An Optical Network Interface Unit for Multichannel Ring Networks

Francis Reichmeyer, Salim Hariri, Wang Song and Kamal Jabbour
Department of Electrical and Computer Engineering
Syracuse University

Abstract

In this paper we present an all optical hypercube-based multichannel network, implemented on a ring topology, using wavelength division multiplexed channels. For a network of size N, only $\log_2 N$ distinct channels are required. We design an Optical Network Interface Unit (ONIU) which employs free-space optics to perform routing and switching functions. The design eliminates the electronic bottleneck at intermediate nodes and is capable of perfoming high speed packet switching. ONIU can be reconfigured in GHz range and provides high bandwidth necessary to perform distributed and parallel computing in a local area network environment.

1 Introduction

Optical fiber has been widely accepted as the interconnection medium for communication networks, due mainly to the advantages offered by fiber such as higher bandwidths, better noise performance and longer distances without repeaters, as compared to electrical media. However, the rate at which current electro-optic devices can access the fiber is on the order of four to five times slower than the bandwidth of the fiber, which can be as high as 30 terahertz. Hence, there has been a great deal of research in designing system architectures to eliviate this so called *electronic bottleneck* and improve the performance of communication networks.

Multichannel lightwave networks have been proposed as a solution to this problem. In multiple channel systems, several users can access the fiber simultaneously thus increasing the overall network capacity. The optical bandwidth is divided into several channels, each capable of transmitting noninterfering signals at distinct wavelengths. The channels are then wavelength division multiplexed (WDM) and transmitted over the network. Examples of such networks are the ShuffleNet Multihop Networks [1, 2, 3] which implement a perfect shuffle interconnection scheme, the Wavelength Division Multiple Access Channel Hypercube (WMCH) [4] which uses tunable electro-optical devices and a passive star-coupled topology and the Hypercube-Based Multichannel Ring Network [5] which requires fewer number of channels and static devices.

This paper presents the design of an all Optical Network Interface Unit (ONIU) to implement routing and switching functions for the hypercube-based local area ring network. Based on the Optical Interface Message Processor [6], the ONIU fully exploits the inherent parallelism and multidimensionality offered by current free-space optical technology to perform switching without conversion of the optical signal to the electrical domain. This design differs from the proposed *hybrid* systems. In the hybrid designs, an interface message processor strips the packet header from the message and converts it into electrical signals in order to control the setup of the electrical switches. These electrical switches, in turn, control the state of the photonic switches that route packets at the input ports to the required output ports [7]. By requiring electro-optical conversions only at the source and destination, the ONIU eliminates the electronic bottleneck at intermediate nodes.

The rest of this paper is organized as follows. Section 2 presents a brief overview of conventional hypercube networks along with the architechture and routing techniques used in hypercube-based multichannel ring networks. In Section 3 we describe the design and implementation of the main functions performed by the ONIU. Finally, in Section 4 we outline our conclusions and future work planned for this area of research.

2 Hypercube-Based Networks

An r-dimensional hypercube can be represented as $N = 2^r$ nodes located at the vertices of a hypercube graph. Each node is represented by a binary r-tuple

0-8186-3670-X/93 $03.00 © 1993 IEEE

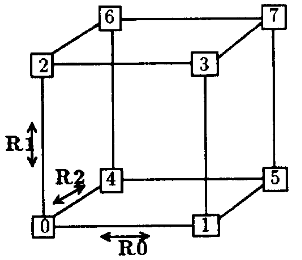

Figure 1: Hypercube Interconnection

Network Size	# Channels Required		
N	ShuffleNet Bus	ShuffleNet Ring	Hypercube-Based Ring
8	16	4	3
32	64	20	5
64	128	12	6
128	256	152	7
1024	2048	356	10
2048	4096	212	11

Figure 2: Comparison of required channels for ShuffleNet and Multichannel Ring Networks

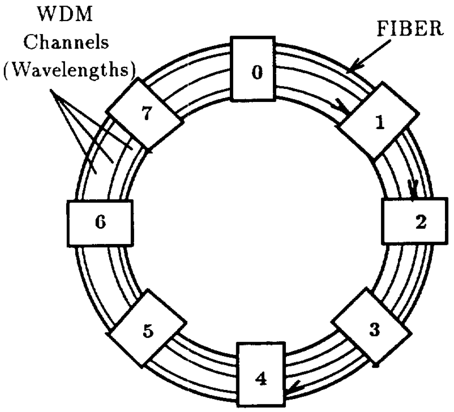

Figure 3: Hypercube-Based Multichannel Ring, $r = 3$

$(p_{r-1}, p_{r-2}, \ldots, p_0)$ and is connected to r immediate neighbors, those with r-tuple addresses differing in exactly one bit position, according to the routing functions $R_0, R_1, \ldots, R_{r-1}$ associated with each bit position. Figure 1 shows a simple hypercube of dimension $r = 3$ along with corresponding routing functions R_0, R_1, R_2.

Function R_i defines a connection between nodes with addresses differing in bit position i. For example, in Figure 1, node 0 is connected to neighboring nodes 1 (001), 2 (010) and 4 (100) by R_0, R_1 and R_2, respectively. In the hypercube, packets are routed, according to an e-routing scheme [8], on R_i if source and destination addresses differ in bit position i.

