OPTICAL EXPANDERS GIVE CONSTANT TIME HOLOGRAPHIC MESSAGE ROUTING USING O(N log(N)) SWITCHES

John Reif[1]
Comp. Science Dept
Duke University
Durham, NC 27706

0. SUMMARY

We describe an electro-optical message routing system for sending N messages between N processors in constant time using 2Nlogn switches. A spatial light modulator (SLM) is used to holographically steer messages directly to their destination processor. The system is unique in that it uses fixed holograms to achieve free space dynamic routing. A small prototype implementation has been already constructed [Maniloff, Johnson and Reif,89]. (An appendix describes practical issues.)

We introduce a new optical technique which we call the **optical expander**. We discuss how an optical expander can be used to solve a key problem, namely the orthogonality of message patterns. In particular, the optical expander system is used to decrease the number of address bits used by the router and to improve separation of distinct address patterns matched by the holograms. We discuss the theory of the optical expander system and give for the first time a rigorous **proof of its correctness and performance**.

---

[1] Supported in part by DARPA/ARO contract DAAL03-88-K-0195, Air Force Contract AFOSR-87-0386, DARPA/ISTO contract N00014-88-K-0458, NASA subcontract 550-63 of primecontract NAS5-30428.

1. INTRODUCTION

1.1 The Potential of Optical-Electronic Systems

The inherent high parallelism and connectivity of optical signal processing lends itself directly to such applications as optical interconnection. (See the recent text of [Feitelson,88]). The recent development of moderately high speed, high dynamic range spatial light modulators has lead to the prototype development of variety of optically based signal processing systems.

1.1 Our Holographic Routing System

Dynamic message switching is the problem of sending a f N messages between N processors, where the destination permutation is given dynamically. In this paper we describe a novel *holographic message routing system* for dynamic message switching. We use a spatial light modulator (SLM) to holographically steer messages directly in free space to their destination processor. An important innovation of our holographic routing system is the use of fixed holographs to do the dynamic message switching. It uses 2N logn boolean switches, which is optimal within a factor of 2. It has a constant time bound to do the routing and uses volume $O(N^{3/2})$. These time and volume bounds are asymptotically optimal with respect to the VLSIO model (this is a theoretical model for optical-electronic computing developed in [Barakat and Reif, 1987])

In brief, our holographic message routing system is a unique architecture which uses N multiple-exposure holograms, each containing N images to connect N processors to N processors, via free space routing. The system uses N spatial light modulators (SLMs), each with 2logn pixels. A column of light illuminates each processor's SLM which is programmed with an encoded address for a destination processor. This optically encoded address is routed directly to the correct processor by a hologram containing N images, each correlated with a particular destination processor. This optical interconnection network is a direct message router taking constant time as compared to conventional fixed interconnection networks which require time delay at least logn. Our holographic message system can be applied to do very high speed message routing for massively parallel machines such as the CONNECTION machine.

While this paper is primarily concerned with theoretical proof and analysis of our optical routing method, we also discuss practical issues(see the Appendix of this paper for more practical details). There is a collaborative Optical Routing Project (funded by DARPA) between theoretical computer scientist, John Reif, at the Computer Science Department, Duke University and optical engineers Kristina Johnson and Eric Maniloff at the Center for Optoelectronic Computing Systems at University of Colorado, Boulder. While Reif initially conceived of the theory of the system, the practical implementation was due to Johnson and Maniloff. Johnson and Maniloff built a 4 by 4 prototype holographic routing system (for implementation details see [Maniloff, Johnson and Reif,89]) at the Center for Optoelectronic Computing Systems at University of Colorado, Boulder. This running prototype implementation was completed in April 1989. Because of the small size of this prototype system, an optical expander system was not required. They have also developed in [Strasser, Maniloff, Johnson, Goggin,89] a , procedure for recording multiple-exposure holograms with equal diffraction efficiency in photorefractive media. Reif has also directed computer simulations of the message routing applications.(The availability of a device which can control light with a high spatial resolution and with a short cycle time is critical to the successful realization of a second generation our system; for this we acknowledge the technical assistance from Derek Lile, Colorado State University, on the development of III-V MQW/CCD SLMs.).

1.2 Outline of paper

We begin in Section 2 with an identification and discussion of the significance of the message communication problem. Next in section 3 we describe our optical routing system in detail. In section 4 we ddescribe and prove optical optical expander system. Finally, in section 5 we prove our optical routing system using the now proved optical expander system. In the Appendix we discuss practical issues

2 IDENTIFICATION AND SIGNIFICANCE OF THE PROBLEM:

2.1 The Message Communication Problem

Message routing is the following task: we assume there are N processors, where each processor has a distinct message with a given distinct message destination address. Simultaneously, each message is to be routed from its originating address to its destination address. The message routing problem is crucial to parallel machines, since this is how information is moved between distant processors. Furthermore, message routing allows a parallel machine with fixed connections to simulate arbitrary interconnects. In the following, we are concerned with the case when the length of each message is of some fixed constant length L

2.2 Electronic Implementations of Message Routing

A considerable amount of theoretical work([Valiant and Brebner,81],[Aleliunes,82],[Upfal,82]) has been done to devise efficient algorithms for parallel message routing on fixed connection networks such as the HYPERCUBE, and BUTTERFLY. These routing algorithms achieve bounds that are asymptotically optimal for routing N messages on fixed connection networks with N processors; in particular they require (in the worst case with high likelihood) time O(log(N)) and require the setting of O(N log(N)) boolean switches.

A number of parallel machines have electronic hardware implementations of fixed connection network message routing.

For example, the BUTTERFLY is a 128 processor MIMD machine constructed by BBN. The BUTTERFLY has a software controlled message routing system requiring many times the basic step time of the machine. The BUTTERFLY machine has 128 processors under MIMD control, where each processor can execute 32 bit wide instructions at a rate of approximately 1 megahertz. Due to contention problems, the BUTTERFLY requires approximately .1 milliseconds for its software to implement a general routing of 128 messages to distinct processors. (The routing time can be much faster only in special cases with very little contention). Thus the time for general message routing is approximately 100 times slower than the basic instruction rate of the BUTTERFLY machine.

The CONNECTION machine is a 65,536 processor bit serial SIMD parallel machine constructed by Thinking Machines, Inc. The CONNECTION machine is only one of a number of highly parallel computers which have been recently built (others include the 16,384 processor MPP machine of NASA Goddard Space Flight Center manufactured by Goodyear, Inc. See [Lerner,85] for a survey of these and other recent parallel machine architectures.). However, it is unique in that it's hardware was the first with a built-in ability to do message routing. This message routing system can simulate arbitrary permutation interconnections.

The message routing on the CONNECTION machine requires 65,536 messages to be routed to distinct addresses. The CONNECTION machine executes instructions at each bit serial processor at a rate of approx 5 megahertz. The hardware message routing system of the CONNECTION machine is by its use of electrical wires and integrated electronics. General message routing on the CONNECTION machine requires approximately 500 microseconds . to complete a general routing of 65,536 messages. (Of course special, restricted message routing problems such as routing only to nearby processors can be solved on the CONNECTION machine much faster than this.) Therefore, general message routing is approximately 2,500 times slower on the CONNECTION machine than the basic instruction step time. While we expect that later models of these machines will have somewhat improved times for message routing, we can *not* expect improvements by orders of magnitude.

