Holographic Routing Network for Parallel Processing Machines

E.S. Maniloff, K.M. Johnson

Center for Optoelectronic Computing Systems, University of Colorado, Boulder
Campus Box 425, Boulder, CO 80309-0425

and

J. Reif

Computer Science Department, Duke University
Durham, North Carolina 27706

ABSTRACT

Dynamic holographic architectures for connecting processors in parallel computers have been generally limited by the response time of the holographic recording media.[1-5] In this paper we present a different approach to dynamic optical interconnects involving spatial light modulators (SLMs) and volume holograms. Multiple-exposure holograms are stored in a volume recording media, which associate the address of a destination processor encoded on a spatial light modulator with a distinct reference beam. A destination address programmed on the spatial light modulator is then holographically steered to the correct destination processor. We present the design and experimental results of a holographic router for connecting four originator processors to four destination processors.

1. INTRODUCTION

Message routing in a parallel machine concerns providing arbitrary interconnections between its processors. The Connection Machine for example, is a 65,536 processor bit serial SIMD parallel machine, requiring 65,536 messages to be routed to distinct addresses. There is a bottleneck in this information transfer mechanism: The routing time in these parallel machines is approximately a thousand times longer than the machine instruction time. Optical hardware provides the potential for high bandwidth, low crosstalk and power dissipation for connecting processors at the board-to-board level.[6] It has also been shown that impedance matching requirements favor optics over electronics for fast data transfer.[7]

Previous work on dynamic optical interconnects has employed spatial light modulators (SLMs) in optical crossbars[8-10], or volume holograms[1-5,11] to reconfigure connections in real-time. These two approaches have disadvantages: The former requires setting N^2 switches to achieve the interconnections, while the latter is limited by the slow response time of photorefractive recording materials.

In this paper we present a novel optical interconnect architecture, which uses spatial light modulators to dynamically control the holographic routing of messages between originator and designation processors. This system is not limited by the response time of the volume holographic recording media, which stores the destination address: The routing is achieved as fast as the optical beam can be modulated by the SLM. Our first prototype holographic router uses ferroelectric liquid crystal (FLC)[12-14] SLMs to connect four originator processors to four destination processors at 10 kHz. We also present preliminary results on reducing the number of switches in the SLM required to route N originator processors to N destination processors in a single time step.

2. HOLOGRAPHIC SELF-ROUTING ARCHITECTURES

The FLC SLM and $LiNbO_3$ photorefractive recording material are the essential components in both architectures. We will first briefly describe the operation of the FLC SLM and the photorefractive effect. The specific architectures will then be presented.

The operating characteristics of surface-stabilized FLCs have been extensively described elsewhere.[12-14] Briefly, the surface stabilized FLC SLM consists of ~2μm of FLC material placed between glass plates coated with patterned transparent conductive oxide (TCO) electrodes. Voltages of opposite sign applied perpendicular to the glass plates cause the macroscopic polarization $\vec{P}$ to align with the electric field. This in turn causes the optic axis $\vec{n}$ to switch between two states, which are oriented $\pm\Psi$ with respect to the layer normal $\vec{z}$ (see Fig. 1). Angle Ψ is referred to as the molecular tilt angle, which is a material property of the FLC and can be approximately 22.5° for many commercially available mixtures [15]. Applying ±10V/μm across the device will rotate $\vec{n}$ by 45°. When the thickness of the material $d = \lambda/2\Delta n$, where Δn is the material birefringence and λ is the wavelength of light in vacuo, the device acts like a programmable half wave plate.

In the holographic router, the FLC SLM is placed between crossed polarizers and operates as an intensity modulator. In the OFF state, the optic axis is aligned with the incident polarizer. No rotation of incident laser light occurs, and no light is transmitted by the output analyzer. Switching the polarity of the applied electric field rotates the optic axis by 45°. Incident vertically polarized light is rotated by 90° to the horizontal and transmitted by the output analyzer. The contrast ratio T_{ON}/T_{OFF} is approximately 25:1 for the SLM

Preprint of paper appearing in Holographic Optics II: Principals and Applications, G. Michael Morris; Ed. SPIE Proceedings Series, Vol. 1136, EPS/EUROPTICA/SPIE International Congress on Optical Science and Engineering, Paris, France, April 1989, pp. 283-289.

used in these experiments.

