x 256	(65,536 total)
Active array size	10.24 x 10.24mm	
Pixel pitch	40μm	
Pixel mirror size	19 x 19μm	
Effective fill factor	22%	
Modulation	Binary	Phase or amplitude
Flatness	Better than λ/2	Target
Operating wavelength	633nm standard	Can be optimised for operation at other customer specified wavelengths
Cell thickness for maximum contrast	$d = \dfrac{(2m+1)\lambda}{4\,\Delta n}$	where: m is an integer λ is the operating wavelength Δn is the liquid crystal birefringence
Interface	PC Digital Input Output (DIO) card	Simple interface is upgradeable for specific customer applications
Software	PC compatible	
Refresh rate	~20Hz	Rate limited by the simple interface for a single frame and PC dependant.
Frame scan time	~300μs @ 25°C	~3kHz SLM device update rate. (Charge balanced operation)
Duty cycle	50%	Unpulsed light source
Pixel contrast ratio	150 : 1	* Target
Area contrast ratio	150 : 1	* Target
Optical power density	To be determined	High levels expected
SLM drive voltage	6V	
SLM power dissipation	~300mW	Estimated
Mounting of SLM	PGA socket onto custom PCB 138 x 63 x 35mm	Un-packaged PCB is designed of ease of mounting to accessory optical base plates. Supplied with removable transit covers.
Custom interface size	190 x 100 x 40mm	Plastic cased PCB
Electrical supply	None required	Drive voltages derived from host PC
Anti-reflection coatings optimised for particular customer specified wavelengths are available		

The Silicon Backplane incorporated in this device was designed at The University of Edinburgh.

For further information contact:	**C R L** Smectic Technology
	CRL Smectic Technology, Dawley Road, Hayes, Middlesex, UB3 1HH, United Kingdom Telephone +44 (0)81 848 6428 Facsimile +44 (0)81 848 6565 E-mail smectic@thorn-emi-crl.co.uk

THORN EMI

Appendix C

Author's publications

[1] D. C. Burns and I. Underwood. Real time grey level generation using a binary mode spatial light modulators. In *Proceedings of the Applied Optics and Opto-Electronics Conference*, pages 214–215, Leeds, UK, 14–17 September 1992.

[2] D. C. Burns, I. Underwood, A. F. Murray, and D. G. Vass. An optoelectronic neural network with temporally multiplexed grey-scale weights. In *Proceedings of the Fourth International Conference on Microelectronics for Neural Networks and Fuzzy Systems*, pages 3–7, Torino, Italy, 26–28 September 1994.

[3] D. C. Burns, I. Underwood, D. G. Vass, and A. O'Hara. Electronically addressed ferroelectric liquid crystal over silicon spatial light modulators for optical computing. In *The International Conference on Optical Computing Technical Digest*, pages 27–28, Edinburgh, UK, 22–25 August 1994.

[4] D. C. Burns, I. Underwood, D. G. Vass, A. O'Hara, and I. Rankin. Electronically addressed ferroelectric liquid crystal over silicon spatial light modulators. In *IEEE/LEOS Summer Topical Meeting Digest on Smart Pixels*, Lake Tahoe, USA, July 1994. Post-deadline addendum.

[5] A. O'Hara, I. Underwood, D. G. Vass, and D. C. Burns and J. T. M. Stevenson. Interlevel dielectric planarisation to optical tolerances using chemical-mechanical polishing in the fabrication of spatial light modulator backplanes. *To be published by the Materials Research Society*, Spring 1995.

[6] A. J. Stevens, J. Gourlay, S. Samus, W. J. Hossack, D. G. Vass, and D. C. Burns. Experimental investigation of free space optical interconnects. In *Submitted to OSA Spatial Light Modulators Spring Topical Meeting*, Salt Lake City, Utah, USA, 12–17 March 1995.

[7] I. Underwood, D. G. Vass, A. O'Hara, and D. C. Burns. Issues in improving the performance of liquid crystal over silicon spatial light modulators. In *Spatial Light Modulators and Applications Technical Digest*, volume 6, pages 110–113. Optical Society of America, Washington, DC, 1993.

[8] I. Underwood, D. G. Vass, A. O'Hara, D. C. Burns, P. W. McOwan, and J. Gourlay. Improving the performance of liquid crystal over silicon spatial light modulators: issues and achievements. *Applied Optics*, 33(14):2768–2774, 1994.

An Optoelectronic Neural Network with Temporally Multiplexed Grey-Scale Weights

D. C. Burns, I. Underwood,
and A. F. Murray
Department of Electrical Engineering
University of Edinburgh
Mayfield Road
Edinburgh, UK, EH9 3JL

D. G. Vass
Department of Physics & Astronomy
James Clerk Maxwell Building
University of Edinburgh
Mayfield Road
Edinburgh, UK, EH9 3JZ

Abstract

We present a high fan-in optoelectronic artifical neural network with temporally multiplexed grey scale weights and a prototype neural-detector backplane. The weights are generated using a high performance liquid crystal over silicon spatial light modulator. The neural-detector backplane is capable of recognising both excitatory and inhibitory activations.

1 Introduction

Optics offers much for neural integration. Optical systems can have massive fan-in, and optical processing has potential speed advantages over a purely electronic system. More pragmatically, an optical system can interface directly to an optically-presented task. In most neural systems, the input layer is the location of the massive analogue parallelism. Subsequent layers are narrower. We describe below an experimental implementation of a modest layered network where the input layer is optical, and the output layer electronic. We believe that this offers an optimal combination of the strengths of both techniques in many recognition applications. Most examples of direct optical-input systems use a (possibly integrated) sensor array for image capture, with subsequent electronic processing. With spatial light modulator (SLM) technology, the first layer of processing is performed optically, under electronic control. We believe this to be a novel and powerful combination. We also present a simple technique to significantly reduce photo-induced charge leakage of the neuron activations which are stored dynamically on capacitors. The SLM technology of ferroelectric liquid crystal over very large scale integrated (FLC/VLSI) sil-icon has matured considerably over the past few years. These high speed, low power SLMs function as electronically-written, optically-read memory devices and are used as reflective optical processing elements in image recognition, miniature displays and optical switching. At present, most SLM silicon backplanes work in a binary mode, as they are based on standard digital circuits. Similarly, FLCs are inherently binary in nature. However, their fast switching times and the high addressing speeds of silicon backplanes permit temporal multiplexing to be used as a means of real-time grey-scale emulation. This renders them suitable as adaptive weight planes in optoelectronic neural networks.

2 256 × 256 pixel FLC/VLSI SLM

The silicon backplane of a FLC/VLSI SLM consists of an 2-dimensional array of pixels, each with a memory circuit and reflective electrode-mirror, and peripheral circuitry such as shift registers and/or decoders to address the array. An electric field can be generated across a thin overlying layer of FLC by applying the appropriate voltage to the electrode-mirror. This can alter the state of the FLC to produce binary phase or amplitude modulation in an incident read-beam.

Many of the medium-resolution FLC/VLSI SLMs designed so far have been based on the single transistor dynamic random access memory (DRAM) type pixel [6, 3, 2]. As a development of previous prototype devices [7, 4, 5], we have designed, fabricated, and tested a 256 × 256 pixel FLC/VLSI SLM that incorporates a static random access memory (SRAM) latch and an exclusive-OR (XOR) gate at each pixel (see Fig.1). Although requiring more area, the SRAM-

XOR pixel has significant performance advantages over the DRAM-type. These include robust data storage under high read-beam intensities, easy FLC charge balancing to prevent chemical degradation, synchronous FLC switching across the whole array, and the ability to switch high spontaneous polarisation FLC mixtures. The SLM was constructed by sandwich-

Figure 2: Cross-section of a FLC/VLSI SLM.
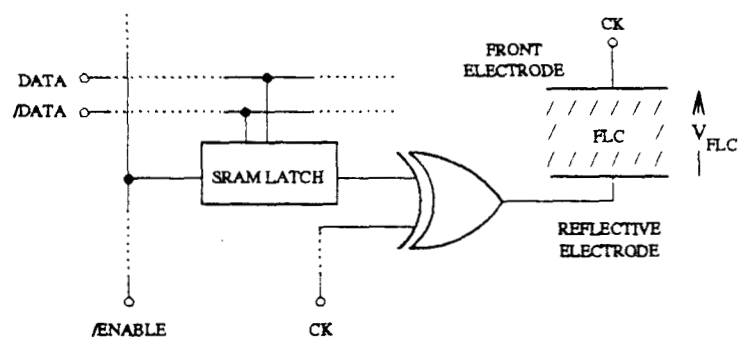
(a)
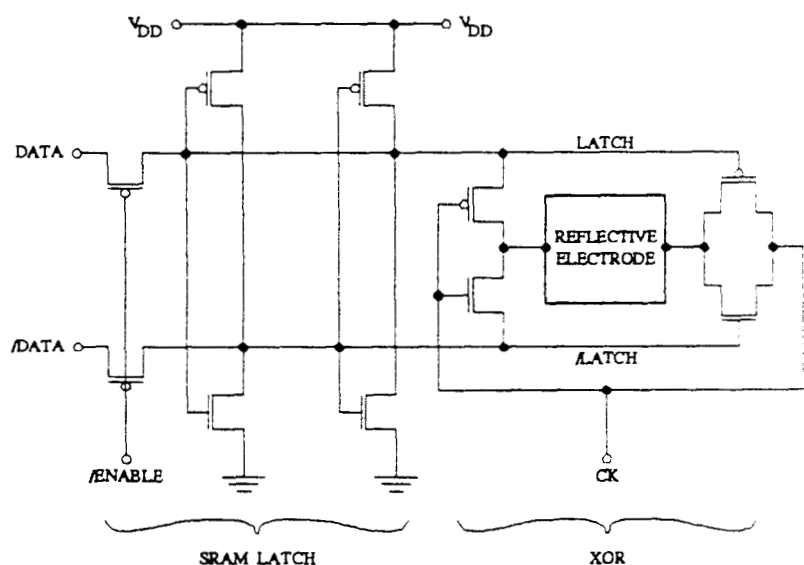
(b)

Figure 1: (a) Schematic and (b) transistor-level diagram of the SRAM-XOR pixel.
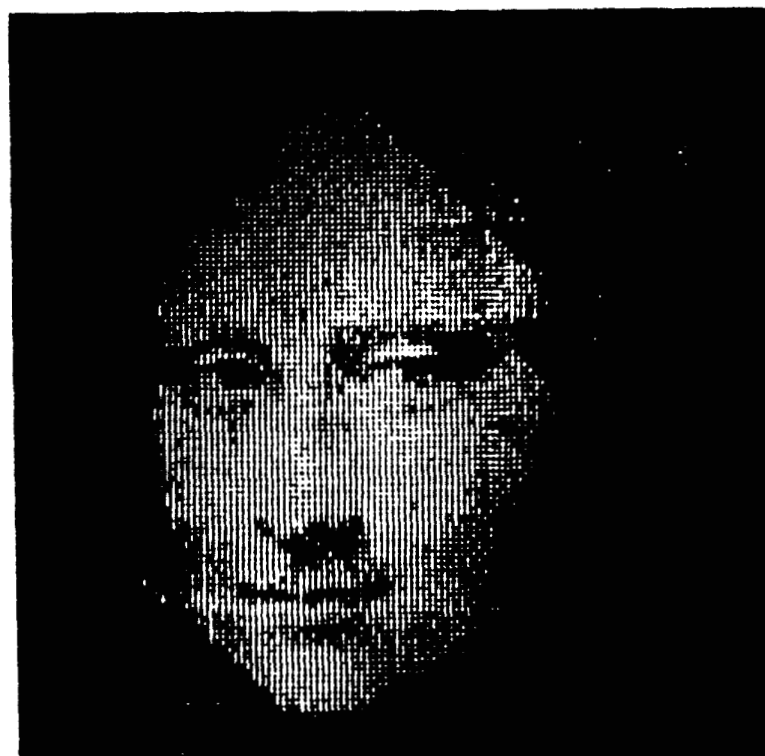
Figure 3: Test pattern on a section of the 256 × 256 pixel FLC/VLSI SLM.

ing a 2-3μm film of FLC between the VLSI silicon backplane and a cover glass coated with a transparent front electrode (see Fig. 2). The backplane was fabricated by Austria Mikro Systeme on a 1.2μm n-well double-metal, 5.5V, complementary metal oxide semiconductor epitaxial process. The 256 × 256 pixels are on a 40μm pitch, resulting in a 10.24mm × 10.24mm active array. A bit-plane is loaded column-sequentially onto the array at up to 24MHz via a 32-bit data bus (see Fig. 3).

The SLM can display up to 3000 bit-planes per second. This high frame rate permits temporal multiplexing to be used to generate real-time grey-scale. Temporal multiplexing involves splitting a grey-scale image into a number of time sequential subframes. Each subframe has an associated bit-plane that is generated by an appropriate encoding algorithm. The bit-planes are sequentially scanned into the SLM so that 'integrating' over a number of subframes gives the desired grey-scale image. The linear encoding algorithm

shown in Fig. 4 is satisfactory for generating a small number of grey levels; however, each extra grey level requires an extra subframe. The frame rate is inversely proportional to the number of grey levels.
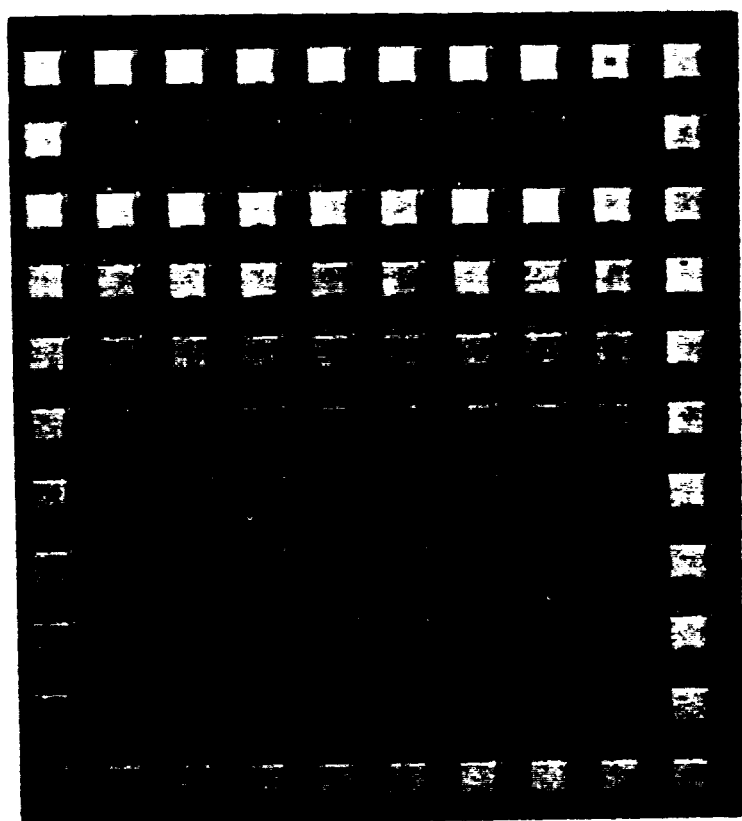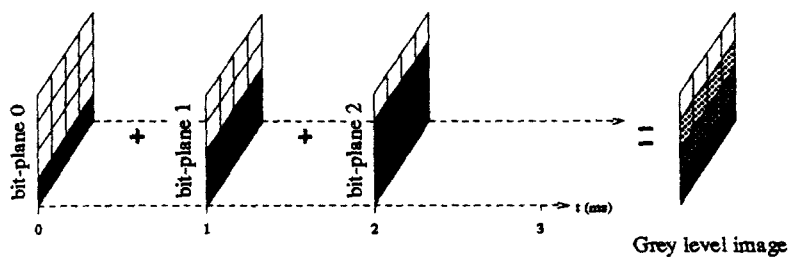

Figure 4: Temporal multiplexing using a linear encoding algorithm.