The goal of our message routing system is to do the message routing required by parallel machines described above, *within the instruction rate* of the machine, keeping up with the rate at which message routing requests can be issued by the processors of the parallel machine. This requires use to implement general message routing *many times faster* than previous electronic interconnects

The key practical problem is to make the routing time approach the instruction step time of these machines. For the BUTTERFLY our goal is to decrease the message routing set-up time to the instruction rate. Note that 32 is the typical length of messages on the CONNECTION machine sent by the bitserial processors. Our goal for the CONNECTION machine is to decrease the message routing time by (at very least) a factor of 50 to 32 instruction steps per message routing; that is to at most 16 microseconds. See section 7 for a discussion of the use of pipelining to help achieve these routing times.

2.3 WHY KNOWN OPTICAL CROSSBARS AND MULTIPLE STAGE NETWORKS DO NOT SOLVE THE PROBLEM

One possible solution to the message routing problem with N messages and N processors is the use of an *N x N optical crossbar.* The input to the optical cross-bar is a vector of N optically encoded messages, one from each processor. These inputs are spread out to columns of an N x N optical switching array (SLM). The device then performs a vector-matrix multiplication of the vector of input messages times this array. If the switch array is set to a permutation matrix defining the required permutation routing destinations, the messages will be permuted corresponding to the destination addresses. (For more details see [Goodman and Johnson, 1981, 1981], [Neff,1985], [McAulay,85], and the survey paper of [Sawchuk, Jenkins, Raghavendra, Varma,85]). The advantage of this optical crossbar is that once the array switches are set, the message data can be transmitted at optical transmission rates, which can be between .2 and 20 gigahertz per message steam, or at total of between .2N and 20n gigahertz for the N messages transmitted in parallel.

The unfortunate drawback of this optical cross-bar scheme is that it requires a very long time to set-up the switches for message routing. In particular, the key problem is that there are N^2 switches in the switch array which must be set. For example in the case of N = 256 processors and N = 256 messages, the are a total of N^2 = 65,536 array switches which must be set. In 32 x 32 optical crossbar systems which have been built to date, setting all these switches has required at least 320 milliseconds. To significantly decrease the set-up time, a method must be devised that avoids the use of an array of N^2 switches.

Another possibility is the use of multiple stage optical switching networks, which are optical versions of the fixed connection electrical routing schemes mentioned above. These have the advantage of requiring only the setting of O(N log(N)) switches, but the disadvantage of requiring clog(N) time, where c is a moderate constant somewhat above 2.

Another possibility is to use *Benes networks* which require only 2log(N) time. Unfortunately, determining the setting of the switches of a Benes routing network requires rather lengthy computations. Therefore, an optical Benes network is only practical in the case of a fixed set of message routing permutations, which is not generally the case for applications on parallel machines such as the CONNECTION machine.

3. OUR HOLOGRAPHIC MESSAGE SWITCHING SYSTEM

The *holographic message routing system* which we propose will require only setting 2Nlog(N) switches to route N messages to N processors. Moreover, for moderate N is uses only constant time. For example, in the case N = 256 we require 2Nlog(N) = 4,096 switches, rather than then the array of N^2 = 65,536 switches required by optical crossbars described previously. Furthermore our switch settings can be determined immediately from the message destinations. Nevertheless, our basic proposed method requires only one stage of routing, and requires no active elements except for the 2N log(N) switches.

3.1 Holographic Matching

The basic idea in our holographic message router is to use holographic associative matching to direct each of the messages to their destination processors. We now give a brief introduction of hologram associative matching.
We recommend the reader review a text on holographic techniques such as [Iizuka,87], [Collier,Burchhardt,Lin,71], or [Hariharan,84]) . Also see section 6 for a discussion of multiple exposure techniques.

3.2 Details of our Router

We now describe our holographic message router in detail.

We shall assume that there are N processors. These N processors are arranged in a $n^{1/2} \times n^{1/2}$ two dimensional array. The holographic message routing system is a unique architecture which uses N multiple-exposure holograms, each containing N images to connect N processors to N processors, via N SLMs each with $2\log_{2N}$ pixels.

A simplified 1D diagram of the architecture is illustrated in Fig. 2 for connecting 4 originating processors to 4 destination processors.

We will now describe the electro-optical components required by each originating processor. A column of light illuminates each processor's SLM which is programmed with an encoded address for a destination processor. This optically encoded address is routed to the correct processor by a hologram containing N images, each correlated with a particular destination processor. This optical interconnection network is a direct message router as compared to conventional fixed interconnection networks which require time delay at least logn.

We will write the *name* of each processor as a fixed 2log(N) bit binary number.

Each processor will have a *message* and that processor will be known as the *originating processor* for the message. The originating processor will decide a *destination* for the message. Each originating processor sets a linear array of 2log(N) optical switches to its message's destination name. This linear array of optical switches will be called the *destination switch array*. Note also that each destination switch array is very short; in fact in the case where N = 256, the array is of length only 2log(N) = 16. Also note that since there are N processors, a total of 2Nlog(N) switches must be set in the entire system to implement a message routing.

The (electrical) message will be encoded onto a single coherent optical beam in Phase or Pulse Position Modulation format. This conversion from electronic signal to optical can be done by conventional use of a GaAs Laser. The resulting single optical beam can transmit the message bits at a rate of between 1 and 10 gigahertz. This coherent beam will be called the *message beam*.

This message beam is transformed (using say suitable anamorphic optics) from a point source of light to a column of light. This column, called the message column, illuminates the 1-D SLM, whose 2log(N) switches have been set to the message's destination. The result is that each originating processor transmits through it's destination switch array a column of at most 2log(N) identical beams, each transmitting the identical message. The destination address of the message is now called the *destination pattern*. The destination pattern then enters an optical expander (see explanation of optical expanders in Section 4 and see [Barakat and Reif, 89] for a more detailed discussion), which produces as output an *expanded destination pattern* of length 2N. By the definition of the optical expander system, each destination pattern is orthogonal to any other expanded destination pattern.

We use holographic associative matching to direct each message to its destination processor, as specified by its expanded destination pattern. In particular, we illuminate a volume or planar hologram (in practice, we use a volume photorefractive crystal of Fe-doped $LiNbO_3$) with the destination pattern, which will be called the *routing hologram*. The routing hologram is simply a fixed multiple-exposure hologram. Each exposure is made by interfering a *reference* beam with an *objective* pattern, where the reference beam consists of a single coherent beam eminating from the destination processor, and the objective pattern consists of the e expanded destination pattern for the destination processor. When this multiple-exposure hologram is fixed (it is fixed for all time in the initial setup of the apparatus), the result is what we call the routing hologram. By the well known properties of multiple-exposure holograms, it has the property that the destination pattern matches (correlates) strongly with one of the exposed objective patterns, and the result is that the optical message is reflected or transmitted by the routing hologram to the destination processor. (Note that false matches are minimized by using an optical expander to insure the expanded destination patterns are orthogonal. This insure the resulting incorrectly matched patterns have extremely low signal to noise ratio). Hence the message is routed by this hologram directly to the processor which is the destination of the message. Note that aside from the initial setting of the optical switches, this holographic routing runs at the speed of light and requires no other active elements. Thus as long as the number N of processors is moderate (see discussion below), the routing time is constant .

Figure 1 gives an illustration of the setup required for each originating processor. That is, each processor requires its own routing hologram. Note that to configure the entire system, the processors will be arranged in a two-dimensional array, facing a two-dimensional array of their routing holograms. As we shall see, this is practical for up to at least 512 processors.

[Barakat and Reif,87] propose a theoretical model for optical-electronic computing, which includes both the standard VLSI model as described in [Ullman,84], as well as reasonable assumptions for the volume and time required by optical components. We assume that VLSIO model in this paper. The entire holographic routing system described here can be placed in a cube with sides of length $n^{1/2}$, and using a total volume of $O(N^{3/2})$., which by the lower bound results of [Barakat and Reif,87], is optimal for permutation routing in the VLSIO model of optical computing .[Tyagi and Reif,89] also prove energy lower bounds for VLSIO and [Reif and Tyagi,90] give some optimal algorithms for the VLSIO model with the DFT primitive.