Light transmitted by the FLC SLM is thus horizontally polarized and interferes with a reference beam, also horizontally polarized. The interference patterns between the N destination addresses and N angularly multiplexed reference beams are recorded in the volume photorefractive media.

The photorefractive effect refers to an intensity dependent change in the index of refraction of a material.[16] In $LiNbO_3$, the photorefractive crystal used in these experiments, photo-generated charge carriers are redistributed from regions of high intensity illumination to low intensity regions. This transport of charges results in a space charge field oriented in the opposite direction. The change in index Δn is proportional to the space charge field,

$$\Delta n = \frac{1}{2} n^3 r_{33} E_{sc}, \qquad (1)$$

where n is the unperturbed index of refraction, r_{33} is the large electro-optic coefficient in $LiNbO_3$ and E_{sc} is the internal space-charge field. The spatially varying field E_{sc} modulates the index of refraction of the material. The dependence of diffraction efficiency on Δn for a hologram stored in a thick lossless media is approximated by the well known Kogelnik expression,

$$\eta = \sin^2\left(\frac{\pi \Delta n d}{\lambda \cos\theta}\right), \qquad (2)$$

where d is the crystal thickness and θ is half the reference object beam angle. Reference [16] contains a complete treatment of dynamic holographic recording in photorefractives.

2.1 ARCHITECTURE 1

A schematic of the first optical interconnect architecture is shown in Fig. 2. A message from an originator processor illuminates a one-dimensional SLM, which spatially encodes the address of the destination processor on the optical wavefront. The wavefront illuminates a multiplexed hologram, which reconstructs a reference beam at an angle θ_n, associated with the address of the nth destination processor. For connecting N processors, an NxN SLM is used with each originator processor illuminating one column. Each column may encode one of N orthogonal destination addresses. A volume holographic recording media is used because of it's high angular selectivity, and hence potential for interconnecting a large number of processors.[17]

The holographic self-router requires first programming the $LiNbO_3$ crystal with N holograms. Each exposure consists of associating a particular reference beam angle with a particular processor address. This N bit address appears on all columns of the SLM. After each exposure, the reference beam is moved such that it illuminates the crystal at a new angle. Another orthogonal pattern is programmed on the SLM, and the next holographic exposure is formed between the new reference beam and the new destination address. N holograms are recorded in this manner. Once the volume media is fully programmed, connections between processors are formed dynamically by displaying different destination addresses on columns of the SLM, which are then holographically steered to the appropriate destinations. The response time of holographic routing is no longer limited by the response time of the photorefractive crystal: It is limited by the response time of the SLM. Using parallel addressed III-V (InP/InGaAs) multiple quantum well SLMs, GHz reconfiguration rates are possible.[18]

The experimental arrangement is shown in Fig. 3. Light from an Argon-ion laser (λ = 514.5nm) is horizontally polarized and split into two beams by a beamsplitter. The object beam is spatially filtered (NRC model 900), and collimated by lens L1, (f = 200mm). The horizontal polarization is then rotated to the vertical by a $\lambda/2$ plate, and illuminates a 6x6 FLC SLM (Displaytech, Inc. model 6X6P). Since only a 4x4 SLM is required to interconnect 4 processors, all but 16 pixels were masked off in the 6x6 SLM. A Galilean telescope with a demagnification factor of six was used to project the SLM patterns into the volume of the $LiNbO_3$ photorefractive crystal.