A more efficient binary weighted encoding algorithm pulsed sampling in the interrogation system can result in much higher frame rates. In this case, n subframes can generate 2^n grey levels. So, for example, 256 grey levels can be generated at 125Hz with eight subframes each 1ms in duration.

3 System Architecture

The optical system (see Fig.5) has been made as simple as possible so that the neural backplane can be

fully tested and characterised. The multiplication of the input pattern by the interconnection weights matrix (IWM) is calculated in software. A more advanced system is planned, where the multiplication will be performed optically by cascading two SLMs. A 4 × 4 array of 16 × 16 input-IWM patterns are placed on the SLM with a mark-space ratio of 1:1, to match the spacing of the photodiodes on the neural backplane. The lens, L1, scales and focusses each of the input-IWM patterns onto the appropriate photodiode in the neural-detector backplane. If a computer generated hologram was used rather than a simple lens, then more of the SLM could be utilised as the input-IWM patterns would not have to be on the same mark-space ratio as the photodiodes.
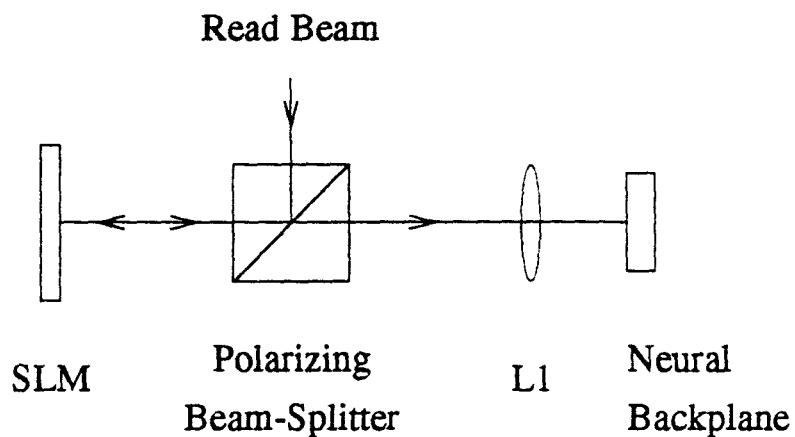

Figure 5: The optical system

3.1 Neural-Detector Backplane

The neural-detector backplane consists of a 4 × 4 array of optical-input, electronic output neurons (see Fig. 6). The four photodiodes at the corners of the array are used to help align the optical system. The outputs are pulse width modulation encoded so that they are compatible with the EPSILON chip set [1]. This will allow us to construct a multi layer perceptron network for more complex image classification.

3.2 Neuron Circuit

Each neuron consists of a $100\mu m \times 100\mu m$ photodiode, current-switching circuitry, an activity storage capacitor, and a two-stage comparator to encode the neuron activity into a pulse width signal (see Fig. 7). The photodiode is operated in the linear photoconductive mode. Its dynamic range is limited to 10^2–10^3 by shot, thermal, generation-recombination, and

flicker noise sources inherent in it and the supporting circuitry. To keep the activity capacitor, C_{ACT}, to a reasonable value (5pF), a sampling scheme was used. A binary weighted pulse train, synchronised with the bit-planes that appear on the SLM, is applied to the sample and hold signal, S/H. The signal O/O defines whether charge is dumped onto or taken off C_{ACT}, whenever S/H is pulsed. Most of the time however, S/H is low so current is sourced or sunk from the reference signal, VSPONGE, depending on O/O. After both positive and negative activation components have been presented to the activity capacitor, the 2-stage comparator encodes V_{ACT} as a pulse width signal by applying the desired transfer function signal, TF. This signal is generated by counting through a block of off-chip RAM that feeds an 8-bit digital-to-analogue converter. Simple triangular ramp or variable-gain sigmoidal functions can be produced simply by loading the RAM with the appropriate bit-pattern.
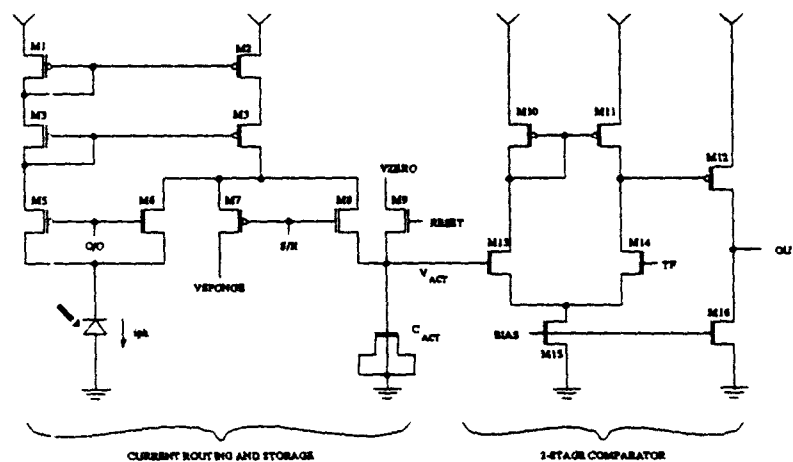
Figure 7: Transistor level diagram of optical input electronic output neuron.

3.3 Activity Storage

Figure 6: Photomicrograph of the 4 × 4 neural-detector backplane. Each neuron occupies 200μm × 200μm.

The neuron activity is stored dynamically on C_{ACT}. The drains of access transistors M8 and M9 are reversed biased n^+p diodes so any photo-induced minority-carriers (electrons in the p-substrate) generated near them could be swept across their depletion regions so discharging C_{ACT}. From a previous test-chip, we found that just covering a storage capacitor and associated access circuitry with second-level metal to protect against direct incident light did not improve the charge storage characteristics by an appreciable amount. Photo-induced minority electrons in p-type Si can travel for up to 300μm before they become localised in a recombination centre and recombine with a hole, so the second-level metal covering the neuron circuitry was extended 400μm beyond the edge of the array.

As most of the photo-induced carriers are generated near the surface, the access transistors were surrounded with a ring of n-well [8] tied to the positive power rail so that minority electrons would tend to be swept into it before they reached the access transistor drains (see Fig. 8). Fig. 9 shows the results of the 1% charge decay time for a 5pF capacitor with and without n-well ring protection. This simple protection device has decreased the decay rate by a factor of over 2000 for light intensities present in the system so that a complete forward pass calculation can be performed without any charge-decay degradation in the result.
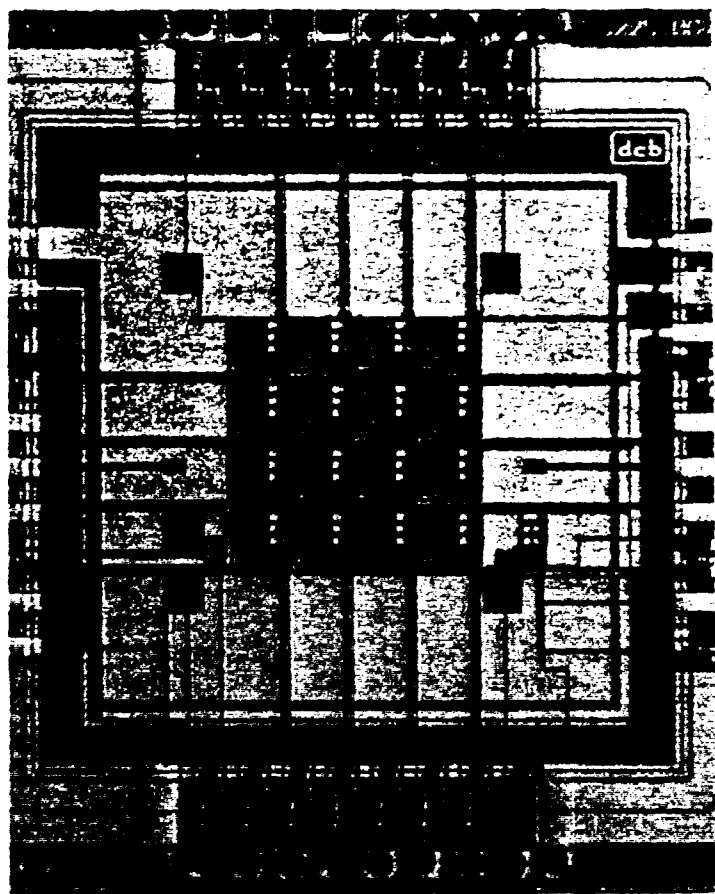
6

Figure 8: Charge storage n-well ring protection to reduce decay.

4 Conclusions

This paper has described a modest system. However, the potential is much greater. FLC/VLSI SLM technology will continue to improve both in frame rate (up to 10kHz) and resolution (up to 1024 × 1024). The neural-detector plane could be expanded to up to a 16 × 16 array before multiplexing would be required on the output pins. Such a system would be capable of performing recognition tasks in the optical domain in a compact and robust form. Also, programmability - the ubiquitous problem with optical systems - is dealt with electronically, thus allowing of optical neural systems that are reprogrammable in situ, and electronically.

Figure 9: Charge storage with and without n-well ring protection.

References

[1] A. Hamilton et al, "Pulse stream VLSI circuits and systems: The EPSILON neural network chipset," *International Journal of Neural Systems*, **4**, 395-405 (1993).

[2] M. A. Handschy et al, "One-transistor DRAM ferroelectric liquid crystal spatial light modulator," *Proceedings of the Society of Photo-optical Instrumentation Engineers*, **2237**, 432–438 (1994).

[3] D. J. McKnight, K. M. Johnson, and R. A. Serati, "Electrically addressed 256 by 256 liquid crystal on silicon spatial light modulator," *Optics Letters*, **18**, 2159–2161 (1993).

[4] D. J. McKnight, D. G. Vass, and R. M. Sillitto, "Development of a spatial light modulator: A randomly addressed liquid crystal over nMOS array," *Applied Optics*, **28**, 4757–4762 (1989).

[5] I. Underwood et al, "Improving the performance of liquid crystal over silicon spatial light modulators: issues and achievements," *Applied Optics*, **33**, 2768–2774, (1994).

[6] I. Underwood et al, "A high performance spatial light modulator," *Proceedings of the Society of Photo-optical Instrumentation Engineers*, **1562**, 107–115 (1991).

[7] I. Underwood, D. G. Vass, and R. M. Sillitto, "Evaluation of an nMOS VLSI array for an adaptive liquid-crystal spatial light modulator," *IEE Proceedings*, Part J, **133**, 77–82 (1986).

[8] R. R. Troutmann, "Epitaxial layer enhancement of n-well guard rings for CMOS circuits," *IEEE Electron Device Letters*, **EDL-4**, 438–440 (1983).
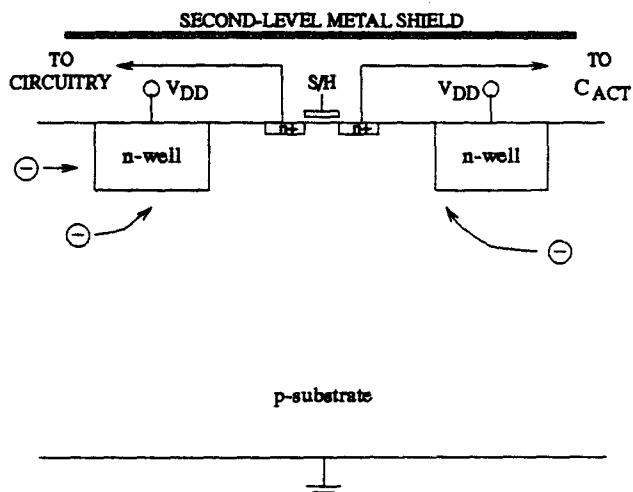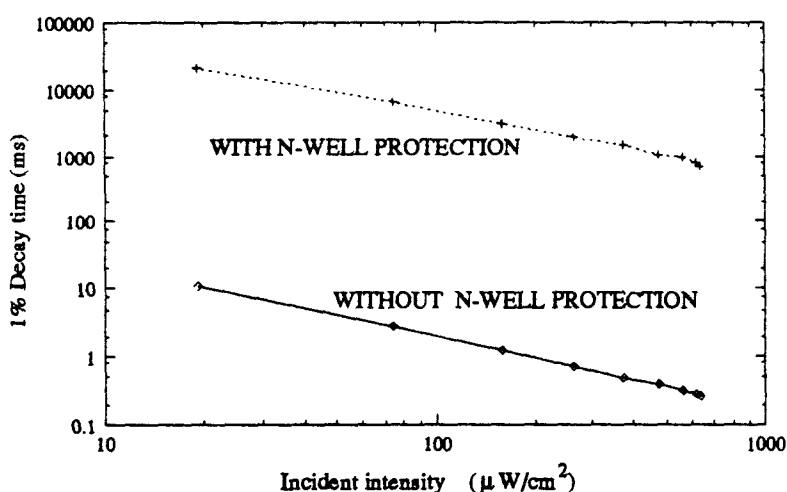
INTERLEVEL DIELECTRIC PLANARIZATION TO OPTICAL TOLERANCES USING CHEMICAL-MECHANICAL POLISHING IN THE FABRICATION OF SPATIAL LIGHT MODULATOR BACKPLANES

A. O'Hara, I. Underwood[†], D. G. Vass, D. C . Burns[†] and J. T. M. Stevenson[†],
Department of Physics, [†] Department of Electrical Engineering,
The University of Edinburgh, The King's Buildings, Edinburgh EH9 3JL, Scotland, UK

ABSTRACT

Chemical-mechanical polishing (CMP) has been used as the planarization technique in the fabrication of the silicon backplanes which form the basis of the opto-electronic devices, electrically addressed spatial light modulators (SLMs). The production of high optical quality devices requires that the interlevel dielectric be planarised to optical tolerances. The CMP process was applied to commercially fabricated wafers as part of a post-processing procedure. The local surface variation of the interlevel dielectric was consistently less than 100Å providing a suitable substrate onto which the top level metal is deposited.