3.3 Limitations to Our System

Ultimately, our routing system is limited by various limitations due to physical devices; for example on the number of multiple exposure holograms (see Appendix A.1). and the switching speed of SLMs(see Appendix A2). (Also see Appendix A.5 for a discussion of the use of pipelining to help overcome these limitations.)
However, the key limitation, which we consider in the following section, is the orthogonality of the address patterns.

4 Optical Expanders

In this section we describe a method for creating a set of N orthogonal vectors each of lengthN by use of an electro-optical device with at most 2logN boolean inputs. We prove the correctness our method.

4.1 IDENTIFICATION OF THE PROBLEM

Vectors X, Y are defined to be **orthogonal** if $(X,Y) = 0$, that is the inner product of X and Y is 0.

An **optical expander** is a electro-optical device that takes as input a boolean vector of length d. We distinguish N input boolean vectors $p_1,\ldots,p_N$ and call these the **input patterns.** Each input pattern will be optically encoded as a spatial array of d of pixels (generated by d coherent 1-dimensional beams) of intensity either 0 or 1.. Given input pattern p_i, the optical expander outputs a spatial **output pattern** $r_i = \mathrm{EXPAND}(p_i)$, which consists of an array of N pixels of intensity either 0 or 1. Each output pattern r_i is viewed as a boolean vector of length N. We require that these output patterns $r_1,\ldots,r_N$ be **orthogonal**, i.e. every distinct pair r_i,r_j has inner product $(r_i,r_j) = 0$. The optical expander is **exponential** if N is an exponential function of d.

4.2 Our results

We first show that no linear optical system can be an optical expander, and thus we must use a nonlinear filter. Then we describe a method for constructing (nonlinear) optical expanders with $d \leq 2\log N$.

Note. We have also investigated [Barakat and Reif,89] another optical expander method using a simpler class of linear filters(i.e., interference patterns), followed by (nonlinear) thresholding, but have not succeeded in rigorously proving this alternative method. (although computer experiments [Barakat and Reif,89] have empirically verified the performance of this alternative optical expander method.)

4.2 Linear Filters do not Suffice

Note that by linear algebra, at most d vectors of length d are linear independent. Our key problem is getting these patterns transformed by a filter EXPAND into N output patterns that are orthogonal. If we generate these output vectors by linear function application, then at most d of the resulting output vectors can be linear independent. Thus at most d of the output vectors can be orthogonal.
Hence no linear filter can act as an optical expander with $d < N$.

4.3 OUR OPTICAL EXPANDER

The main idea is to put the input image through a N x d matrix multiplier followed by a nonlinear filter which thresholds the intensity over the each of N subpatches of the resulting image. The effect of the thresholding is that the resulting output images are orthogonal.

The Input Pattern Code Words

Let N be a power of 2. (Throughout this paper log N is be taken with base 2). To construct the input patterns, we proceed as follows. For any even $k \leq \log N$, let S_k be the set of boolean vectors (which we will call **candidate patterns**) of length k where exactly k/2 of the values in any vector are 1. We can find each such a S_k by enumeration; we simply compute a list of all boolean vectors of length k that have exactly k/2 ones(there are at most N^2 such vectors). Let d be the minimum even k such that S_k has size at least N. Note that if $k = 2\log N$ then S_k contains as a size N subset the set $\{x.x' \mid xy$ is a boolean string of length log N and y is the binary complement of x $\}$. Thus we can set $d \leq 2\log N$. (In fact, we can do somewhat better than this. By a counting argument, we estimate the probability that a random boolean vector of length d has exactly d/2 ones. Then we can show with probability limiting to 1, S_k has N vectors for $k = b \log N$ for a constant b below 2.). For the rest of the paper, we assume for simplicity that we have constructed the set S_d. with $d = 2\log N$.

Our input patterns $p_1, \ldots, p_N$ are defined to be the smallest d N elements of S_d. Each pattern p_i is viewed as a pixel array of even length d, each pixel being either ON(intensity 1) or OFF(intensity 0). We have insured that each p_i has exactly d/2 pixels ON, so the total intensity of each input image summed over all the pixels is the same, namely d.

4.4 A Mathematical Solution to the Optical Expander Problem

We first give a precise mathematical description of the optical expander, and then describe the optics required.

We define a function EXPAND which maps from boolean vectors of length d to boolean vectors of length N, where given a boolean vector $X = (x_1, \ldots, x_d)$, then EXPAND(X) is a boolean vector $(y_1, \ldots, y_N)$. For $j = 1, \ldots, N$ we define $y_j = \text{threshold}_t(f_j(X))$ where f_j is the linear function $f_j(x) =$ the inner product (X, p_j), and threshold_t is a threshold function (i.e., threshold_t is 1 if its input at least threshold value t, and otherwise 0). We will set the threshold parameter $t = d/2$ so that if $X = p_j$ then the resulting vector $f_j(p_j) = (p_j, p_j) = d/2$ and so $y_j = \text{threshold}_t(f_j(X)) = 1$, and otherwise $f_k(p_j) < d/2$ and so $y_k = \text{threshold}_t(f_k(X)) = 0$

for k not equal to j. Thus EXPAND(**X**) is a boolean vector with a unique 1 value at position j (with 0 in the rest of the positions). This implies that each of N of the output patterns $\mathbf{r_i}$ = EXPAND($\mathbf{p_i}$) is orthogonal for i = 1,...,N.

4.5 An Optical Implementation of our Solution to the Expander Problem

Next we describe the optics required by our optical expander.

The filter we propose has 2 stages:

STAGE 1:

Note that the inner products can be computed by a N x d matrix multiplication times the input pattern $\mathbf{p_i}$. The required matrix **P** has $\mathbf{p_j}$ in its jth row, for j = 1,...,N. The matrix multiplier systems described and implemented in [Goodman and Johnson, 1981, 1981], [Neff,1985], [McAulay,85], [Sawchuk, Jenkins, Raghavendra, Varma,85] can be used here. Since the matrix **P** is fixed, we do not suffer with respect to set up time for **P.** The output of the matrix multiplier is the vector $\mathbf{q} = \mathbf{Pp_i}$ of length N.

STAGE 2:

Next use a nonlinear threshold device that takes as input this large vector $\mathbf{q} = \mathbf{Pp_i}$ and computes a threshold on each of it's individual pixels. In particular, we choose a fixed constant parameter t, which is the threshold parameter .

There are a number of electro-optical devices that operate at high rates and can be used for the required threshold array. A straight forward approach is a transformation back to electronic by detectors, and then threshold and back to optical signal. Note that the threshold device need only have resolution of N bits over it's total area.

We define t = d/2 as above noting that only one of the pixels of **q** can be intensity t. Then we set $y_i = \text{THRESHOLD}_t(q_i)$ by replacing each pixel of intensity $\geq$ t with a pixel of intensity 1 and replacing each pixel with intensity < t with a pixel of intensity 0. The result is a pixel pattern that encodes EXPAND(X), the boolean vector $(\mathbf{y_1},...,\mathbf{y_N})$ as defined above.

5 Applications of Optical Expanders

5.1 Proof of our Holographic Routing System using Optical Expanders

Assuming we have an optical expander system as described above, we can now give a detailed construction and proof of our holographic routing system using O(Nlog(N)) boolean switches with constant time.