The pattern for a particular destination address is programmed on each column of the SLM as opaque and transmitting pixels, representing 0's and 1's, respectively. Each destination address interferes with a single reference beam in the $LiNbO_3$ crystal forming a phase hologram. After each exposure, a new destination address appears on each of the 4 SLM columns, and interferes with a new reference beam at a new angle θ_n in the crystal. The recorded holograms are multiplexed at angles much greater than the angular selectivity (.1mrad) of the 1 cm thick $LiNbO_3$ crystal.[19] The reference-to-object beam ratio K, for these experiments was approximately 2:1, with the reference intensity $I_R = 3.4 mw/mm^2$.

Four orthogonal destination addresses were recorded in the $LiNbO_3$ crystal using this procedure. To test architecture 1, the SLM was programmed to sequentially display the four addresses on all of it's columns. Detected power in each of the four locations is shown in Figure 4. Each destination address programmed on the SLM reconstructs a unique reference beam, laterally displaced such that it can be individually detected. The diffraction efficiencies are optimized using a previously published procedure[20] for storing multiply-exposed holograms in photorefractive $LiNbO_3$. The architecture described here succeeded in connecting 4 processors to 4 processors, but requires a 4x4 SLM to accomplish this task. This is because SLM addresses with bits in common significantly reduce the signal-to-noise ratio in the detection plane. In the next section, we present a second optical interconnect architecture, which uses non-orthogonal destination addresses and therefore reduces the number of SLM pixels required for routing.

2.2 ARCHITECTURE 2

The feature that allows the use of non-orthogonal patterns in the second architecture is the distribution of system crosstalk. Architecture 1 required orthogonal addresses, since crosstalk from different input channels can add up at a single detector. System 2 permits the use of non-orthogonal addresses by insuring that crosstalk from different input channels will necessarily go to different detectors.

In this implementation each processor controls a horizontal one-dimensional SLM with M pixels for N processors, a two-dimensional NxM SLM is used. The routing holograms are recorded in a similar manner as described in Architecture 1. The same destination address is displayed on each row of SLM, and this address is spatially encoded on a collimated laser beam. The $LiNbO_3$ crystal lies in the Fresnel region, as illustrated in Fig. 5. The encoded object beam and a collimated reference beam form a hologram, again matching a particular destination address to a unique horizontal angle. A multiple-exposure hologram is recorded, mapping N addresses to N horizontal angles.

During reconstruction, light is diffracted into a particular horizontal angle when the corresponding pattern is displayed on any row of the SLM. However, the vertical position of the diffracted light is determined by the row on which the pattern is displayed. Since there are N rows and N addresses, an N^2 detector array is required. Because different input channels cannot diffract light to the same detector, overlap in the destination address bits is permissible.

To test this system, a second 4x4 interconnection network was formed. Non-orthogonal intensity modulated addresses were used. Each of the addresses had two bits *ON* (1's) and two bits *OFF* (0's). To increase the vertical separation of reconstructed beams, a cylindrical lens with power in the vertical direction was used.

After recording the four holograms, the SLM was programmed to sequentially display eight of the sixteen possible input/output combinations. Detected power at one detector location is shown in Fig. 6a. Fig. 6b shows the result of using the same non-orthogonal patterns in architecture 1. A comparison of Figures 6a and 6b illustrates the reduced crosstalk in the optical interconnection architecture 2.

Since these are parallel networks, the signal power must be compared to the sum of the crosstalk power. This experiment shows that it is possible to reduce the SLM size, by increasing the dimensionality in the detector plane.

3. DISCUSSION

Both of the architectures presented require recording angularly multiplexed holograms. If angularly multiplexed holograms are recorded with equal exposure energies in a photorefractive crystal, they reconstruct with unequal diffraction efficiencies.[20,21] Holograms recorded later in a multiple exposure reconstruct more brightly than earlier exposures. A compensation technique has been presented which allows the recording of 20 equal diffraction efficiency holograms.[20] For detectable signals, it is necessary for the lowest diffraction efficiency to be brighter than the crosstalk. Fluctuations in diffraction efficiency pose more of a problem in the second architecture than in the first, since architecture 1 uses orthogonal addresses, which automatically minimize crosstalk.