INTRODUCTION

Spatial light modulators (SLM) are opto-electronic devices used in the fields of optical computing, image processing and image displays [1]. One type of SLM is the liquid crystal (LC) over silicon device which combines these two mature technologies [2]. The SLM silicon backplane consists of an array of pixels similar to DRAM or SRAM devices but each with an electrode whose voltage is determined by the associated pixel control circuitry. The voltage on the electrode dictates the state of the overlying liquid crystal which in turn is used to modulate light incident onto the device. The electrodes must also operate as mirrors and with the standard fabrication of the backplanes two problems are encountered. First, the metal surface is optically very poor due to the distortion of the film during metal sinter. Secondly, the pixel mirrors have a very low fill factor (i.e. active mirror area relative to overall pixel area). Both of these problems can be overcome by adding a further level of metal which is dedicated to the fabrication of the electrodes thus allowing optimisation of this level optically [3]. To ensure a flat surface onto which the top level metal is deposited, the interlevel dielectric must be planarized to the high optical standard required for devices illuminated by coherent light.

The level of planarization required is set not by the photolithography and step coverage considerations but by the more demanding optical requirements of the device. The surface variation must be $<\lambda/10$ (~ 500Å for visible light) to provide the necessary optical quality for a standard mirror with stricter requirements for higher quality mirrors. These specifications are more severe than the ~ 2700Å requirement [2] typically set by the fabrication of sub 0.5μm geometries. CMP is shown to be capable of realising these unique demands and has been incorporated into a post-processing procedure which is applied to commercially fabricated wafers.
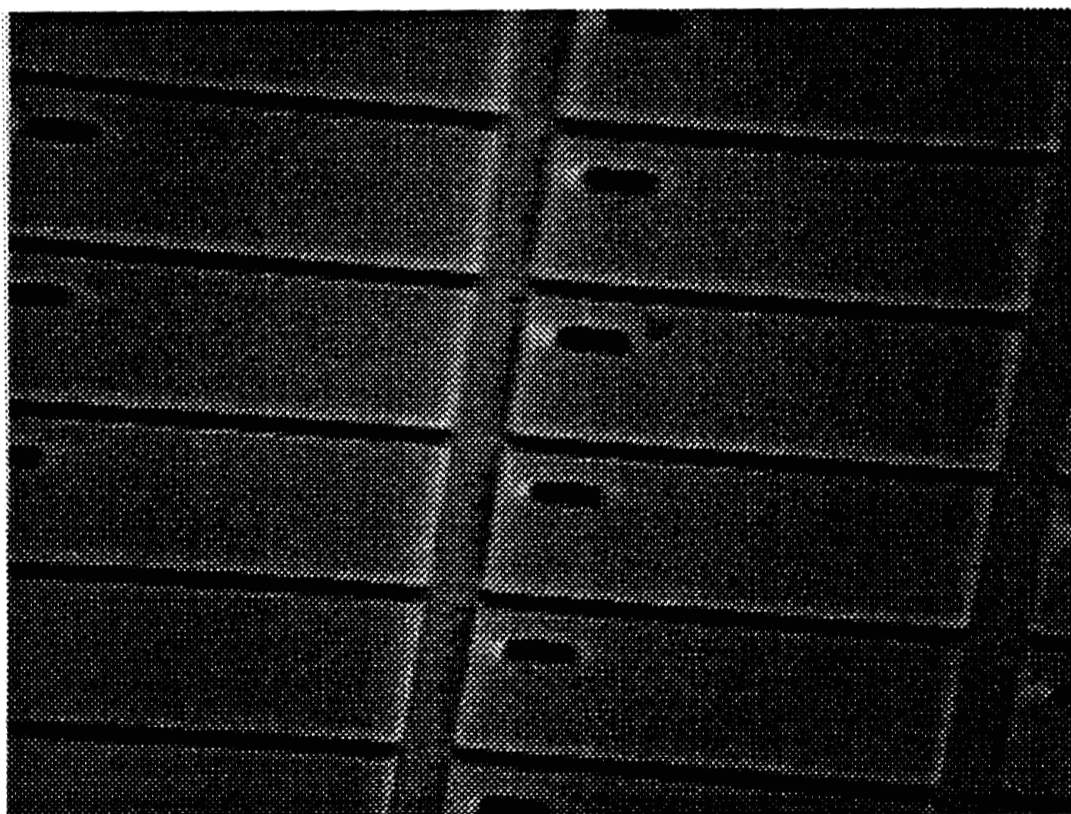
POST-PROCESSING PROCEDURE

The post-processing procedure was used to add a further level of metal to commercially fabricated wafers from AMS in Austria. This technique has been applied to two sets of wafer designs. A 176^2 DRAM SLM backplane fabricated in $3\mu m$ p-well CMOS technology and a 256^2 SRAM SLM backplane fabricated in $1.2\mu m$ n-well technology. The wafers were removed from the normal process prior to passivation to improve the normal functionality of the devices as designed, which is associated with the LC switching speed. Our custom post-processing began with the deposition of a thick ($\sim 4\mu m$) layer of SiO_2 using ECR-PECVD which conforms to the underlying circuitry. The topography of the surface at this stage was $\sim 1\mu m$ peak to peak. This partially sacrificial layer requires to be thick to compensate for the effect of erosion during CMP, thus allowing the polishing stage to achieve the required surface finish while leaving sufficient material to provide a suitable inter-level dielectric. Also, since the underlying metal layer is populated with hillocks a thicker dielectric is desirable to reduce the chance of shorts between the metal layers. The wafers were then polished using a Logitech PS2000 polishing system. The polishing pad is made from expanded polyurethane and on the Logitech system the pad curvature can be varied to fine tune the polish rate uniformity. During polishing the wafer surface is analysed periodically to estimate the polishing rate and uniformity. The process is stopped when the desired surface finish is obtained. The remaining procedures are 'standard' processing steps first to produce the vias and then to deposit and pattern the top level metal.

RESULTS

The level of flatness achieved by the CMP process was independent of the starting wafer. The resultant local surface variation was consistently less than 100Å ($250\mu m$ scans) with some measurements as low as 40Å r.m.s., figure 1. This degree of flatness ensured that the interlevel dielectric provided a suitable substrate onto which the second or third level metal was deposited for the 176 DRAM SLM or the 256 SRAM SLM respectively. From an optical point of view the mirrors created are superior to the original mirrors in both surface finish and pixel mirror fill factor. Figure 2. shows a portion of the 256 SRAM SLM array before and after planarization displaying the high quality of the mirrors obtained and also giving a visual indication of the planarity achieved.
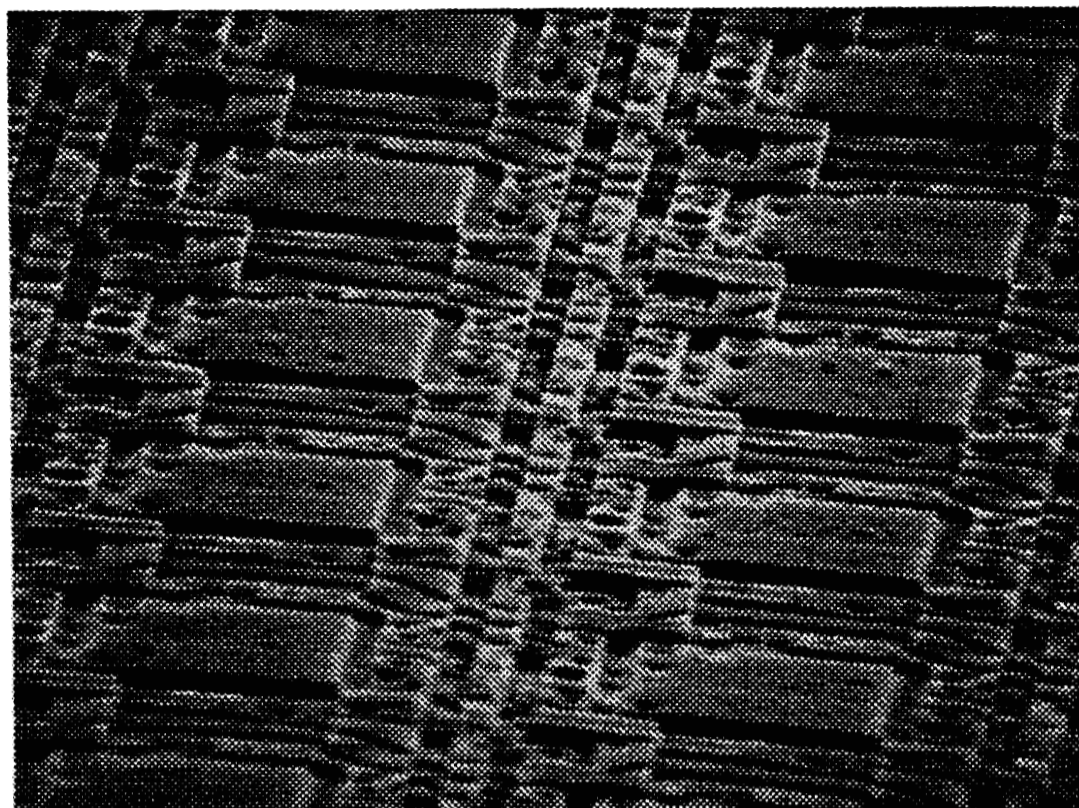
(a)

(b)

Figure 1: Figure 2. SEM micrographs of 256^2 SLM top surface (a) after standard fabrication process and (b) following additional custom post-processing procedure. The pixel pitch is 40μm. The micrographs were taken with the same magnification
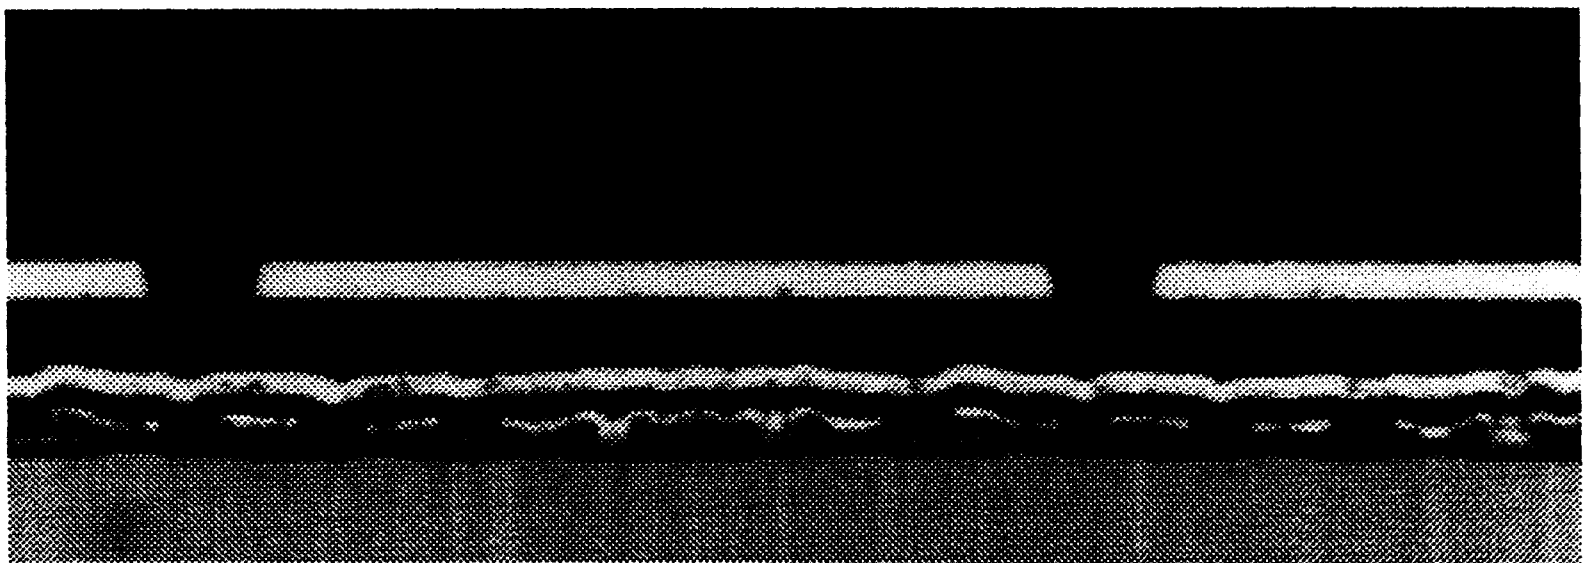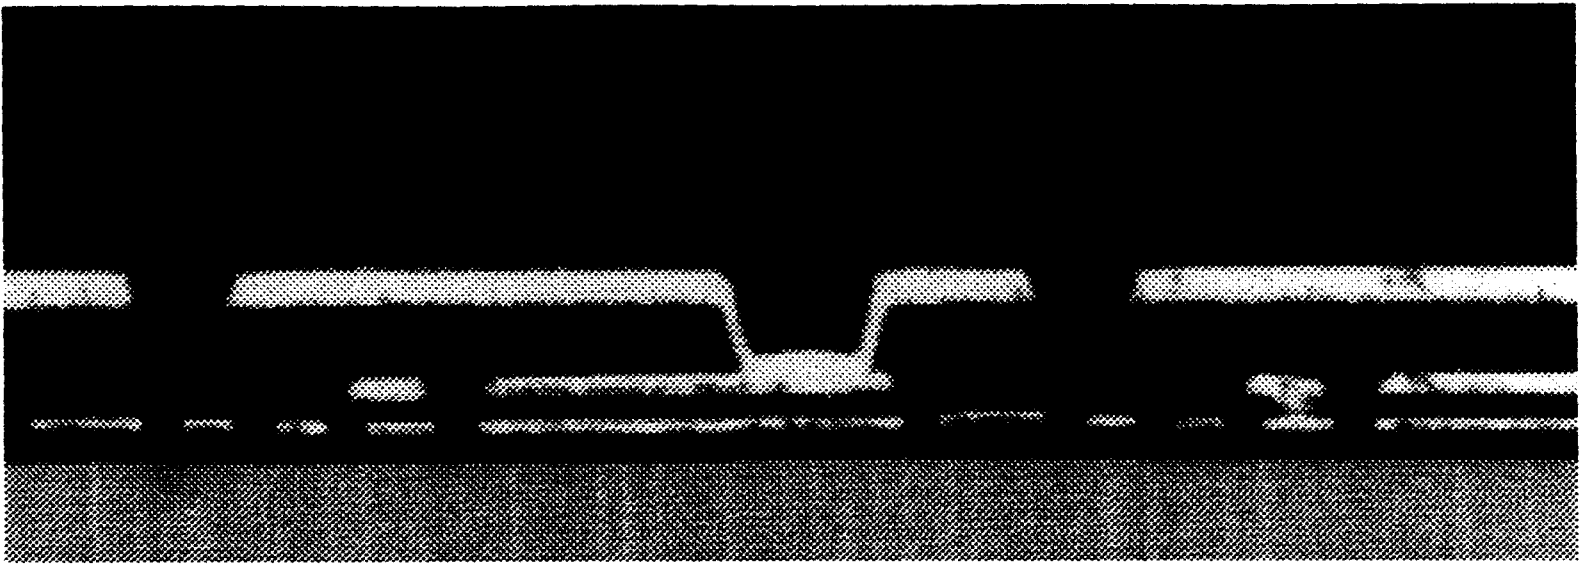
Figure 2: Figure 3. Cross-sections of 256^2 SLM through the pixel array displaying the planarity of the 3^{rd} level metal.