We use the N orthogonal output patterns $r_1,...r_N$ to address holograms. Suppose we have exposed a hologram for each pattern r_i The error intensity resulting from entering a pattern r_i into an incorrect hologram (i.e. a hologram set to another distinct orthogonal pattern r_j) is at most c/N^2 for some constant c in 3 dimensional optical systems.

By the assumption that the patterns r_i are orthogonal, we get that the error intensity from any incorrect match is c/N^2 We will input the filtered images r_i= EXPAND(p_i) to the addressing holograms.
We can assume the intensity of each filtered image r_i is 1.

Now we show we can achieve given **signal to noise** ratio N:c. Recall we have set the length of the output vectors to be N . Each incorrect match gives intensity $\leq c/N^2$. The total of N-1 possible simultaneous incorrect matches (from error intensities) is at most $(N-1)c/N^2 \leq c/N$, which is the total error intensity. In contrast, the correct match intensity is 1. Thus we detect the correct message with a signal to noise ratio $1/(c/N) = N/c$. Hence the correct signal would be detectable as long as c (which optical theory implies should be a constant) remains below say N/4

5.2 Electrical addressing of Holograms: Holographic Storage and Retrieval

The above argument implies that we use input patterns with at most 2log N bits to address any of N holograms in a multiple image hologram. In particular, we us our optical expander to expand the $d = 2\log N$ bits to N bits. The optical expander system outputs N pattern vectors $r_1,...r_N$ which are orthogonal. This restriction allows us to guarantee the storage holograms can distinguish N distinct pattern vectors. The above analysis in 5.1 also implies that the signal to noise ratio is N:c.

5. REFERENCES

R.,Aleliunas Randomized parallel communications. In Proceedings of the ACM Symposium on Principles of Distributed Computing (Ottawa, Ont., Canada). ACM, New York, 1982, pp. 60-72.

Anderson et al., "Dynamic Optical Interconnects: Volume Holograms as Two-Port Operators", Appl. Opt. 26, 5031, 1987.

Barakat R. and Reif, J. Lower bounds on the Computational efficiency of optical computing systems. Journal of Applied Optics, vol. 26, March 15, 1987, pp. 1015-1018.

Barakat, R and Reif, J. A Discrete Convolution Algorithm for Matrix Multiplication with Application in Optical Computers. Journal of Applied Optics, vol. 26, 1987, pp. 2707-2711.

R. Barakat and J. Reif, "Optical Expanders", Oct 1989, Aiken Comp Lab, Harvard University, Cambridge, Mass.

Batcher, K. Sorting networks and their applications. In AFIPS Spring J Joint Computer Conference, vol. 32. AFIPS Press, Reston, Va., 1968, pp. 307-314.

Collier, R.J., Burckhardt, C.B., and Lin, L.H. Optical Holography. Bell Telephone Laboratories, Murray Hill, New Jersey, 1971.

W. T. Cathey and K. M. Johnson, "High Concurrency Data Bus Using Arrays of Optical Emitters and Detectors", Appl. Opt. 18, 1687, 1979.

N.H. Farhat and D. Spaltis, "Optical Implementations of Associative Memmory Based on Models of Neural Networks".

Feitelson, D. Optical Computing, A Survey for Computer Scientists. The MIT Press, Cambridge, Mass., 1988.

J Ganar, Optical Communication Systems, Englewood Cliffs, NJ, 1984.

J. W. Goodman and K. M. Johnson, "Optical Computing Research", Final Report to AFOSR No. L722-6, March, 1981.

J. W. Goodman and K. M. Johnson, "The Incoherent Matrix-Vector Multiplier", Final Report to AFOSR No. L701-1, April, 1981.

Hariharan, P. Optical Holography; Principles, Techniques and Applications. Cambridge University Press, Cambridge, Mass. 1984.

Iizuka, K. Engineering Optics, Second Edition. Springer-Verlag, Berlin, Heidelberg, Germany, 1985 and 1987.

K. M. Johnson, L. Hesselink, and J. W. Goodman, "Holographic Reciprocity Law Failure", Applied Optics, Vol 23, p. 218, Jan 15, 1984.

K. M. Johnson et al., "Multiple-Multiple Exposure Hologram", Appl. Opt. 24, 4467, 1985.

K. M. Johnson et al, "Optical Computing and Image Processing with Ferroelectric Liquid Crystals," Opt. Eng, 26, 385, 1987.

K. Johnson and J. Reif, "Very High Speed Holographic Message Routing For Parallel Machines", (funded DARPA proposal) fall 1987.

H Kressel, Semiconductor Devices for Optical Communication, Springer Verlag, 1980.

E J. Lerner, "Parallel Processing Gets Down to Business", High Technology, July 1980.

E.S. Maniloff, K.M. Johnson, and J. Reif "Holographic Routing Network for Parallel Processing Machines"
EPS/Europtica/SPIE International Congress on Optical Sceince and Engineering, Paris, France, April 1989.

A. McAulay, "Optical Crossbar I3.2
nterconnected Digital Signal Processor with Basic Algorithms", Real Time Signal Processing Conf VIII, 29th SPIE,Aug. 1985.

A. McAulay, Opt. Eng., June 1986.
Suematsu, Yusuhavu, "Introduction to Optical Fiber Communications", Wiley Pub, New York, NY, 1982.

L. Messick, D. A. Collins, and D. L. Lile, IEEE Elec. Device Lett. EDL-7, 680 1986.

J. A. Neff, "Optical Crossbar Switch to be Developed for Strategic Computers", International Solid State Circuits Conference, Feb 14, 1985.

J.R. Reif and A. Tyagi, "Optimal Algorithms Optical Computing with the DFT Primitive", Jan, 1990, Comp Science Dept, UNC, Chapel Hill, NC.

A.C. Strasser, E.S. Maniloff, K. M. Johnson, S.D.D. Goggin, "Procedure for recording Multiple-exposure holograms with equal diffraction efficiency in photorefractive media, Optics

Letters, Vol 14, Page 6, Jan 1, 1989.

D. L. Staebler et al., "Multiple Storage and Exposure of Fixed Holograms in Fe-doped LiNbO3", Appl. Phys. Lett., 26, 182 (1975).

A. Tyagi and J.H. Reif, "Energy Lower Bounds in Optical Computing", Oct, 1989, Comp Science Dept, UNC, Chapel Hill, NC.

Upfal, E. Efficient schemes for parallel communication. In
Proceedings of the ACM, New York, 1982, pp. 53-59.

J. Ullman, Computational Aspects of VLSI, Computer Science Press, Maryland, 1984.

L. Valiant, "A Scheme for Fast Parallel Communication," SIAM J. Computing, 11, 2, 350-361 (1982).

Valiant, L.G. and Brebner, G.J. Universal schemes for parallel communications. In Proceedings of the 13th Annual ACM Symposium
on Theory of Computing (Milwaukee, Wis., May 11-13). ACM,
New York, 1981, pp. 263-277.