The second architecture has inherent crosstalk due to the overlap between addresses. If crosstalk from the highest diffraction efficiency hologram is more powerful than signal from the lowest diffraction efficiency hologram, it becomes impossible to distinguish signal from noise. The ability to record a large number of equal diffraction efficiency holograms is necessary for the development of a large system. This is a topic of current research.

Lithium Niobate was used in these experiments because of it's angular selectivity. However, as with any dynamic medium, erasure will take place during readout. A ND3 filter was used to minimize this erasure. For a practical implementation of this architecture, a static multiple-exposure hologram with high angular selectivity is necessary.

Another limitation on the size of these networks is the SLM size. Clearly, to implement a N processor network using architecture 1, an NxN SLM is required. The SLM size for architecture 2 is less straightforward, since it depends on the acceptable Hamming distance, which is determined by fluctuations in diffraction efficiencies of the reconstructed holograms.

All of the addresses used in architecture 2 should have the same number of pixels in the *ON* state. For crosstalk less than one-half, addresses with two bits *ON* were chosen. There are $\frac{(M)\ (M-1)}{2}$ such addresses of length M. It can be shown that this results in reducing the address length to approximately $\sqrt{2N}$.[22]

4. CONCLUSION

In this paper we present the design and experimental results for connecting 4 processors to 4 processors using two dynamic holographic routing architectures. Neither architecture 1 nor architecture 2 are limited by the speed of the photorefractive recording materials. The dynamic reconfiguration is currently controlled by a ferroelectric liquid crystal SLM, at a rate of 10kHz. The advantage of architecture 2 over 1 is that it reduces the number of SLM switches required to route N messages from N^2 to approximately $N\sqrt{2N}$. future work includes using III-V SLMs to encode the destination addresses, and analyzing the limits on multiple exposure storage in photorefractive media.

5. ACKNOWLEDGEMENT

We acknowledge the support of ARO contract DAAL 03-880k-0195, NSF/ERC CDR 8622236, Presidential Young Investigator Eng8451485 and useful discussions with Kelvin Wagner, Karl Gustafson, Mark Freeman, and C.C. Mao.

6. REFERENCES

[1] D.Z. Anderson and D.M. Lininger, "Dynamic optical interconnects: volume holograms as two-port operators," Appl. Opt. 26, 5031 (1987).

[2] J. Wilde, R. McRuer, L. Hesselink, and J.W. Goodman, "Dynamic holographic interconnections using photorefractive crystals," Proc. Soc. Phot. Inst. Eng., 752, 200 (1987).

[3] G. Pauliat and G. Roosen, "Large scale interconnection networks using dynamic gratings," Phot. Inst. Eng. 700, (1986).

[4] R. McRuer, J. Wilde, L. Hesselink, and J.W. Goodman, "Two-wavelength photorefractive dynamic optical interconnect," SPIE 881, 192 (1988).

[5] P. Yeh , A. Chiou, and J. Hong, "Optical interconnections using photorefractive dynamic holograms," Appl. Opt 27, 2093 (1988).

[6] J.W. Goodman, F.I. Leonberger, S.Y. Kung, and R.A. Athale, "Optical interconnections for VLSI systems," Proc. IEEE 72, 850 (1984).

[7] D.A.B. Miller, "Optics for low-energy communication inside digital processors: quantum detectors, sources, and modulators as impedance convertors," Opt. Lett. 14, 146 (1989).

[8] A.R. Dias, R.F. Kalman, J.W. Goodman, and A.A. Sawchuck, "Fiber-optic crossbar switch with broadcast capability" Opt. Eng. 27, 955 (1988).

[9] E. Marom and N. Konforti, "Dynamic optical interconnections", Opt. Lett. 12, 539 (1987).

[10] M. Kranzdorf, J. Bigner, L. Zhang, and K.M. Johnson, "An optical connectionist machine with bipolar weights", accepted Opt. Eng. (1989).