REFERENCES

1. See special issue on spatial light modulators, Applied Optics 31, 3857 (1992).

2. K.M. Johnson, D.J. McKnight and I. Underwood, IEEE J. Quantum Electron. 29, 699 (1993).

3. A. O'Hara, J.R. Hannah, I. Underwood, D.G. Vass and R.J. Holwill. Applied Optics, 32, 5549 (1993).

4. F. Moghadam, SRC Topical research conference on chem mechanical polishing for planarization. (1992).

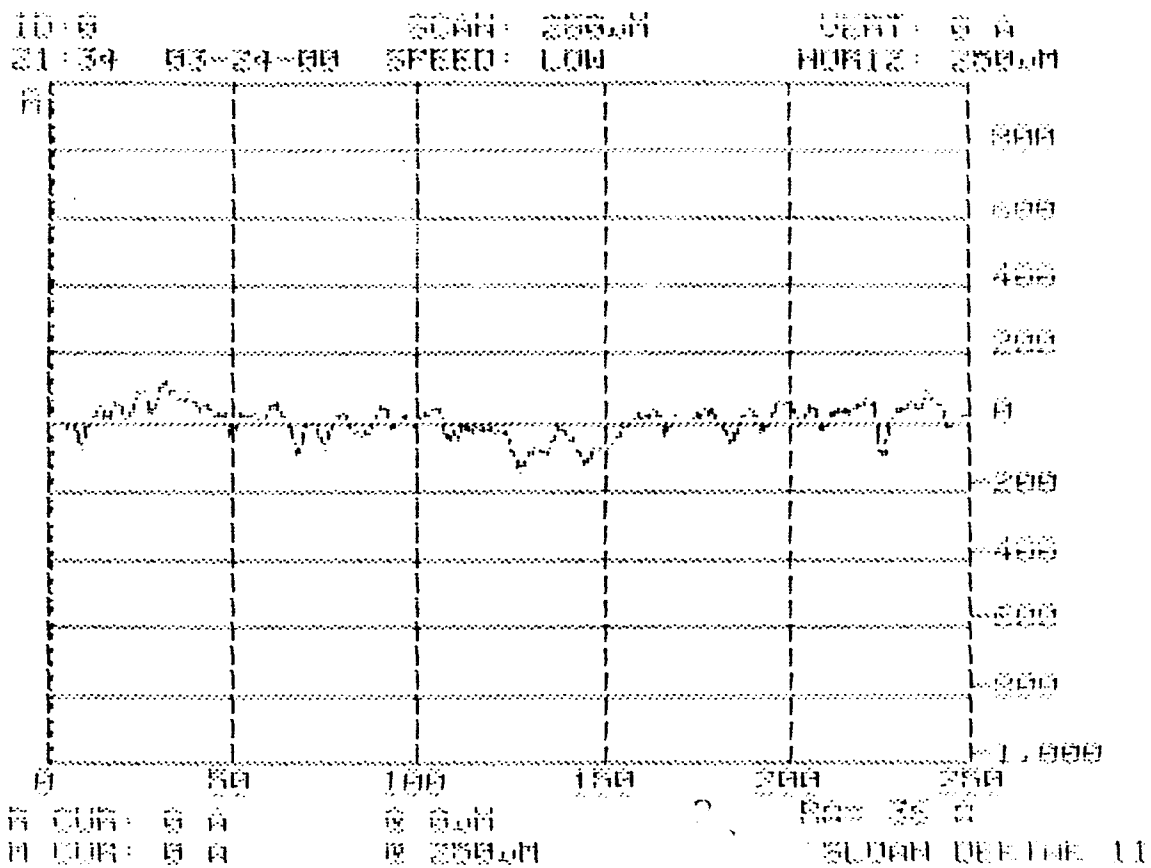Figure 1. Topographical scan (250μm) of a 256^2 SLM device following CMP with the surface variation reduced to 36Å r.m.s..

CONCLUSION

The multi-level metallization of SLMs, to produce high optical quality mirrors, places high demands on the planarization technique which can be satisfied using CMP. A post-processing procedure incorporating CMP was applied to commercially fabricated wafers to add a further level of metal to SLM silicon backplanes, figure 3. This further metal level was optimised optically to produce high quality mirrors. The planarity achieved using CMP demonstrates the capability of this technique to produce surfaces flat to 100Å r.m.s. which will satisfy the future requirements of silicon microfabrication.

Electronically Adressed Ferroelectric Liquid Crystal over Silicon Spatial Light Modulators

D.C. Burns, I. Underwood, A. O'Hara† and D.G. Vass†

Department of Electrical Engineering († Dept of Physics)
The University of Edinburgh, Edinburgh EH9 3JL, UK.
tel +44 31 650 5652, fax +44 31 650 6554, email slm@uk.ac.ed.ee

Abstract

We discuss recent advances in the design and fabrication of electronically addressed ferroelectric liquid crystal over silicon spatial light modulators. We summarise the prospects for further advances in the near future.

BACKGROUND

The Spatial Light Modulator (SLM) is a key component in many optical computing systems. The SLM technology of Ferroelectric Liquid Crystal over Very Large Scale Integrated (FLC/VLSI) silicon has matured considerably over the past few years. Electronically Addressed SLM's (EASLM's) of medium resolution (128^2 pixels or more) have been reported by several labs [1]; systems containing multiple devices have been reported [2]. We are now designing custom devices with system-specific performance characteristics.

INTRODUCTION

FLC/VLSI SLMs are produced by sandwiching a thin layer of FLC between a custom designed silicon backplane and a cover glass coated on the inside with a transparent conductive electrode. Most EASLM backplane designs are based on a pixel circuit consisting of one active element - a Metal-Oxide-Semiconductor Field Effect Transistor (MOSFET) which acts as a switch to control the amount of charge stored on a capacitive storage element - a small metal mirror on the surface of the silicon. The voltage thus generated produces an electric field which alters the state of a thin overlying layer of FLC to produce a binary phase or amplitude modulation in an incident wavefront. This scheme provides the smallest possible pixel and thus the highest density of pixels; it is analogous to the purely electronic Dynamic Random Access Memory (DRAM) cell. This type of pixel suffers from light induced charge leakage which is manifest as a reduction in contrast ratio with incident light intensity thus limitting the maximium light level at which it can be operated. There is also a limit to the spontaneous polarization, P_s, of the FLC material which can be used - the pixel capacitor must store enough charge to switch it. The drive to reduce the pixel size is compromised by the need to maintain an optically-flat metal area to act as the reflective aperture (or mirror). Underlying circuit elements such as transitors or interconnect cause undulations in the overlying part of the mirror; these can cause non uniform optical contrast across a mirror, losses due to scattering and, in coherent systems, phase variations. A flat fill factor (flat mirror area / pixel area) of around 25% has become a de facto minimum.

CURRENT DESIGNS

We have previously reported a 176 x 176 DRAM-type pixel array [3] .The frame rate was limited by the RC time constant of the relatively high resistance polysilicon row access lines. A 512 x 512 pixel array based upon the original 176 x 176 device has been designed. The primary modification has been the use of aluminium to replace the polysilicon row access lines within the pixel array. The finished design has been fabricated by Austria Mikro System and is undergoing electrical testing.

An alternative to the single transistor pixel design above is based around an enhancement of the six transistor Static RAM (SRAM) cell. The enhancement involves inserting a simple logic gate between the memory and the mirror. This allows some of the addressing requirements of the FLC to be met more easily. The SRAM design overcomes all of the above disadvantages of the DRAM pixel at the expense of increased transistor count leading to increased pixel area and decreased chip yield. In particular it maintains its state indefinitely, allows the use of the fast-switching, high P_s, FLC materials and shows no variation of contrast ratio with incident light intensity over a wide range of input intensity. We have demonstrated a fully working 256 x 256 pixel array built in 1.2um CMOS technology. We have also demonstrated the principle of grey scale on this binary device by means of temporal multiplexing of sequential frames.

SILICON FABRICATION ISSUES

We have succesfully applied a post-processing backplane planarisation technique to the 176 x 176 DRAM device which has allowed a flat mirror to be placed on top of the existing circuitry. The technique involves the deposition of a thick dielectric layer which is subsequently polished flat; this allows the deposition of a further metal layer which forms a flat mirror covering almost the entire pixel. The planarisation technique is described in detail elsewhere [4]. LC cell construction has been carried out succesfully on planarised backplanes. The increased flat fill factor of the backplane increases the light throughput of the finished SLM and reduces the stray light reaching the substrate.

SUMMARY AND FORWARD LOOK

Table 1 summarises the current situation. High resolution devices, based on both DRAM and SRAM pixels, have been demonstrated. The pixel designs lend themselves to use in different applications areas. The planarisation process, which significantly enhances device performance, has been demonstrated on one design. It is, in principle, equally applicable to all of the devices of Table 1. (Clearly, to be optically useful, the 512 DRAM device requires to be planarised.)

Given the current (minimum feature size and maximum die size) limitations of using commercial silicon vendors it should be possible to pursue the DRAM technology to 1024 x 1024 and the SRAM to 512 x 512 with frame rates comparable to, or better than, those of Table 1. Beyond that there is a likely need for access to specialised memory fabrication processes.

Device	176 DRAM	512 DRAM[1]	16 SRAM[2]	50 SRAM	256 SRAM
Date	1989	1994	1986	1988	1994
Silicon process	3μmCMOS	3μmCMOS	2μmCMOS	1.5μm nMOS	1.2μm CMOS
Pixel pitch (μm)	30	30	200	72	40
Nominal Mirror size (μm^2)	235 (not square)	145 (not square)	110 x 110	40 x 40	19 x 19
Mirror size (μm^2) after planarization	26 x26	26 x26	196 x 196	NA	36 x 36[3]
Flat fill factor[4] (%) before / after planarisation	26 / 75	16 / 75	30 / 96	31 / NA	23 / 81
Approx electronic address time (μs)	250	1000	1	NA	85
Frame rate (Hz)	1000	250[5]	5[6]	20[6]	4000

[1] undergoing electronic testing at time of writing
[2] currently undergoing re-design from NMOS to CMOS
[3] planarization not yet completed
[4] does not include deductiond for contact or via holes
[5] predicted from SPICE simulations
[6] measured using slow nematic liquid crystal

Table 1. Summary of silicon backplanes for FLC/VLSI EASLM's

REFERENCES

1. K.M. Johnson, D.J. McKnight, I. Underwood, "Smart spatial light modulators using liquid crystals on silicon", IEEE Journal of Quantum Electronics, Vol 29, 699-714, 1993.
2. R. Turner, D.A. Jared, G.M. Sharp and K.M. Johnson, "Optical correlator using very-large-scale integrated circuit / ferroelectric liquid crystal electrically adressed spatial light modulators", Appl. Opt. Vol 32, 3094-3101, 1993
3. 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, S.G. Latham, "A high performance spatial light modulator", Proc. S.P.I.E., Vol 1562, 107-115, 1991
4. A. O'Hara, J.R. Hannah, I. Underwood, D.G. Vass and R.J. Holwill, "Mirror quality and efficiency improvements of reflective mode spatial light modulators by the use of dielectric coatings and chemical-mechanical polishing", Appl. Opt., Vol 32, 5549-5556, 1993.

Improving the performance of liquid-crystal-over-silicon spatial light modulators: issues and achievements

Ian Underwood, David G. Vass, Antony O'Hara, Dwayne C. Burns, Peter W. McOwan, and James Gourlay

The performance of liquid-crystal-over-silicon spatial light modulators has advanced rapidly in recent years. Most progress has centered around new device designs with increased bandwidth. In this paper we report on a number of techniques to improve the optical quality; these have applications in both current and future devices.

Key words: Spatial light modulators, optical quality, ferroelectric liquid crystal, mirror morphology, gray scale, planarization.