J. S. Weiner, D.B.A. Miller and D. S. Chemla, Appl. Phys. Lett, 50, 842, 1987.

J. Wilde et al., "Two Wavelength Photorefractive Optical Interconnect", Proc. SPIE Vol. 881, p 31, 1987.

T. H. Wood, R. W. Tkach and A. R. Chraplyvy, Appl. Phys. Lett 50, 798, 1987.

Appendix: Practical Issues

A.1 Storage and Retrieval in Multiple exposure holograms

Our holographic method for message routing is constrained by the number N of multiple-exposures that can be made on a hologram. (This, of course is the limit N on the number of destinations we can allow our the holographic message routing.) Studies by the [Johnson, Hesselink, Goodman,84] detail methods for storing multiple-exposure holograms that reconstruct with equal diffraction efficiency in silver halide emulsions. Staebler et al. [75] showed that at least 512 multiple holographic exposures can be recorded in three-dimensional media, as long as the distinct reference beams enter at angular displacements of at least pi/1000. Therefore, our holographic message routing system appears to be practical in the case of N = 512 messages and N = 512 processors. (Recently IBM has developed optical materials which appear to allow us to increase the bounds on N to a few thousand)

Many researchers have made multiple-exposure holograms on single silver halide photographic plates and real-time holographic recording materials such as $LiNbO_3$. Researchers at IBM, Yorktown have shown 256 images can be stored in holographic media, and have reconstructed an individual image with relatively low crosstalk between images. [Johnson et al, 1985] have shown that incoherent superposition of holograms reduces the object-object intermodulation terms , but results in a bias buildup. [Johnson et al,1989] has recently investigated the noise between superimposed holograms and in particular, to analyze how the holographic storage scales with size. Nevertheless, a detailed crosstalk analysis needs to be done.

. Researchers are currently investigating many holograms can be recorded and reconstructed in a real-time photorefractive media such as BGO and $LiNbO_3$ with reasonable signal-to-noise ratio. McRuer, Wilde, Hesselink, and Goodman [Wilde et al, 1988] showed storage and retrieval of five images in BG0 crystals without pushing the state-of-the-art. but where limited by the angular selectivity profile which has two read-out states.

In contrast, in $LiNbO_3$ crystals, the matching condition has a one-to-one correspondence with the read out image. . [Johnson et al,89] have been able to record and individually retrieve up to 128 holograms in $LiNbO_3$ crystals, and furthermore Johnson and others feel that many more holograms can be recorded. [Anderson and Linginger, 1987] have investigated phase, polarization and intensity encoding schemes to separate the patterns to be matched by the holograms so that crosstalk is reduced and have demonstated experimentaly storage and retrieval of over ten 1-D vectors in $LiNbO_3$ using phase encoding techniques to pseudorandomize the stored data.

[Johnson et al, 1985] showed it was feasible to make multiple-exposure holograms, each containing a small number of images of CT medical data (10^6 points/images) on a single silver halide photographic plate.

Technical issues currently being addressed in this part of the project include:

(1) Analyzing the noise between recording incoherent and coherent superimposed holograms in silver halide and real-time holographic recording materials. We will analyze how the holographic storage scales with size. Can 256 images be stored in holographic media, and reconstruct an individual image with low crosstalk between images. Incoherent superposition of holograms reduces the object-object intermodulation terms [Johnson et al, 1985], but results in a bias buildup. We pursuiting a detailed crosstalk analysis.

(2) These holograms can be made in a real-time photorefractive media like $LiNbO_3$. Technical issues to be addressed here include how many holograms can be recorded and reconstructed with reasonable signal-to-noise ratio. This analysis has been started by McRuer, Wilde, Hesselink, and Goodman [Wilde et al, 1988] for applications to dynamic optical interconnects in BG0 crystals. The $\vec{k}$ space arguments in the latter work shows storage and retrieval of five images without pushing the state-of-the-art. BG0 is somewhat problematic because of the angular selectivity profile has two read-out states. In $LiNbO_3$ crystals, where the Bragg matching condition has a one-to-one correspondence with the read out image, it should be possible to record and individually retrieve much more than 5 holograms.

(3) We are investigating phase, polarization and intensity encoding schemes to further orthogonalize the holograms such that crosstalk is reduced. Anderson [Anderson and Linginger, 1987] has shown storage and retrieval of over ten 1-D vectors in $LiNbO_3$ using phase encoding techniques to pseudo-orthogonalize the stored data.

(4) In the present optical implementation, N holograms, each containing N images, are required to connect N to N processors. We are exploring reducing this number, via multi-layer holographic structures which should yield a more compact sized routing system.

(5) Lastly, we are analyzing the size and contrast ratio demands placed on the spatial light modulator by the angular selectivity and noise in the silver halide and dynamic holographic storage media.

A.2 Fast Switching SLMs

Recall that each destination switch array is very short; in fact in the case where N = 256, the array is of length only 2log(N) = 16. There are under development within the Center for Optoelectronic Computing Systems linear III-V CCD/SLM arrays of this length and greater that can switch at rates of 1.0 gigahertz.

We have observed above that the time to set the destination switches is 16 nanoseconds, depending on whether conventional or very high speed electro-optical switches are used. This gives the ultimate limit on the total time to set the switches for a general routing of 512 messages.

A.2.1 Spatial Light Modulator Technology

We are using ferroelectric liquid crystal (FLC) and III-V CCD SLMs to demonstrate the message router. The FLC SLMs are currently available with switching speeds on the order of 10 to 100 microseconds [Johnson et al, 1987]. This allows the holographic designers to work with a 1-D array while the fabrication of fast switching III-V CCD SLMs is ongoing at Colorado State University under the direction of Professor Derek Lile.

A.2.2 InP/InGaAs CCDs for High Speed Spatial Light Modulation and Beam Steering

Two types of SLM are being investigated for this application, both based on the

InP/InGaAs system and both an outgrowth of ongoing electro-optic device development at the Center for Optoelectronic Computing Systems. One device, which relies on electro-absorptive and electro-refractive excitonic effects within gas source MBE grown InGaAs multiquantum wells, is programmed electrically using a CCD gating scheme. Messick, Collins, and Lile [86] have demonstrated that such CCDs, with 20 x 120 μm pixels can operate to frequencies in excess of 1 gigahertz and as a result these SLM arrays will clearly meet the switching speed requirements of the presently proposed message router. The second device, which will be much easier to implement and hence offers the prospect of much larger array sizes, is optically addressed and consists of a regular pattern of MQW mesas etched on the surface of the InP substrate. Using the mechanisms proposed by Wood et al. [87] and Weiner et al [87], we would expect to implement beam steering an beam modulation in response to selective optical addressing of these discrete mesa areas. Initial structures would be based on 100 μm pixel sizes on 200 μm centers. Later devices could be scaled in either direction to meet machine requirements. In either case both devices will be fast, with switching speeds well within the requirements of this application.

A.3 The Prototype System

Workers in areas of computer architecture and electronic design (Duke University), holographic storage and retrieval (University of Colorado), and III-V SLM fabrication (Colorado State University) have come together to build a proof-of-principle holographic message routing system. A currently running small prototype system which does the routing of 4 messages among 4 processors at moderate speeds. For implementation details see [Maniloff, Johnson and Reif,89]. An upgraded system using faster SLMs would use 4 to 10 nanoseconds per routing. At a slower rate of approximately 1 to 10 microseconds per routing, our prototype holographic routing system has the potential to route 512 messages among 512 processors . The actual construction of the optical components of this prototype system is done by Kristina Johnson and Eric Manalof at the Center for Optoelectronic Computing Systems, University of Colorado and Colorado State University.

A.4 Overcoming Limitations by Pipelining

Although our basic holographic routing system currently has a limit of approximately N = 512, it can be augmented to service a system with a much large number of processors (say N = 65,000) by the use pipelining and the use of a number of small conventional fixed connection routing systems. The idea is that since each of these processors processes message requests at a very high rate, they can be made to operate in a pipelined mode to service a large number of additional message routing requests. For example, to implement routing of 65,536 messages on the 65,536 processor CONNECTION machine, we can randomly partition the processors into 512 groups each of size 512. Then we proceed with a series of routing rounds, where on each round we route 512 messages, one from each group to the 512 groups. Even with packet destination conflicts, it can be shown using the techniques of [Valiant,1982] that less than 1024 rounds will expect to be required. The final stage required is local routing of 512 messages within each group of destination processors. This can easily be implemented in time at most few microseconds by conventional electronic integrated circuits

Therefore, the total time required by our full scale holographic message routing system to implement a general routing of 65,536 messages on 65,536 processors will be approx 16 microseconds using state of the art SLMs (say III-V CCD/SLMs).