[11] D. Psaltis, D. Brady, K. Wagner, "Multilayer optical learning networks,"Appl. Opt. 27, 1752 (1988).

[12] N.A. Clark and S. Lagerwall "Submicrosecond bistable electro-optic switching in liquid crystals", Appl. Phys. Lett. 36, 899 (1980).

[13] N.A. Clark, M.A. Handschy, and S.T. Lagerwall, "Ferroelectric liquid crystal electro-optics using the surface stabilized structure", Mol. Liq. Cryst. 94, 213 (1983).

[14] K.M. Johnson, M.A. Handschy, and L.A. Pagano-Stauffer, "Optical computing and image processing with ferroelectric liquid crystals," Opt. Eng. 26, 385, (1987).

[15] FLC/SLM available from, Displaytech, Inc., 2200 Central Avenue, Boulder, CO 80301, 303-449-8933.

[16] See for example, T.J. Hall, R. Jaura, L.M. Connors, and P.D. Foote, "The photorefractive effect - a review," Progr. Quant. Electr. 10, 17 (1985).

[17] D. Staebler, W.J. Burke, W. Phillips, and J.J. Amodei, "Multiple storage and erasure of fixed holograms in Fe-doped $LiNbO_3$," Appl. Phys. Lett. 26, 182 (1975).

[18] L. Messick, D.A. Collins, and D.L. Lile, "A 64-bit 800 MHz insulated gate CCD on InP," IEEE Elec. Device Lett. EDL-7, 680 (1986).

[19] H. Kogelnik, "Coupled wave theory for thick hologram gratings," Bell Sys. Tech. J., 48, 2909 (1989).

[20] A.C. Strasser, E.S. Maniloff, K.M. Johnson, and S.D.D. Goggin, "Procedure for recording multiple-exposure holograms with equal diffraction efficiency in photorefractive media," Opt. Lett. 14, 6 (1989).

[21] W.J. Burke and P. Sheng, "Crosstalk noise from multiple thick phase holograms,"J. Appl. Phys. 48, 681 (1977).

[22] E.S. Maniloff and K.M. Johnson, to be published.

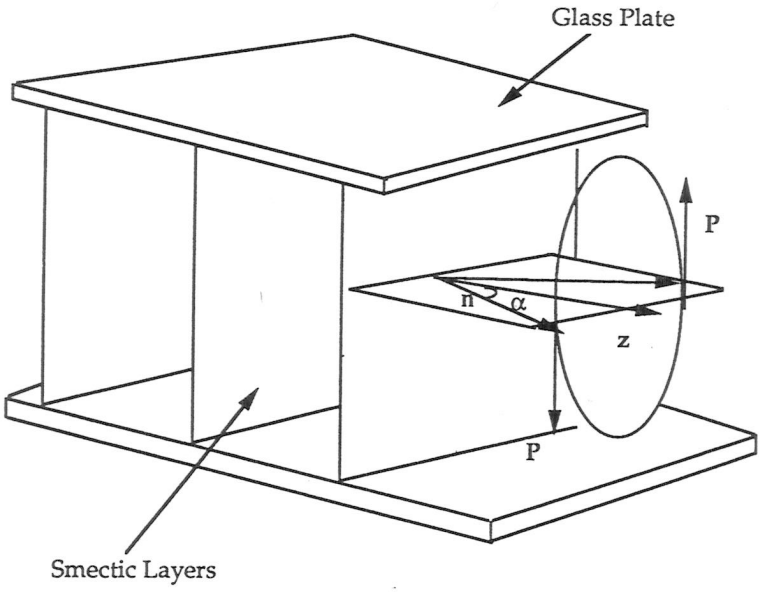

Figure 1 Structure of surface stabilized ferroelectric liquid crystal (SSFLC). n refers to the molecular direction, α the tilt angle, and P the macroscopic polarization.