1. Introduction

The hybrid technology of liquid-crystal-over-silicon has proved successful within the wider field of spatial light modulators (SLM's). The primary attractions are, first, that both component technologies have a high degree of inherent compatibility, and second, that the industries built around the component technologies, i.e., the Si and the liquid-crystal (LC) industries are each progressing rapidly so that the performance capability of SLM's has increased and can be expected to continue to increase significantly purely as a consequence of improvements in the component technologies. The latter is illustrated clearly by the progress over the past ten or so years in Si, from large-scale integration to very-large-scale integration (VLSI) and beyond, and in LC's, from nematic and related materials to the much faster switching smectic materials and in particular the surface-stabilized ferroelectric liquid-crystal (SSFLC) effect.[1] The attendant progression in SLM capabilities is shown in Table 1.

In addition to the SLM performance gains derived from the mainstream of the two component technolo-

gies, further performance gains are derived from SLM-specific (added-value) research such as custom SLM-specific VLSI (circuit and layout) design. This includes dynamic random-access memory (DRAM) and static random-access memory (SRAM) pixels as well as smart pixels that incorporate local processing or photodetectors. The development of custom SLM-specific VLSI fabrication methods, SLM-specific FLC materials, Si backplane packaging for SLM's, and FLC cell fabrication techniques that are compatible with the Si backplanes permits the enhancement of functionality, performance, and quality.

Most current SLM designs are based on Si backplanes of ~ 1- to 3-μm complementary metal-oxide semiconductors coupled with SSFLC's. We henceforth refer to the technology as ferroelectric liquid crystal over very-large-scale integration, FLC/VLSI. The Si backplanes are generally digital in nature and thus match the inherently binary nature of the SSFLC. The current status of SLM technology is examplified by the 176 × 176 DRAM device shown in Fig. 1.

In judging the performance of FLC/VLSI SLM's, particularly the functionally simple SLM's based on SRAM and DRAM, two main issues merit individual inspection: bandwidth and optical quality. The literature indicates that most progress has been made in the former by increasing the pixel count, the frame rate, or both. In this paper we report on a number of issues that have a direct impact on the latter.

We consider the following topics. In Section 2 we look at recent progress in customizing the fabrication of the Si backplane. The benefits include a high-

Year	Group/Reference	Backplane Technology	Liquid Crystal	Resolution ($n \times n$)	Frame Rate (Hz)	Pixel Circuit
1986	Univ. of Edinburgh[2]	6-μm nMOS[a]	Guest host	16	5	SRAM[c]
1988	Univ. of Edinburgh[3]	1.5-μm nMOS	Twisted nematic	50	60	SRAM
1989	Univ. of Edinburgh[4] STC Technology GEC-Marconi Research Centre	3-μm CMOS[b]	FLC	176	1 K	DRAM[d]
1990	Displaytech[5]	3-μm CMOS	FLC	64	8 K	SRAM
1992	Univ. of Colorado[6] Boulder Nonlinear Systems	2-μm CMOS	FLC	128	5 K	DRAM

[a]nMOS, n-type metal-oxide semiconductor.
[b]CMOS, complementary metal-oxide semiconductor.
[c]SRAM, static random-access memory.
[d]DRAM, dynamic random-access memory.

quality optical finish that aids the alignment of the FLC layer and permits flat pixel mirrors coupled with a high fill factor. In Section 3 we examine a method of producing gray scale by temporal multiplexing on a binary device, thus enhancing its application potential. In electronical addressing of the device, the charge (or dc) balancing condition is maintained to extend the lifetime of the liquid crystal. Finally, in Section 4 we look at the potential of the above described techniques to promote further improvements.

## 2. Custom Silicon Fabrication

A. Procedure

The poor optical quality of the metal that forms the mirrors on the Si backplanes is due primarily to the electronically necessary sinter procedure during the metallization stage of microfabrication. Following the metal-deposition, photolithography, and metal-etch stages of fabrication, the wafers must be sintered to produce good ohmic contacts and improved adhesion of the Al film to the underlying SiO_2.[7] A typical example of this procedure involves the wafers being inserted into a furnace at 435 °C for 20 min in the presence of a forming gas ($40\%H_2/60\%N_2$). During the heating and cooling phases of the operation, stresses are induced in the films, and to relieve these stresses, the Al film distorts, which results in hillocks and depressions being formed.[8,9] The optical efficiency of the mirror is therefore substantially reduced because these irregularities scatter the light incident upon the mirror surface.

To protect the Al film during the sinter procedure, we applied a low-temperature SiO_2 coating to the wafer by using electron cyclotron resonance plasma-enhanced chemical-vapor deposition.

In the course of the anneal operation, the SiO_2 film acts as a constraining layer and, when the anneal is complete, the SiO_2 is removed. Initial work was carried out on test structures[10]; this technique has now been incorporated into the fabrication procedure of two batches of SLM wafers, both fabricated at the Edinburgh Microfabrication Facility (EMF). This

additional procedure permits high electrical performance to be achieved without a reduction in the optical quality of the mirrors. An example of the improvement on mirror quality when the sinter protection layer is used is shown in Fig. 2.

Increasing the pixel mirror fill factor requires multilevel metallization in which the mirror is deposited as the last metal layer on top of the underlying circuitry. Further, this structure eliminates spurious switching of the LC layer that is due to electrical signals on exposed interconnect lines and partially protects the underlying circuitry from incident light that would otherwise lead to increased leakage currents in the circuit. In order to provide a flat surface onto which the mirror can be deposited, the intermediate dielectric must be planarized. Various planarization techniques are used in the semiconductor industry; however, the only method capable of producing the standard of optical flatness necessary here is chemical-mechanical polishing[11] (CMP). Although it is only recently that CMP has been used at this stage in the processing sequence it is already being used in the manufacture of leading-edge technology chips.[12,13] We have applied this technique to enhance our SLM's by adding it as a backend step on fully fabricated wafers. Studies have been carried out on electrically inactive 176 × 176 SLM backplane wafers in which the finished wafer is coated with ~4 μm of electron cyclotron resonance plasma-enhanced chemical-vapor deposited SiO_2 and is then polished to produce an optically flat surface. Metal is then deposited to produce high-quality mirrors. Figure 3(a) shows an area of a 176 × 176 SLM that is then coated with Al to enable the surface flatness to be measured by the use of interferometry [Fig. 3(b)]. We have demonstrated that, by using this technique, we can obtain the surface flatness on the intermediate dielectric that is essential for the fabrication of high-quality mirrors. The main difficulty associated with CMP is maintaining the uniformity of the polishing rate across the wafer.[10] Experimentation has allowed us to improve the control of the polishing
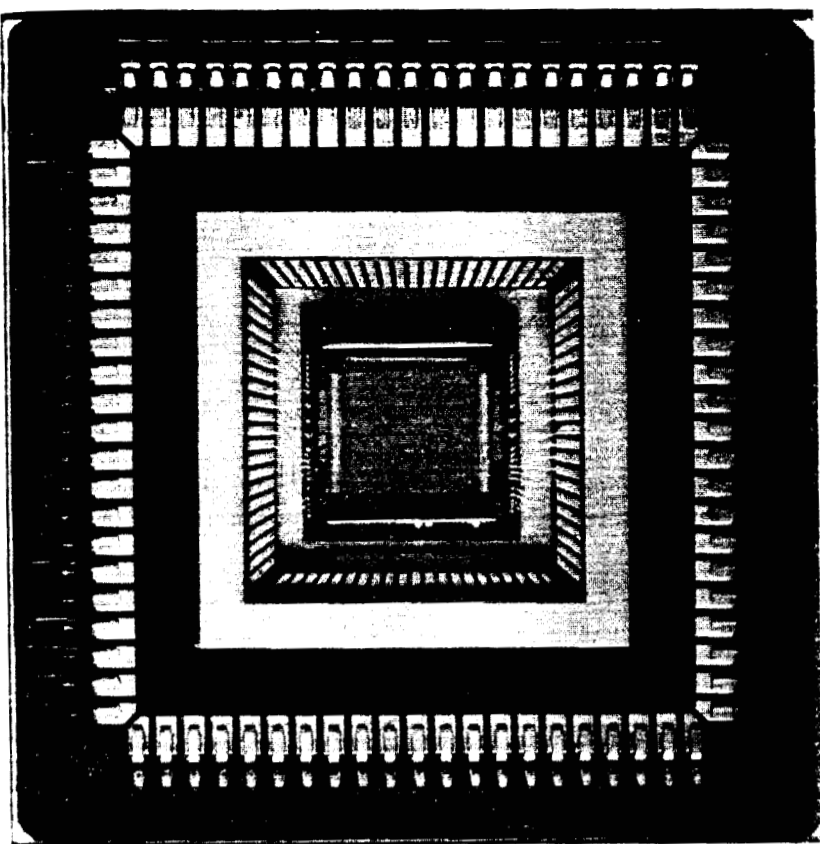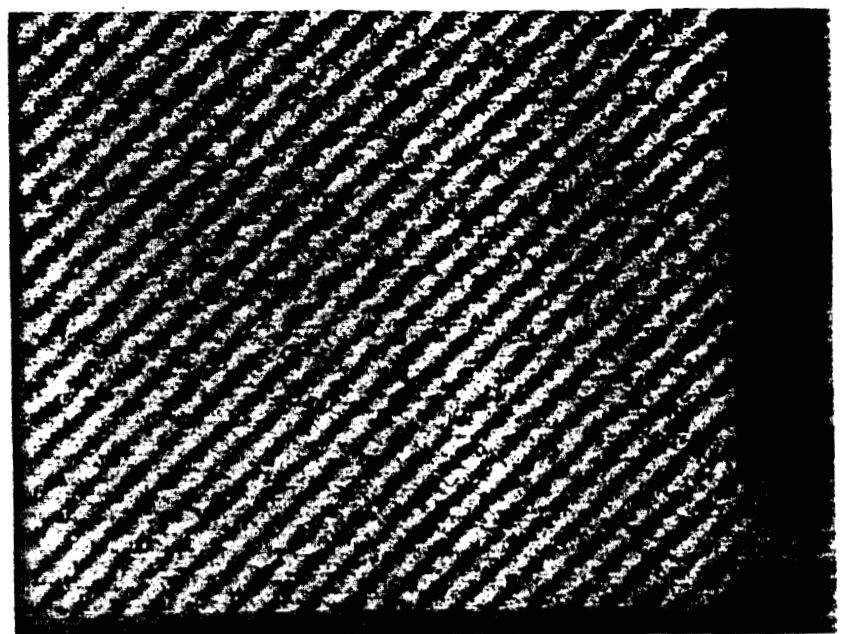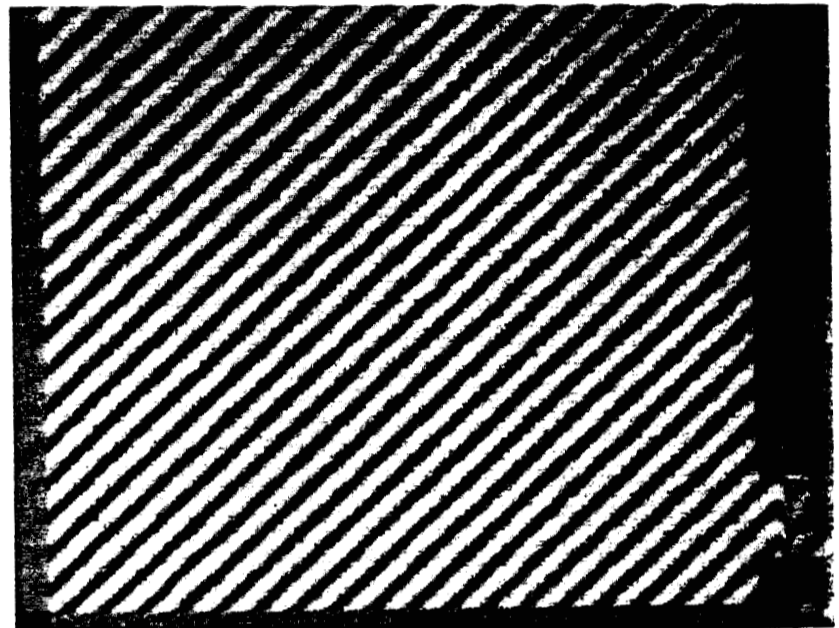
(a)

(b)

Fig. 1. 176 × 176 DRAM SLM (a) device, (b) displayed image courtesy of GEC-Marconi). The backplane of this device is untreated.

(a)

(b)

Fig. 2. Interferograms of pixel mirrors fabricated by the use of (a) standard procedure, (b) metal protection technique.

rate to a point where we are now, for the first time, incorporating CMP into the fabrication sequence of fully operational wafers.

The improvements to the Al mirror surface quality have implications beyond improving the mirror efficiency. Previously, unevenness of the Al surface has had an adverse effect on LC alignment. By improving the mirror surface, the LC alignment has been enhanced (see Subsection 2.B). Our initial experience with cell fabricated by the use of backplanes produced with the hillock suppression technique sug-

gests that LC layer thickness uniformity is greatly enhanced. We propose that this arises because the presence of hillocks causes tilting of the cover glass and a consequence wedging of the LC layer. Further investigation is required before reaching a firm conclusion.