A.5 Work on Electronic Routing

The purely electronic subsystem which will be required to do message routing within a group of 64 processors in time approximately 1 to 10 microseconds. CMOS VLSI technology is sufficient to enable us to achieve these rates. For the scaled down planned system, we require a total of 4 of these systems. (For the full scale 65,536 processor system, we will eventually require a total of 128 of these subsystems, each of which will do message routing for a group of 512 processors.)

We require an electronic subsystem which combines the 64 messages received from a group of 64 processors, and generate a single very high speed stream of these messages. Each resulting high rate binary message stream will be fed into an electro-optical transducer (such as a GaAs Laser). In the scaled down prototype version, there will be only 4 of these streams input to our prototype holographic routing system. (In the full scale system there will eventually be 128 of these message streams, each of which will be an input to our holographic routing system.). We are investigating the issues by use of software simulation. Currently, we have developed a software simulation system that does the routing on the BLITZEN system, which is a 4k processor SIMD parallel machine.

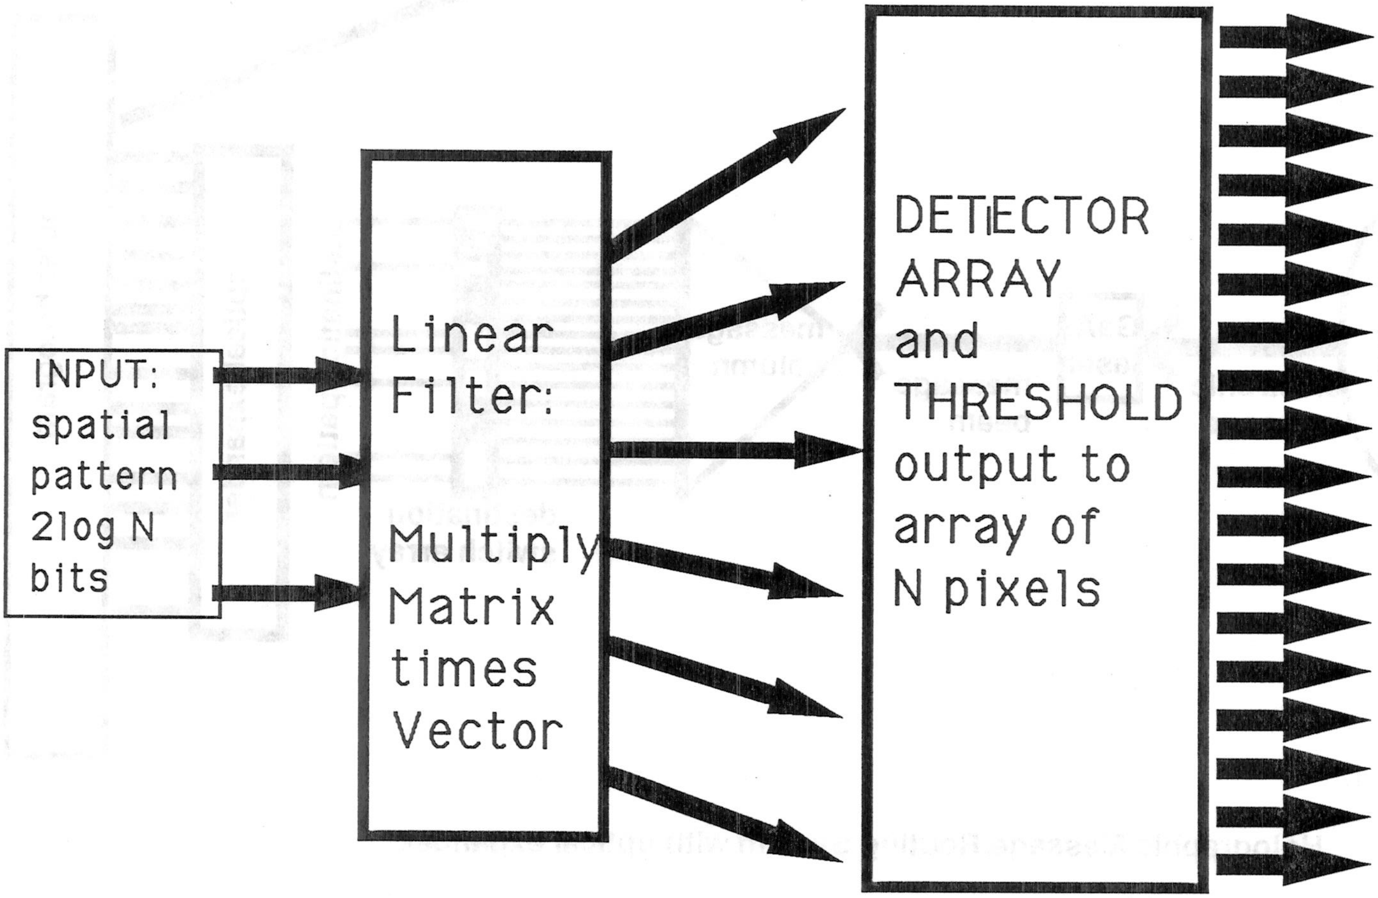

Optical Expander System

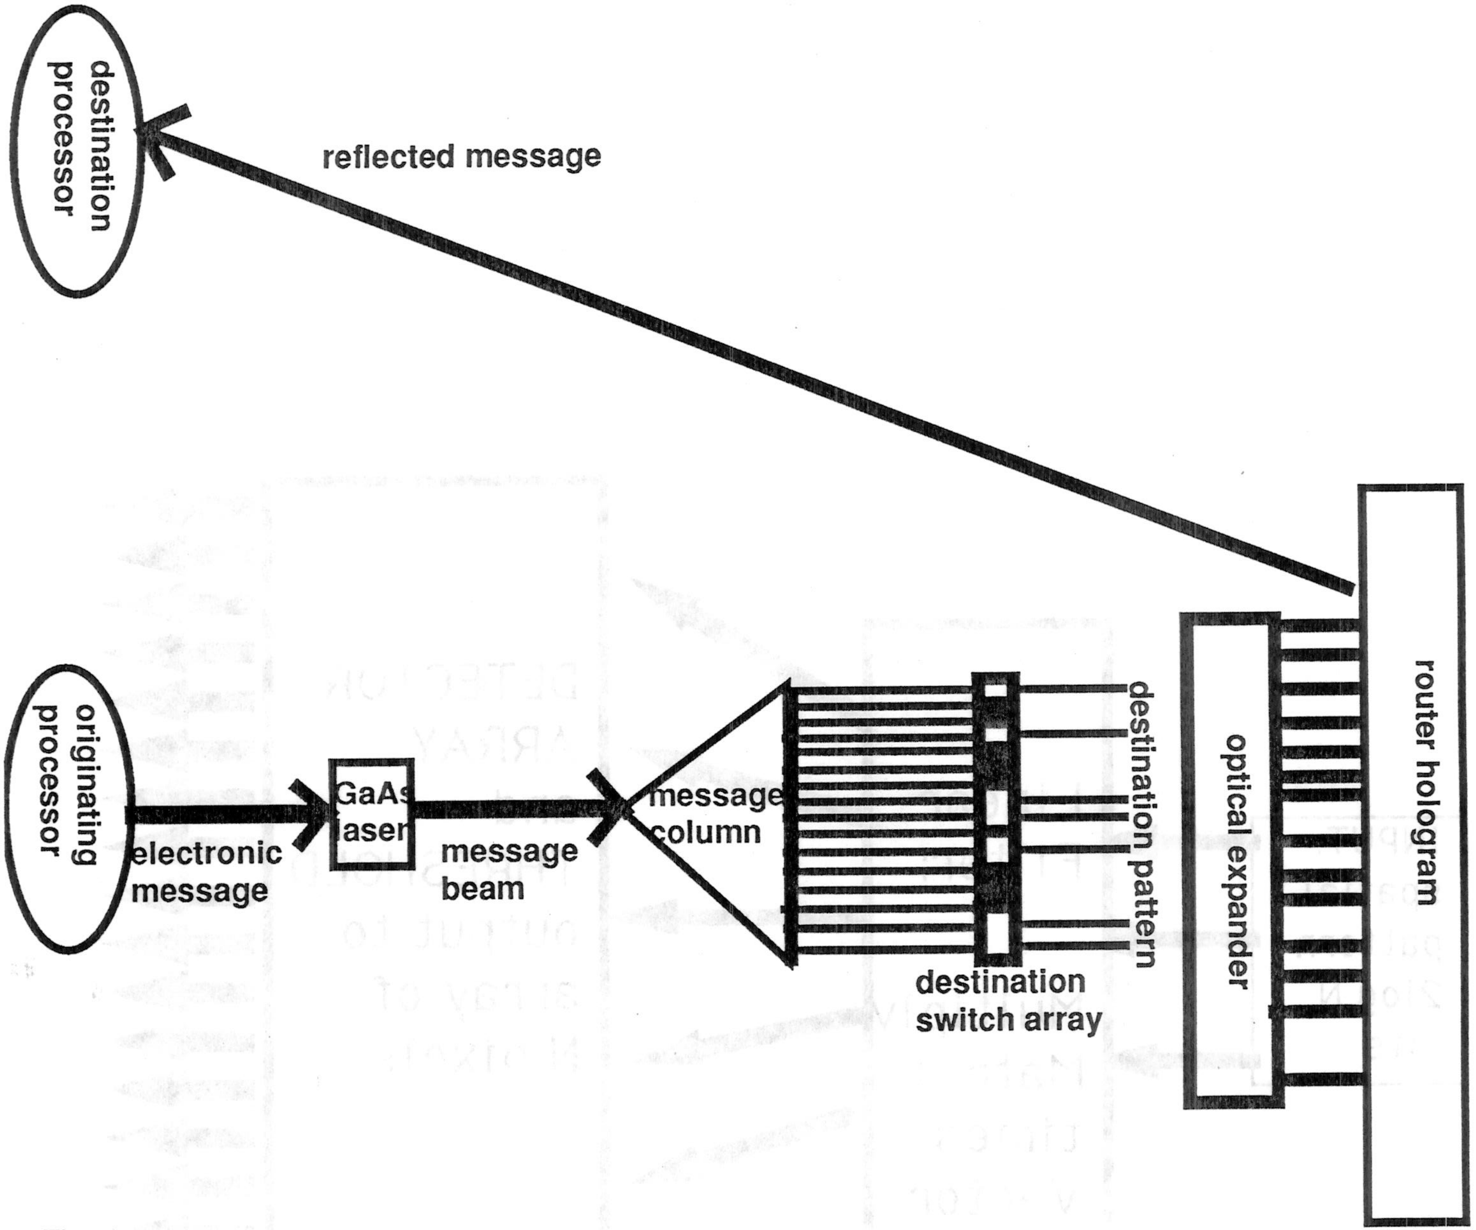

Fig. 1.

Holographic Message Routing System with optical expander:

Illustration of the devices required by each originating processor.