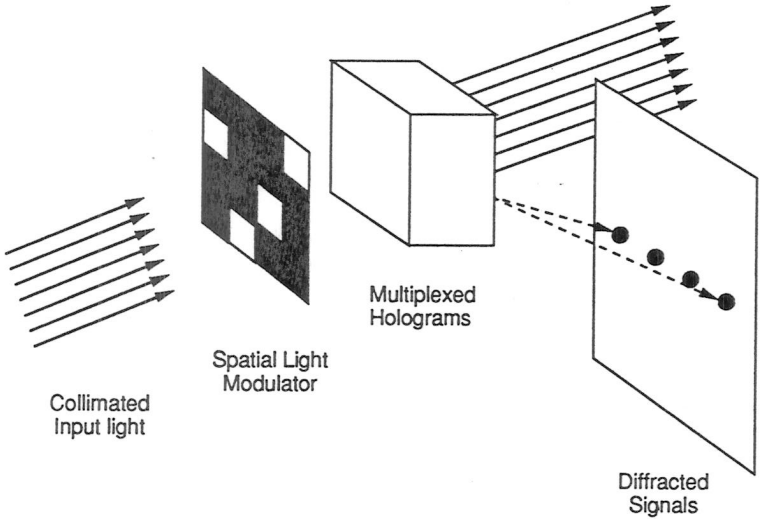

Figure 2 Schematic of architecture 1, addresses are encoded on SLM columns.

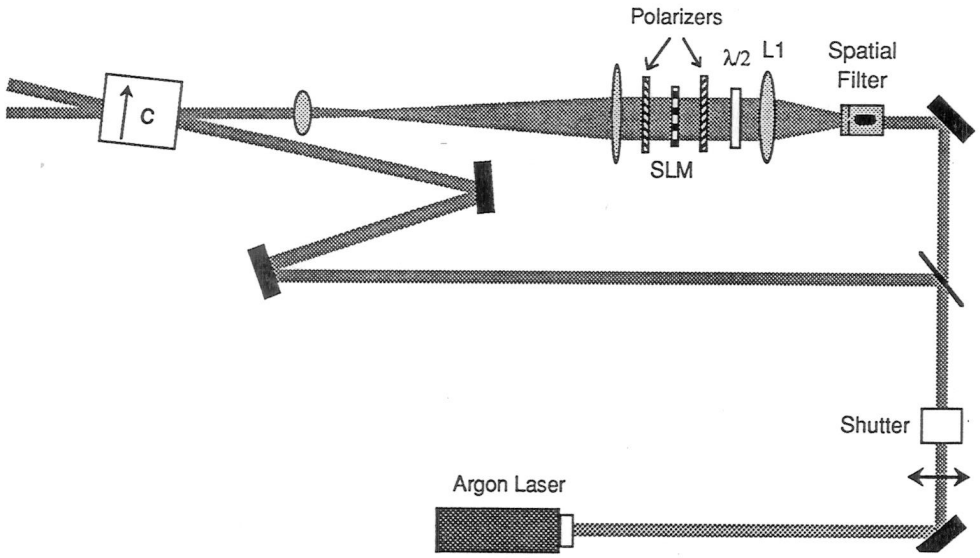

Figure 3 Experimental arrangement for recording the holographic routing network.