B. Mirror Morphology

Good mirror quality is crucial to enable highly aligned FLC structures. The alignment of FLC molecules is highly dependent on the bounding surface morphology. Irregular hillocks and depressions locally alter the sensitive boundary conditions required of the SSFLC device structure. They disturb the formation of parallel smectic planes, introduce variations in surface molecular tilt angle, and randomly induce all manner of defects from the device ideal.[14] The boundary conditions for the formation of the SSFLC device structure are generally obtained by anisotropic treatment of the bounding plates. Surface energy is transferred to the bulk FLC medium by elastic forces. The surface treatment usually takes the form of the
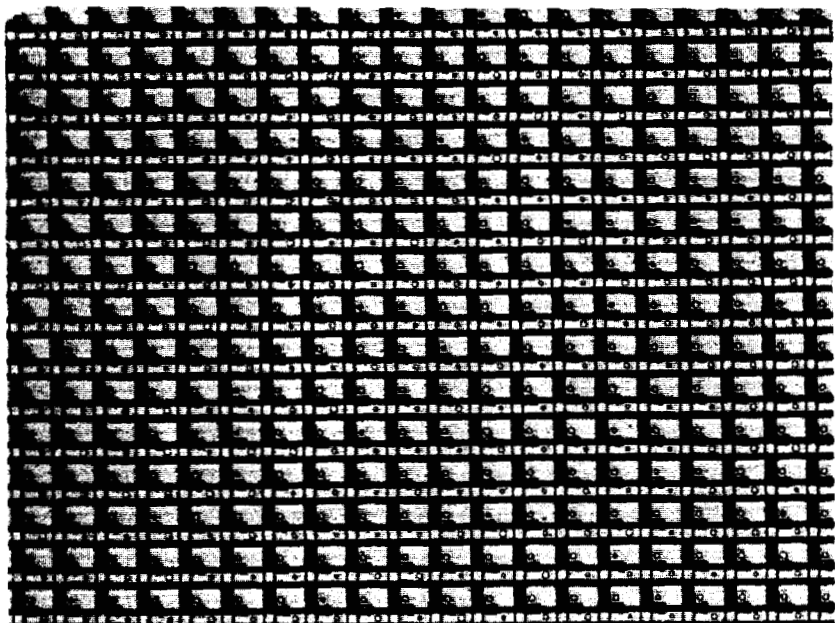
(a)
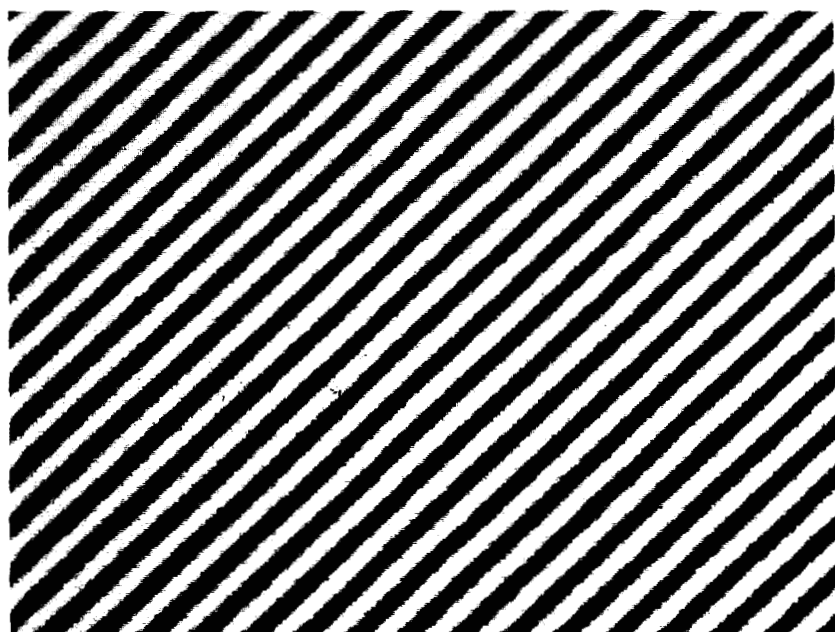
(b)

Fig. 3. Effect of CMP technique for backplane planarization: (a) Part of 176 × 176 SLM backplane with an area of ~600 μm × 450 μm, (b) interferogram of part of a polished wafer following metallization. Magnification is the same as (a).

Fig. 4. SEM picture of alignment structure produced by oblique SiO_x evaporation.

creation of microrelief structures. A common technique for the generation of such structures is oblique evaporation of thin films such as SiO_x that can generate a variety of anisotropic surface structures such as ridges and tilted columns.[15] This surface energy should be uniform across the bounding plates of the FLC device. If large enough, any perturbation from this uniformity affects the molecular alignment in the bulk FLC material. Therefore mirror quality directly determines the defect structure. It is unlikely that surface treatments can compensate for large fluctuations in the surface energy caused by poor mirror quality.

Defect-free FLC structures should be obtainable when high-quality mirrors can be constructed on VLSI Si backplane SLM devices. Such improvements are significant steps toward increases in contrast ratio, device uniformity, modulation efficiency, and reflectivity. This improvement is demonstrated by a comparison of FLC devices fabricated on a standard (unprotected metal) VLSI process and those
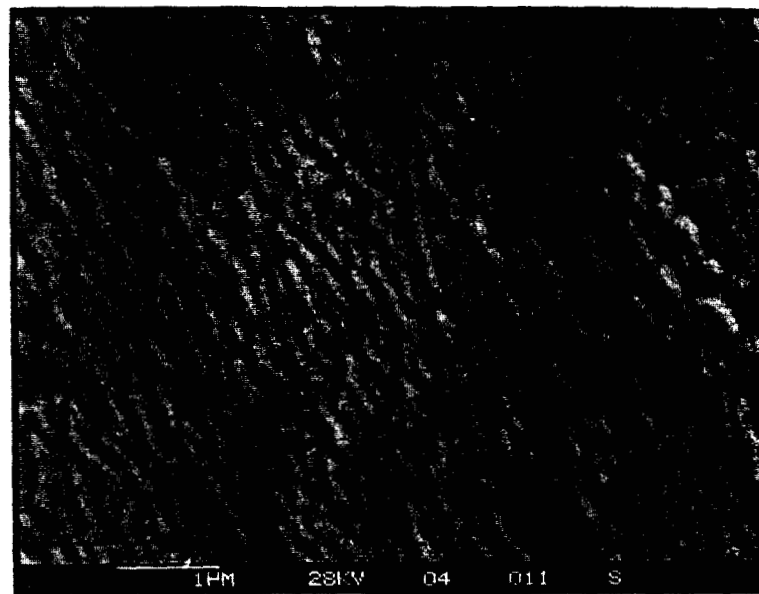
fabricated on a protected metal process, as described above. The protected metal process results in a reduction of hillock and depression formation, although some of these undesirable features can still be observed. It was expected that such an improvement to the mirror quality would result in more favorable surface conditions to attain good FLC alignment. This has been observed experimentally. SSFLC cells were constructed on two SLM devices from the same VLSI process. The SLM backplane used was a 16 × 16 test-bed device.[2,16] The wafers were processed at the EMF. One device originated from a wafer that used the standard, unprotected metal process. The other originated from a wafer fabricated with the extra metal protection processing step described at the beginning of Subsection 2.A. The SSFLC device alignment layers were obliquely evaporated with SiO_x to a thickness of 50 nm on both the SLM chip and the front cover glass with an indium tin oxide transparent–conducting electrode. The SiO_x was obliquely evaporated at an angle of 30° to the substrate. This is commonly referred to as medium-angle deposition.[15] This alignment layer technique results in a ridge structure with a period of ~100 nm running perpendicular to the direction of evaporation, which has been observed by a scanning electron microscope (SEM) and is shown in Fig. 4. This structure gives relatively strong surface anchoring of the FLC molecules. A zig-zag defect structure is usually observed in FLC test cells, suggesting the presence of the chevron defect. Small-angle deposition has been shown to give improved alignment with FLC devices.[15] But oblique evaporation of SiO_x at ~5° will have major shadowing problems because of the topography of the surrounding circuitry. Only fully planarized SLM backplanes would permit problem-free oblique evaporation at small angles on pixel mirrors. Alternatively the small-angle deposition technique could be applied to only the front cover glass electrode, and this would result in asymmetric boundary conditions on the FLC device. A study into such asymmetric alignment layers would be of considerable interest.
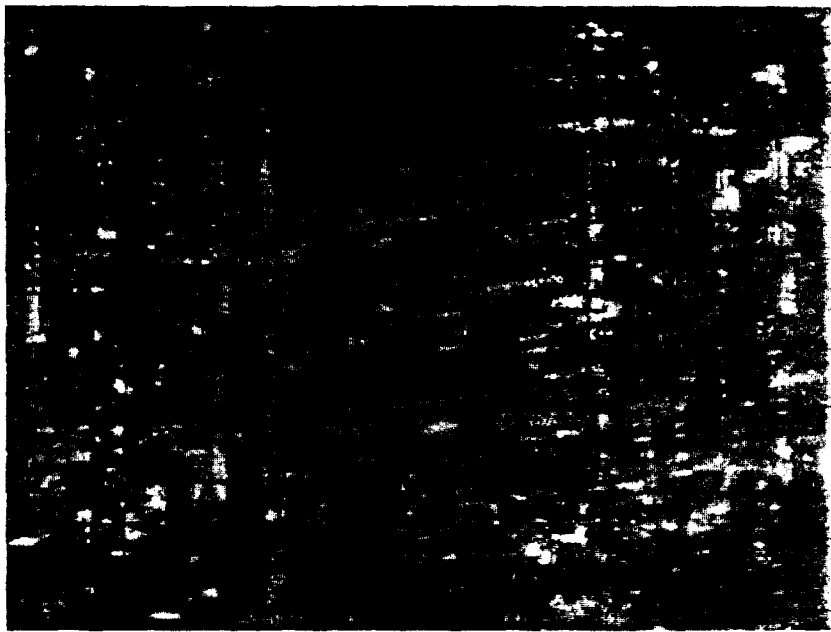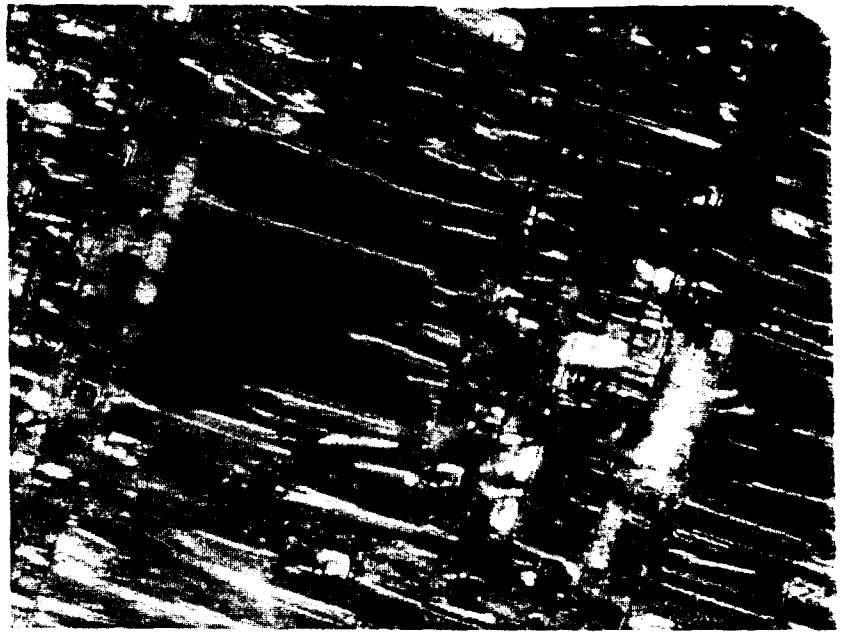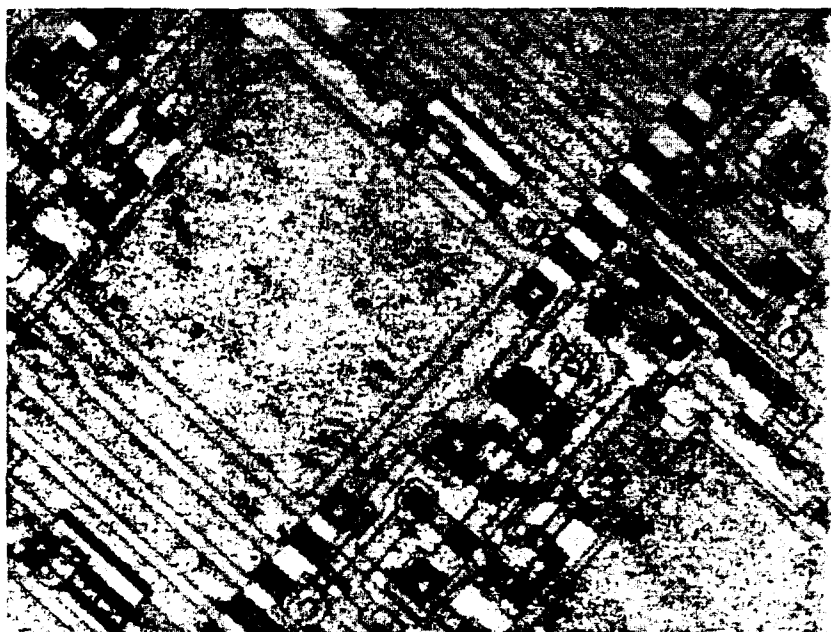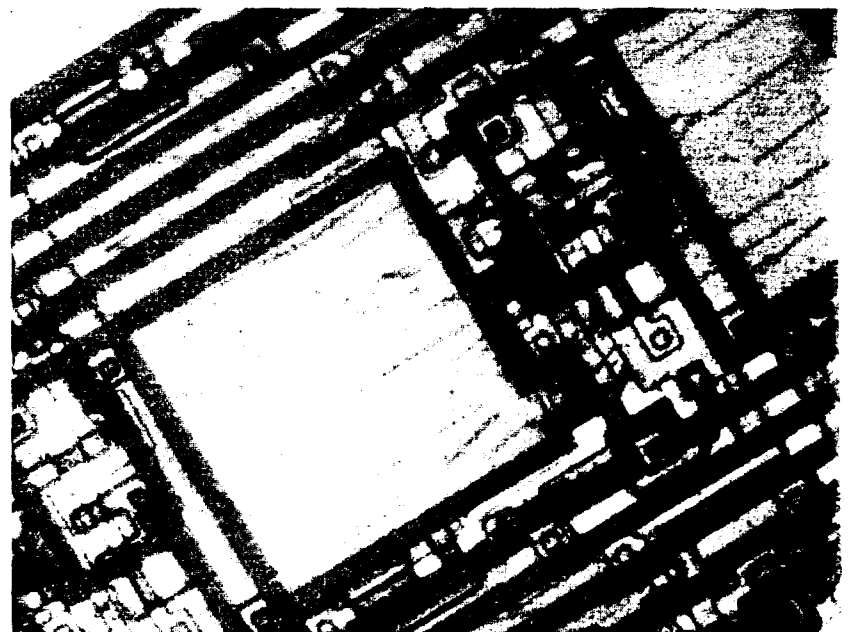
Fig. 5.   LC structure on pixel (a) untreated. OFF; (b) untreated. ON; (c) treated. OFF; (d) treated, ON.