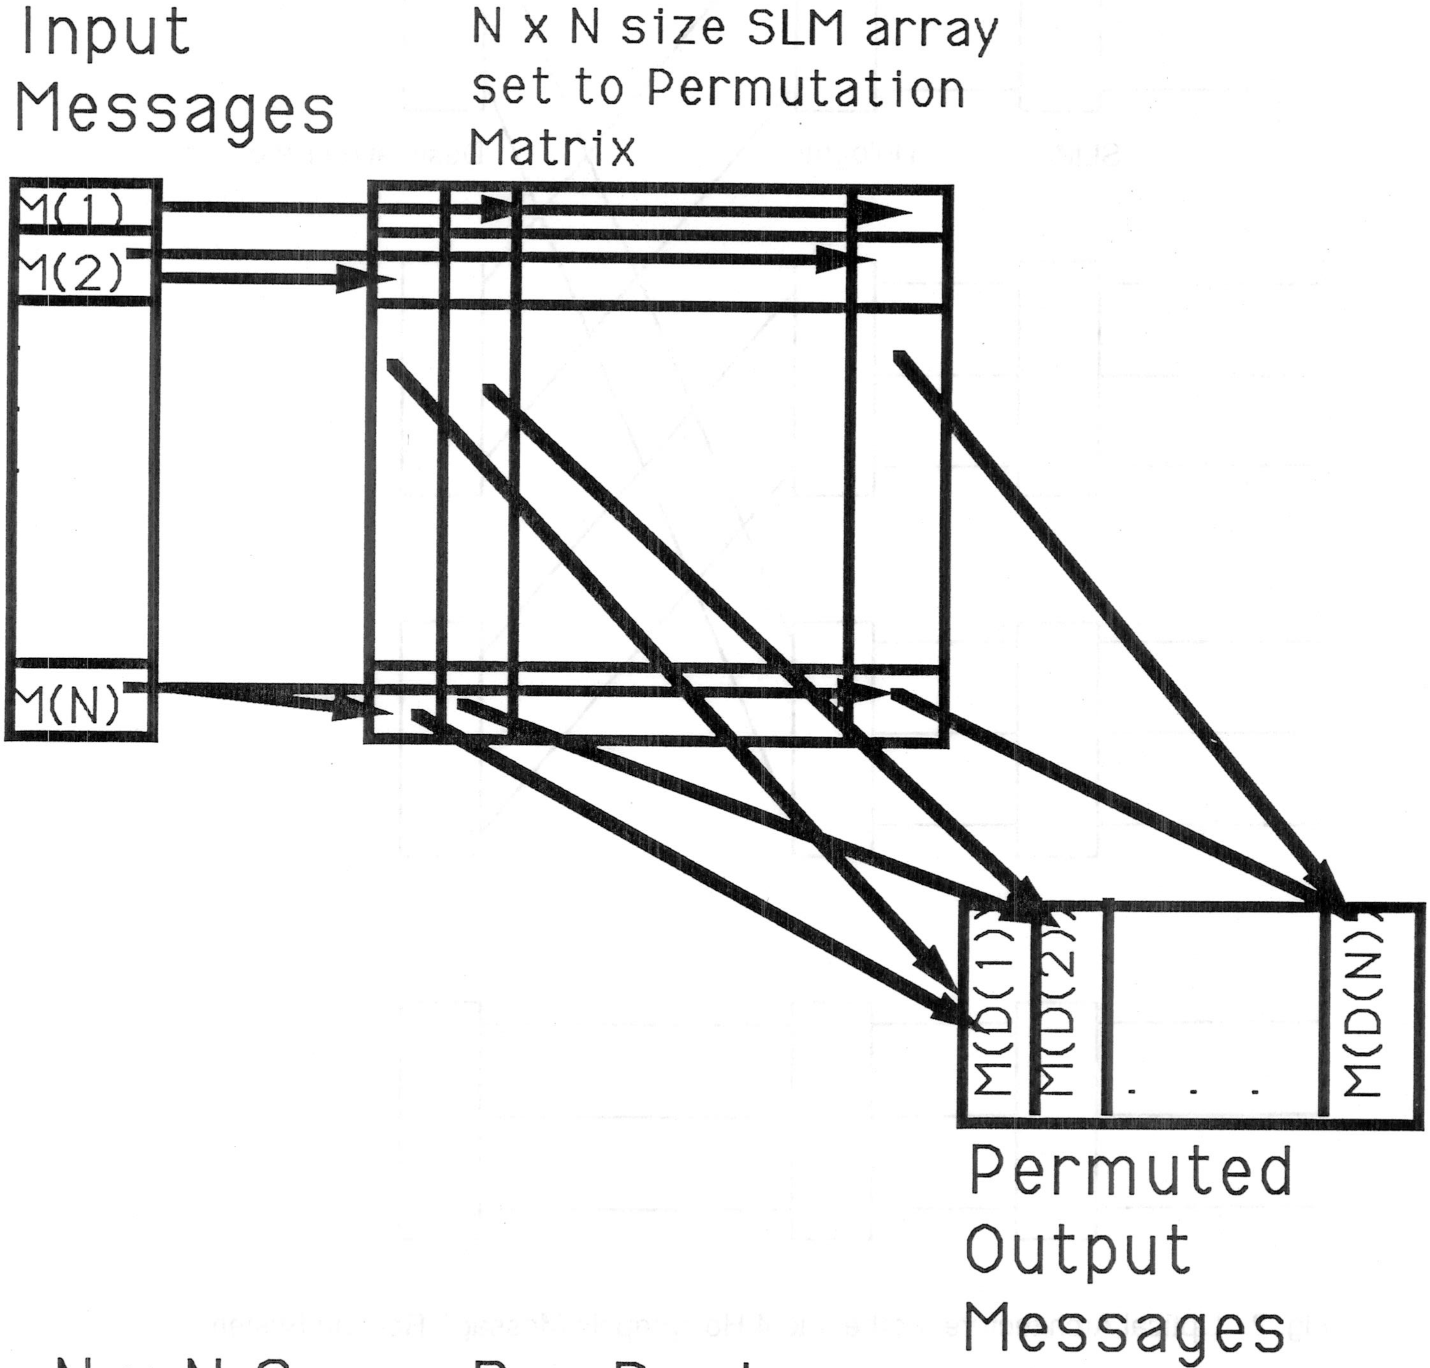

N x N Cross Bar Router using a N x N Matrix, Vector Multiplier