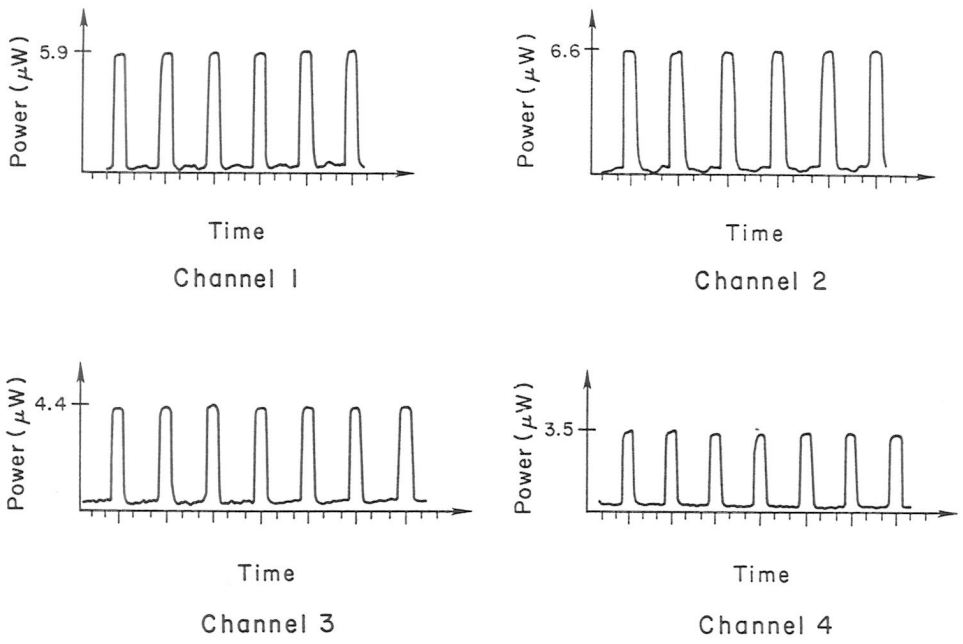

Figure 4 Detected power at each of four locations of a 4x4 interconnect using orthogonal addresses

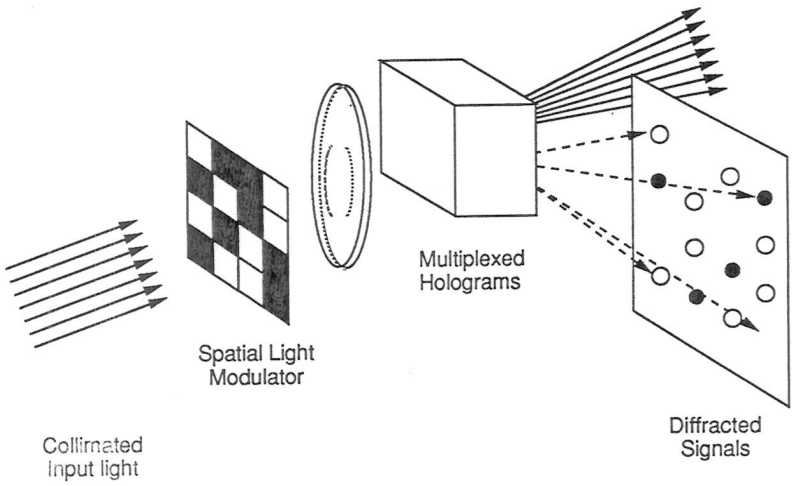

Figure 5 Schematic of architecture 2, addresses are encoded on SLM rows.

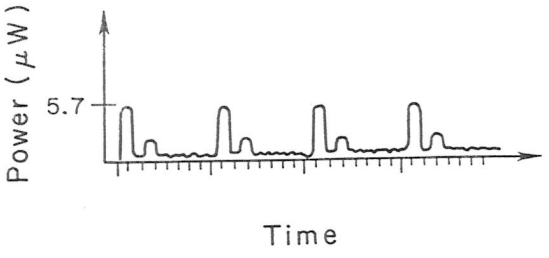

Figure 6 (a) Detected power at one location for a 4x4 interconnect using nonorthogonal addresses and crosstalk distribution.

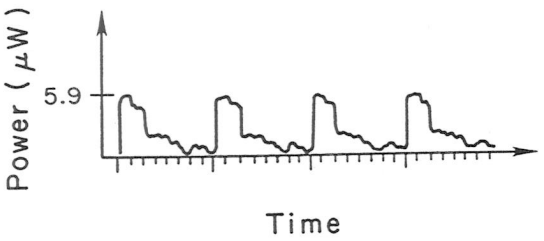

(b) Detected power at one location using the same addresses without crosstalk distribution.