An FLC cell was constructed on the SLM chip with a 1-μm cell gap between the mirrors and the conducting front electrode. The device is therefore optimized to operate in reflection as a switchable half-wave plate with smectic C* FLC material exhibiting a 22.5° cone angle for light with a wavelength of approximately 633 nm (He–Ne). The device was filled with SCE13 FLC material (Merck–British Drug House) under vacuum at an elevated temperature of 120 °C.   Both protected metal and unprotected metal devices were constructed under the same conditions. Therefore these devices permitted comparisons to be made to show any improvement to the FLC characteristics.   Figures 5(a) and 5(b) show the unprotected metal device illuminated with polarized light and observed through cross polarizers, respectively. Figures 5(c) and 5(d) show the protected metal device under similar circumstances.   On the unprotected metal device one observes a lower reflectivity and apparently no domain structure.   There are domains, however, but they are so numerous and fine (because of the poor metal quality) that they appear uniform.   This is not a desirable feature as they will

still adversely affect performance and be difficult to remove.   In the metal protected sample, domains are large and clearly visible.   This is actually encouraging as such distinct domains are observed in glass-on-glass FLC test cells.   Therefore the alignment of the FLC in the metal protected device is clearly approaching the level attainable in high surface quality test cells.   We conclude that the high performance of FLC structures observable in test cells should be possible through improved metal quality VLSI backplane SLM's.

3.   Temporal Multiplexed Gray Levels

Gray-level SLM's have many potential uses, including display applications and as adaptive weight planes in optoelectronic neural networks.   As described in the introduction, most current FLC/VLSI SLM's operate in a binary mode.   However, the microsecond switching times of the FLC and the high addressing speeds of Si backplanes permit temporal multiplexing to be used as a means of real-time gray-level generation.
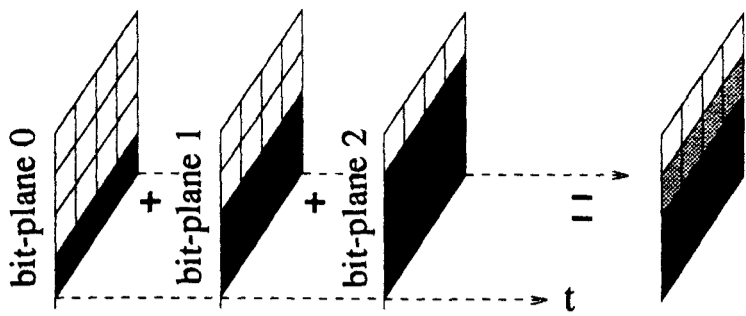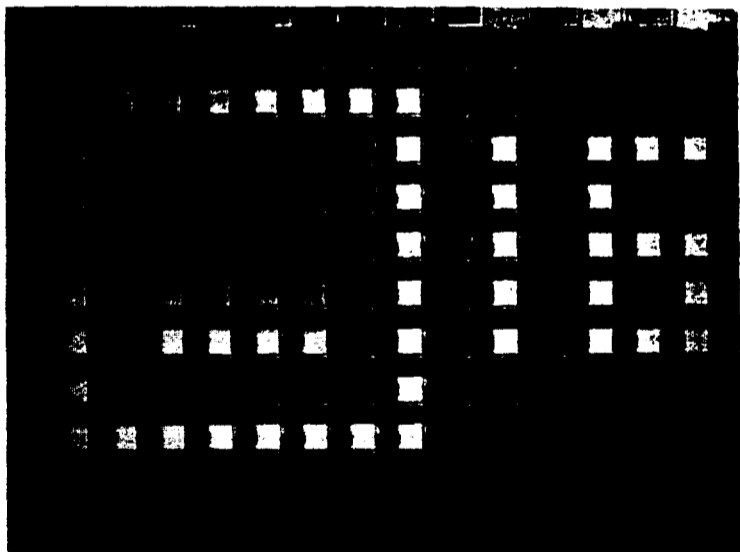
Fig. 6. Four-gray-level image obtained by the use of a linear encoding algorithm.

Temporal multiplexing involves splitting each gray-level image into a number of time-sequential sub-frames. Each subframe has an associated bit-plane that is generated by an approximate encoding algorithm. These bit-planes are then sequentially scanned into the SLM to produce the desired gray-level image.

A simple linear encoding algorithm has been used to create the bit-planes. Figure 6 shows how a four-gray-level image is built up. For the 16 gray levels illustrated in Fig. 7, 15 subframes were required. Each subframe lasts 500 μs, resulting in a frame rate of $\sim$ 133 Hz.

The linear encoding algorithm above is satisfactory for generating a small number of gray levels; however, each extra gray level requires an extra subframe. The frame rate is inversely proportional to the number of gray levels. A more efficient binary weighted encoding algorithm with a pulsed light source, as illustrated in Fig. 8 is now being investigated. In this case, n subframes can generate 2^n gray levels. So, for example, 256 gray levels can be generated at 125 Hz when a 1-ms subframe duration is used.

Fig. 7. Photo showing 16 linearly encoded temporal multiplexed gray levels. The key shows the gray level present in the central square of 4 × 4 pixels.
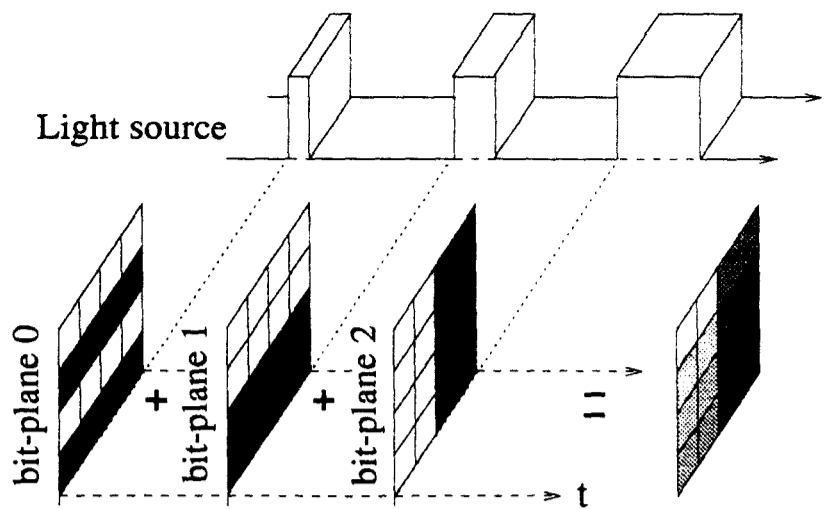

Fig. 8. Eight-gray level image obtained by the use of a binary encoding algorithm and a pulsed light source.

A. Advantages of Surface-Stabilized Ferroelectric Liquid Crystals for Producing Gray Scale

There are, in principle, several advantages to producing gray scale as described above:

(1) The fast switching speed of the FLC should eliminate the problems of shadowing and ghosting of a fast moving pattern that are associated with slow switching nematic-LC displays.

(2) In contrast to the gray levels produced by applying *analog* electrical signals to nematic LC's, there should be no phase variations associated with the gray levels produced by temporal multiplexing in which the applied electrical signal is *digital*. This makes the proposed scheme particularly suitable for use in coherent optical systems.

B. Direct-Current Balancing

It is widely accepted that to prevent chemical degradation of the FLC layer, the electric field across it should be properly dc balanced. In DRAM-type devices, this has been achieved by scanning in a bit-plane; then, after the FLC is allowed to settle and the read beam is pulsed, the exact inverse of the bit-plane is scanned in with the front electrode signal toggled.[4] The resulting image duty cycle was less than 5%.

The functionality of some SRAM-type pixels, such as those on the 16 × 16 SLM, is enhanced by the inclusion an XNOR gate.[2] The inputs to the XNOR are the pixel latch and a global clock signal. One can achieve dc balancing simply by continually toggling both the global clock and the front electrode. The resulting image duty cycle is 50%. A further advantage over the DRAM pixel is the availability of an effectively unlimited amount of charge to switch the FLC. This is discussed elsewhere.[6]

4. Conclusions and Future Developments

In this paper we began by identifying that improving the optical performance of FLC/VLSI SLM's is a critical issue if their success is to be continued.

We have demonstrated several techniques that, when applied to the Si backplane, boost the optical quality of finished devices. Custom sintering of the

metal layer has facilitated much improved FLC alignment and a consequent increase in optical contrast and uniformity. CMP has allowed the planarization of electrically inactive Si backplanes to optical flatness. This makes possible, for what is, to our knowledge, the first time, pixels that exhibit both optical flatness and a high fill factor. This research will continue with the application of the techniques to fully operational Si wafers, thereby permitting the improvements to be translated into fully working FLC/VLSI devices.

We have demonstrated a means of producing gray scale on an inherently binary device by means of temporal multiplexing. The drive scheme for the device maintains the dc balance condition that is essential for device longevity. The next steps here are to apply a binary weighting scheme to the subframes to permit the production of 2^n gray levels from n subframes and to prove by experiment the hypotheses made in Subsection 3.A, namely, that such a system eliminates shadowing and ghosting and that the gray levels produced are phase free.

The techniques described are applicable to all FLC/VLSI backplanes and will be used to enhance the performance of the next generation of SLM's based on backplanes, recently designed at The University of Edinburgh and currently in fabrication, which include a 256 × 256 SRAM backplane and a 512 × 512 DRAM backplane.

References

1. N. A. Clark and S. T. Lagerwell, "Submicrosecond bistable electro-optic switching in liquid crystals," Appl. Phys. Lett. **36**, 899–901 (1980).
2. I. Underwood, D. G. Vass, and R. M. Sillito, "Evaluation of an nMOS VLSI array for an adaptive liquid-crystal spatial light modulator," Proc. Inst. Electr. Eng. Part J **133**, 77–82 (1986).
3. D. J. McKnight, D. G. Vass, and R. M. Sillitto, "Development of a spatial light modulator: a randomly addressed liquid crystal over nMOS array," Appl. Opt. **28**, 4757–62 (1989).
4. I. Underwood, D. G. Vass, R. M. Sillito, G. Bradford, N. E. Fancey, A. O. Al-Chalabi, M. J. Birch, W. A. Crossland, A. P. Sparks, and S. G. Latham, "A high performance spatial light modulator," in *Devices for Optical Processing,* D. M. Gookin, ed., Proc. Soc. Photo-Opt. Instrum. Eng. **1562**, 107–115 (1991).
5. L. K. Cotter, T. J. Drabik, R. J. Dillon, and M. A. Handschy, "Ferroelectric-liquid-crystal silicon-integrated-circuit spatial light modulator," Opt. Lett. **15**, 291–293 (1990).
6. K. M. Johnson, D. J. McKnight, and I. Underwood, "Smart spatial light modulators using liquid crystals on silicon," IEEE J. Quantum Electron. **29**, 699–714 (1993).
7. S. Wolf, *Silicon Processing for the VLSI Era* (Lattice, Sunset Beach, Calif., 1990), Vol. 2, pp. 110–116.
8. P. Chaudhari, "Hillock growth in thin films," J. Appl. Phys. **45**, 4339–4346 (1974).
9. D. S. Herman, M. A. Schuster, and R. M. Gerber, "Hillock growth on vacuum deposited aluminum films," J. Vac. Sci. Technol. **9**, 515–519 (1972).
10. 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," Appl. Opt. **32**, 5549–5556 (1993).
11. S. Sivaram, H. Bath, R. Leggett, A. Maury, K. Monnig, and R. Tolles, "Planarizing interlevel dielectrics by chemical-mechanical polishing," Solid State Technol. **35**, 87–91 (1992).
12. H. Landis, P. Burke, W. Cote, W. Hill, C. Hoffman, C. Kaanta, C. Koburger, W. Lange, M. Leach, and S. Luce, "Integration of chemical mechanical polishing into CMOS integrated-circuit manufacturing," Thin Solid Films **220**, 1–7 (1992).
13. W. J. Patrick, W. L. Guthrie, C. L. Standley, and P. M. Schiable, "Application of chemical-mechanical polishing to the fabrication of VLSI circuit interconnections," J. Electrochem. Soc. **138**, 1778–1784 (1991).
14. N. A. Clark and S. T. Lagerwall, "Surface-stabilized ferroelectric liquid crystal electro-optics: new multistate structures and devices," Ferroelectrics **59**, 25–67 (1984).
15. D. Armitage, "Ferroelectric liquid crystal alignment by oblique evaporation of SiO_x," Ferroelectrics **122**, 239–252 (1991).
16. J. Gourlay, P. McOwan, D. G. Vass, I. Underwood, and M. Worboys, "Optical engineering aspects of ferroelectric liquid crystal over silicon spatial light modulators," in *Proceedings of the Conference on Applied Optics and Opto-electronics* (Institute of Physics, London, England, 1992), pp. 211–213.

Issues in Improving the Performance of Liquid-Crystal-over-Silicon Spatial Light Modulators

I. Underwood, D.G. Vass[†],A. O'Hara[†] and D.C. Burns

Edinburgh Microfabrication Facility
Department of Electrical Engineering
(† Applied Optics Group
Department of Physics)
The University of Edinburgh
The King Buildings
Mayfield Road
Edinburgh EH9 3AF

1. Introduction

The marriage of the two technologies of Very Large Scale Integrated (VLSI) silicon backplane active drive and Ferroelectric Liquid Crystal (FLC) light modulating layers has attracted much attention as a promising means of implementing compact high performance, smart and advanced Spatial Light Modulators (SLMs)[1]. Two of the attractions of the hybrid technology are (i) that the two component technologies show a significant degree of compatibility, and (ii) that vast resources are being poured into these component technologies so that the performance of VLSI/FLC SLMs can be expected to continue to improve at a useful rate simply as a consequence of improvements in the component technologies[2].

The ready availability of both user friendly Computer Aided Design (CAD) tools for VLSI design and relatively cheap silicon fabrication for prototype quantities of devices have combined to make the technology accessible to a wide range of potential users so that, to date, many new devices have been reported[3]. In fact, most reports concentrate on the description of either larger/faster arrays of simple pixels or new (smart) pixel designs; progress in these areas has been rapid while, in the authors' opinion, the rate of improvement in the quality of the SLMs as measured by factors such as contrast ratio and uniformity has failed to keep pace.