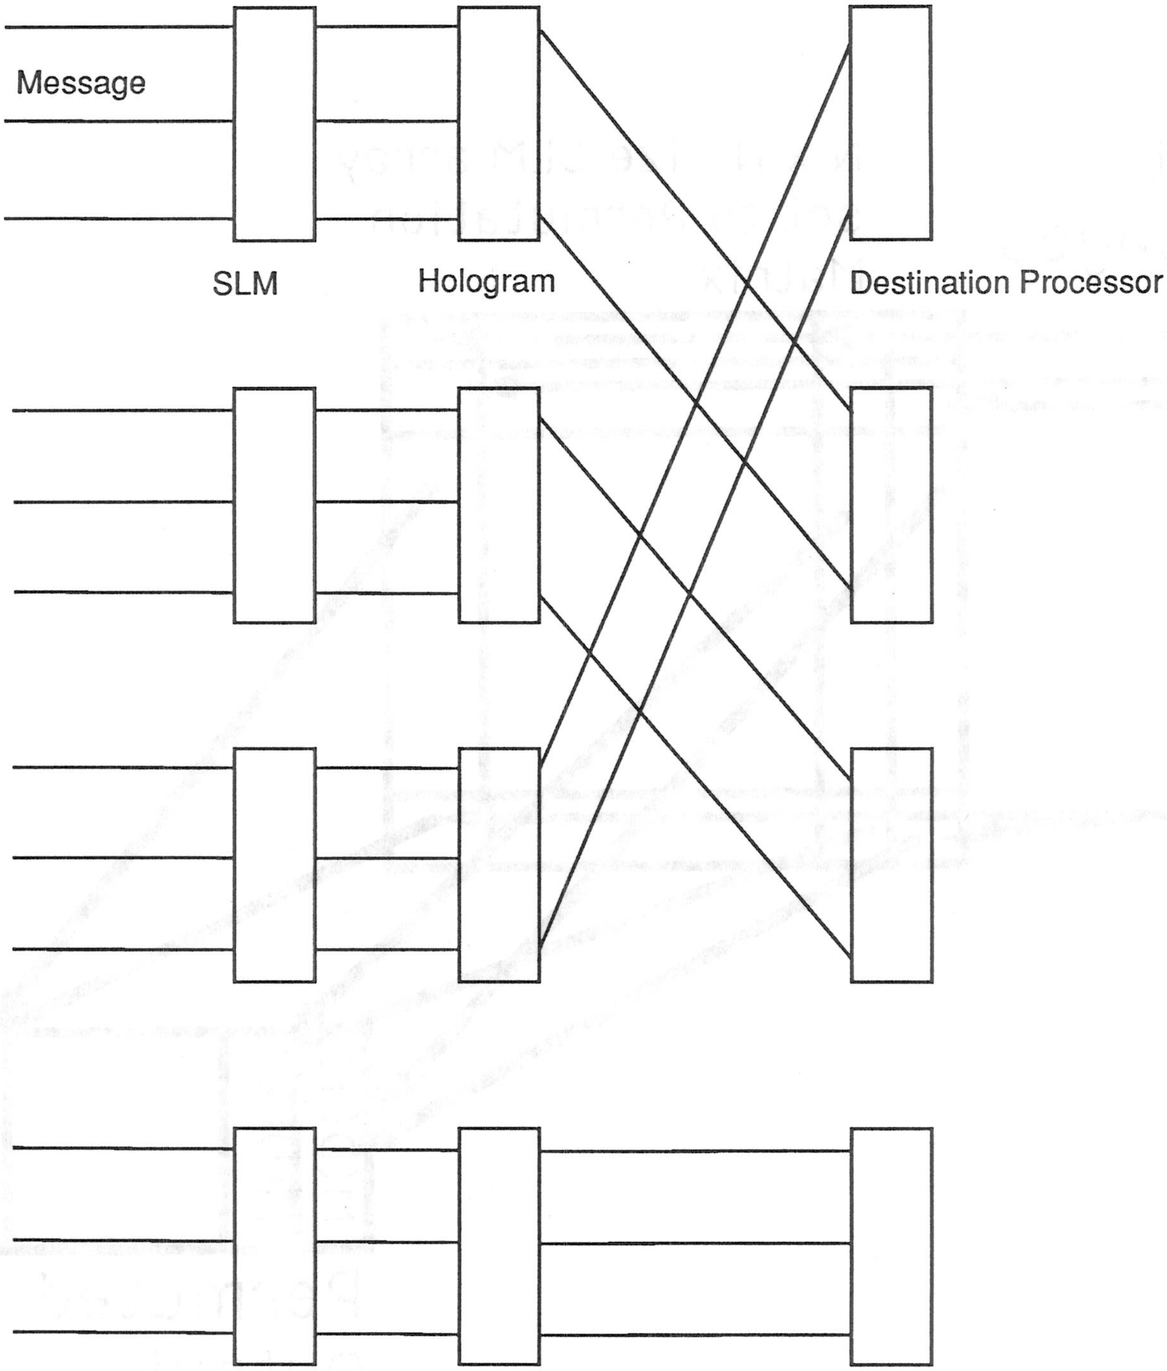

Fig. 2. Optical Architecture for the 4 to 4 Holographic Message Routing System