Some of the issues which must be addressed in order to maximise the optical quality of VLSI/FLC SLMs include (i) optimisation of the optical properties of the silicon backplane , (ii) consistent repeatable production of a high quality Surface Stabilised Ferroelectric Liquid Crystal (SSFLC) layer, (iii) optimal electronic addressing of the SSFLC by the backplane, and (iv) the production of grey scale from the inherently binary SSFLC effect. In this paper we discuss the initial approach we have taken, at The University of Edinburgh, to tackling these issues.

2. Custom Processing of the Silicon Backplane

Here, we present a brief description of two of the primary deficiencies of "standard" silicon when applied to driving an overlying FLC layer; we discuss the consequences of the deficiencies and present results of our efforts, within the Scottish Collaborative Initiative in Optoelectronic Science (SCIOS), to overcome them.

1. The surface profile of the layers making up a pixel circuit typically shows layer boundary step heights of up to the order of 1μm. This leads to intra-pixel variations in the amplitude and phase of light reflected from a metal pad covering the whole pixel and to variations in the thickness of the FLC layer The alternative approach of sharing the pixel real estate between un-flat circuit and flat pad leads to a lower optical fill factor. The preferred solution is planarisation of the dielectric layer underlying the metal pad. Three options have been explored - spin on polyimide , Electron Cyclotron Resonance (ECR) oxide deposition and chemical-mechanical polishing. The last of these techniques has been tried succesfully[4] on an electrically defective SLM wafer. An optical interferogram of part of a polished wafer is shown in Figure1(B). Operational SLM wafers are currently being polished; it is anticipated that these will produce "optically flat" working pixel arrays with a high pixel fill factor (>90%).

2. The granular nature of the thick metal layer gives rise to light scattering and leads to poor FLC alignment and bistability."Hillocks" which are present in the metal layer also scatter light and can be high enough to span the FLC layer causing short-circuits between the pad and the ITO counter electrode; furthermore they may be the cause of uneven LC layer thickness. The electrically necessary step of sintering the aluminium is largely responible for the poor optical quality of the metal. We have explored the use of dielectric over-layers during sintering which have resulted in a significant improvement in the smoothness of the final sintered metal surface as shown in Figure 2. Recent results also indicate an improvement in LC alignment and layer thickness uniformity as a result of the high quality metal finish.
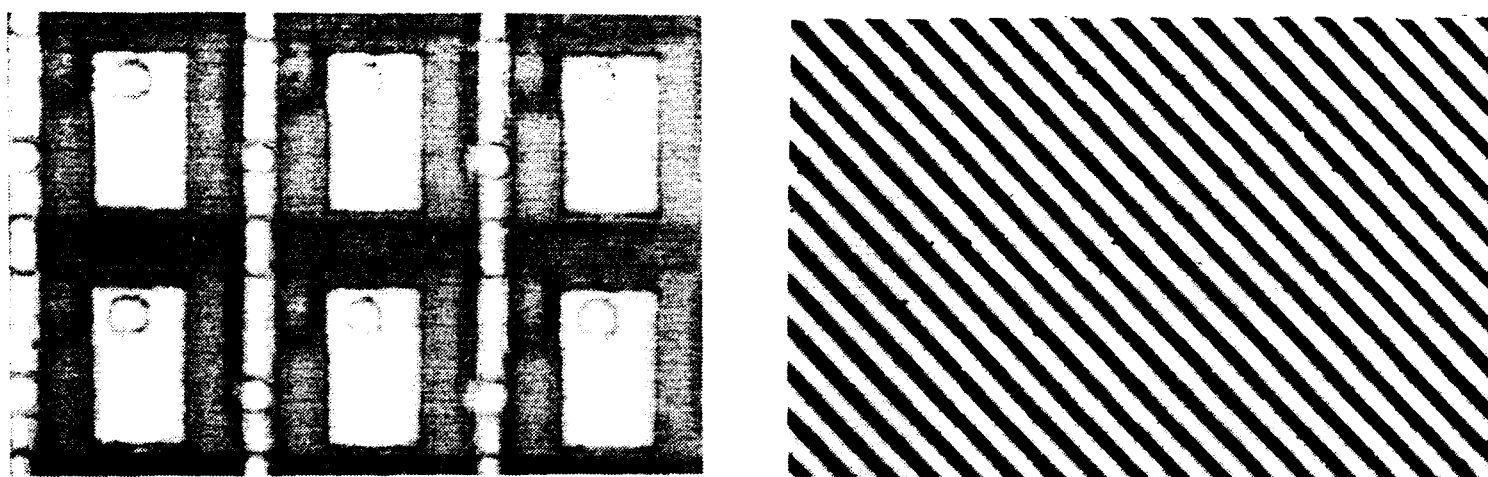

Figure 1. Left - photograph of a portion of a 176^2 pixel array (30 μ m pixel pitch).
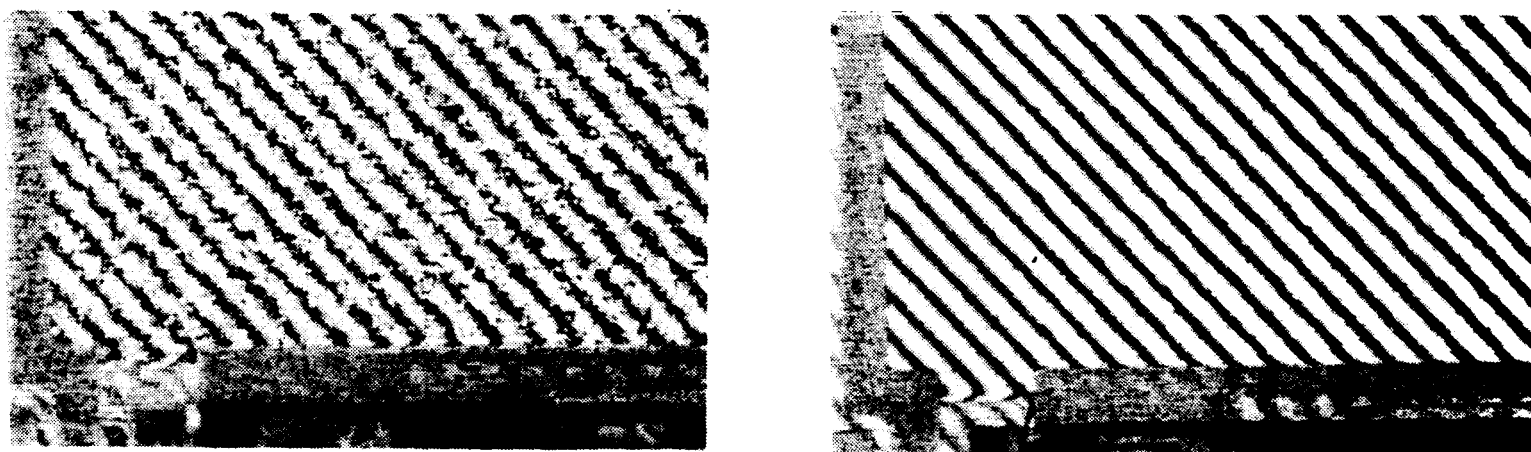**Right** - optical interferogram (λ =564nm) of portion of the array after planarisation

Figure 2. **Left** - Optical interferogram (λ=564nm) of part of a metal pad (approx 74x50 μ m^2) following standard sintering. **Right** - Optical interferogram of similar metal pad after modified sintering operation. An ECR oxide layer of thickness 0.3 μ m was present during sintering.

3. Electronic Addressing of the SSFLC layer

It is generally accepted that LC longevity requires that no net long term DC field is allowed across the layer. This leads to complications in addressing the bistable SSFLC based SLMs. The High Performance (HP)SLM 176^2 device[5] solves this problem by building each frame from two sequential complementary sub-frames in a system where a pulsed light source illuminates only the positive sub frame. We have devised (in collaboration with BNR Europe) a 2 transistor pixel[6] which is capable of both DC balancing and a high (approaching 100%) duty cycle. We have also developed drive schemes[7] which allow a modified-SRAM pixel to supply a very stable drive signal to a SSFLC layer as demonstrated in §4 below.

4. Grey scale in SSFLCs

Although inherently binary in nature, the switching speed of SSFLCs is such that effective grey scale can be achieved by temporal multiplexing of sequential bit planes while maintaining a significant speed advantage over the older analogue nematic LCs. Several options exist, including
      1. Equal weight subframes
      2. Binary weighted sub-frame duration with constant illumination intensity
      3. Binary weighted illumination intensity with constant sub-frame length
      4. A combination of 2 and 3 (for increased dynamic range).

We are currently active in investigating the first three and present here, in Figure 3, preliminary results from the first of these[8]. The microdensitometer trace shows a fairly linear range of grey scale on part of a 16x16 pixel SLM.(It also shows significant non uniformity in transmitted intensity across individual pixels - an example of the problem we are tackling as described in §3 above.
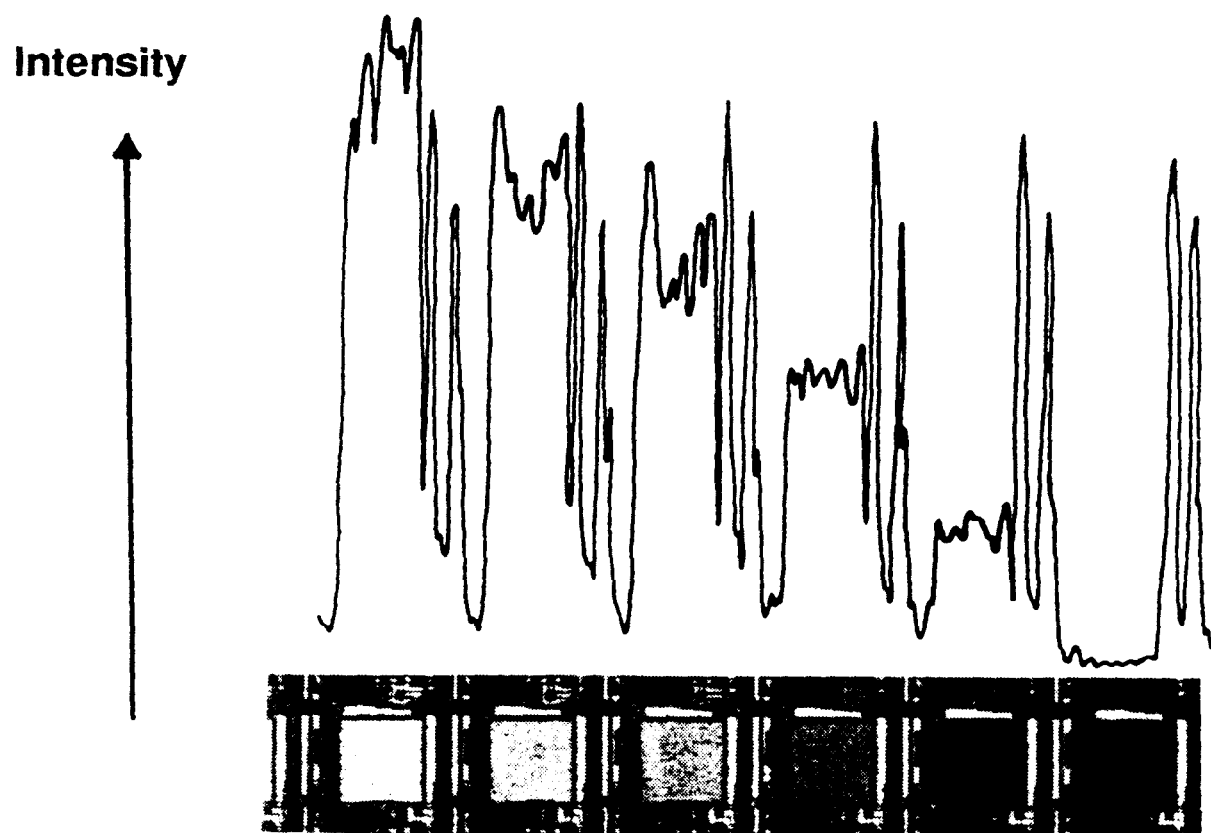
Intensity

Figure 3. Top - Photograph of a line of pixels (pitch 200μ m) showing grey scale achieved by time multiplexing **Bottom -** corresponding microdensitometer trace.

5. Conclusions

We have presented an analysis of some issues which must be addressed in order to maximise the potential of the VLSI/FLC SLM technology. Preliminary results are very encouraging and indicate the potential for significant performance improvements.

6. References

1. N Collings, WA Crossland, PJ Ayliffe, DG Vass and I Underwood, *Evolutionary Development of Advanced Spatial Light Modulators,* Appl Opt , Vol **28**,1989, pp4740-7.
2. KM Johnson, DJ McKnight and I Underwood, *Smart Spatial Light Modulators Using Liquid Crystal on Silicon,* JQE 1993, to be published.
3. See for example Optical Society of America Conference on Spatial Light Modualtors, Lake Tahoe, Nevada, Septemeber 1990, OSA Technical Digest Series 14.
4 A O'Hara, RJ Hannah, RJ Holwill, I Underwood and DG Vass, *Improvements in the Mirror Quality of Reflective Spatial Light Modulators using Dielectric Coatings and Mechanical Polishing,* in Optical Interference Coatings Technical Digest 1992, (Optical Society of America, Washington DC, 1992) Vol **15**, pp274-6.
5. I Underwood, DG Vass, RM Sillitto, G Bradford, NE Fancey, AO Al Chalabi, MJ Birch, WA Crossland, AP Sparks and SG Latham, *A high performance Spatial Light Modulator,* in Devices for Optical Processing, DM Gookin, Ed, Proc SPIE 1562, pp107-15.
6. WA Crossland, MJ Birch, AB Davey and DG Vass, *Active Backplane Spatial Light Modulators using Chiral Smectic Liquid Crystals,* Proc SPIE/IS&T Symp on Electronic Imaging Science and Technology, San Jose, CA, February 1992, paper 1665-36, p143.
7. I Underwood and G Bradford, *Driving FLC SLMs using SRAM backplanes,* Edinburgh University Internal Report, January 1990 (available from authors).
8. DC Burns and I Underwood, *Real Time Grey Level Generation using a Binary Mode Spatial Light Modulator,* In Applied Optics and Optoelectronics Technical Digest 1992, Leeds, England, September 1992, paper OE P6, pp214-5.