2.1 Hypercube-Based Ring Architecture

In [5] we introduced a local lightwave network, based on a hypercube interconnection, implemented on a multichannel ring topology. It has been shown that the same operations used in the perfect shuffle interconnection, on which the ShuffleNet systems are based, can be used to construct a logical hypercube network [8]. However, in the ring network, concurrent communication is achieved over fewer wave-division multiplexed channels than are needed by similar size ShuffleNet networks. By using fewer channels, the cost of the network will decrease and the available bandwidth will be utilized more efficiently. Figure 2 compares the number of channels required for similar size ShuffleNet and Hypercube-Based Multchannel Ring Networks.

The multichannel ring based on the r-dimensional hypercube has $N = 2^r$ Optical Network Interface Units (ONIUs) connected by a fiber ring containing $r = \log_2 N$ distinct channels, $C_0, C_1, \ldots, C_{r-1}$. The WDM channels transmit messages at wavelengths $\lambda_0, \lambda_1, \ldots, \lambda_{r-1}$ respectively. The ONIUs, like the hypercube nodes, have r-tuple addresses $(n_{r-1}, n_{r-2}, \ldots, n_0)$ and the channels provide point-to-point connections according to the routing functions defined in the hypercube. That is, channel C_i connects ONIUs whose addresses differ in bit position i. Figure 3 shows an implementation of a multichannel ring based on the 3-dimensional hypercube of Figure 1. It has $N = 8$ (2^3) ONIUs and requires $r = 3$ distinct channels. We see the point-to-point connections for ONIU 0, supplied by the channels C_0, C_1, C_2 which transmit at wavelengths λ_0, λ_1 and λ_2. In the remainder of the paper, we use the 3-dimensional hypercube-based ring as a running example to explain the design and implementation of a high performance distributed system based on optical technology. Extensions to larger size networks can be made in the obvious way.

2.2 D-Routing in Hypercube-Based Ring Networks

In order to implement a logical hypercube on the physical multichannel ring we employ a routing scheme similar to that described above for the hypercube network, exploiting the point-to-point connections established by the WDM channels $C_i, i = 1, 2, \ldots r-1$. Recall that in e-routing, packets or messages are routed using address information (node addresses) to choose a channel to transmit the message. In the multichannel ring, we use a routing scheme called ***d-routing*** (distributed e-routing) where if addresses differ in bit position i then channel C_i is chosen.

One difference between routing in the conventional hypercube and that of the multichannel ring is that in the hypercube, links are independent of one another and so more than one node can transmit in a particular dimension at the same time. That is, node 0 can send to node 4 on R_2 at the same time node 2 sends to node 6 on R_2 without a collision (see Figure 1). However, the same transmissions on the ring result in a collision because these connections share a single multiaccess channel C_2. Therefore, in d-routing, routing tables for each ONIU are constructed based on address bit comparisons performed in different orders. For example, in ONIUs 0-2, the least significant bit is compared first, then the next significant bit, then the most significant bit; in ONIUs 3-6, the second significant bit is compared first, then the most, then the least; and in ONIUs 7-8, the most significant bit is compared first, then the least, then the second. By changing the order of the bit comparisons made at the ONIUs, the traffic entering the network is more evenly distributed over all channels and hence the number of resulting collisions can be reduced.

When constructing the routing tables, our d-routing scheme also accounts for the physical topology of the ring network. The routing path of a message will consist of intermediate ONIUs whose addresses differ in exactly one bit position (according to the hypercube, Figure 1). If an ONIU is not in the routing path, it will be ***bypassed***, that is the switch is configured so as to output the message on the same wavelength as it was received. If it is in the routing path then the ONIU switch will be set to transmit the message on the appropriate wavelength, each wavelength corresponding to a dimension in the hypercube.

We also found that the physical ring topology could be exploited to reduce the number of hops (either physical or logical) needed by a message between certain pairs of ONIUs. For example, going from node 7 to node 0 in the hypercube requires three hops, becuase the addresses differ in three bit positions (Hamming distance = 3). In the multichannel ring this is reduced to one physical hop. As another example, consider a message going from node 1 to node 6. Again, three hops are required in the hypercube structure (e.g. $1 \rightarrow 0 \rightarrow 6$). But in the ring (assuming unidirectional clockwise transmission) one ***logical*** hop can be used. Following the hypercube routing path, $1 \rightarrow 0 \rightarrow 6$ the message reaches the destination ONIU 6 while trying to reach the imtermediate ONIU 0. So, by requiring only one logical hop, the number of frequency translations needed is reduced and thus so is the delay.

Since switching and routing functions will be performed in the optical domain by the ONIUs, the delay reduction is not as significant as it would be in an electrical switching system. However, this example shows that by exploiting the ring topology, some ONIU routing tables will not have to support certain source/destination pairs. For example, since ONIU 0 is not reached when ONIU 1 transmits to ONIU 6, then ONIU 0 needs not have an entry in its routing table for this pair. Thus due to the ring topology, the size of the optical routing table in each ONIU is reduced.

The d-routing scheme discussed above can be implemented using either packet switching or circuit switching. In both cases routing tables are constructed for each ONIU and used in routing messages. In the remainder of this paper we assume packet switching is used. Figure 4 shows sample routing tables created for the hypercube-based mutichannel ring ($r = 3$) using the d-routing scheme. Depending on the source/destination pair in the packet header, one of the three wavelengths (represented by a one in that bit position) is selected to transmit the packet. In the following section, we discuss how the ONIUs use these routing tables to implement optical routing functions.

3 Optical Network Interface Unit

In the previous section we have discussed the architecture and design of a hypercube-based multichannel ring network and presented the d-routing scheme to implement a logical hypercube on the physical ring topology. We now discuss how the routing functions can be performed in the optical domain, requiring *o/e* conversions only at the packet source and destination, thus eliminating the electronic bottleneck. This is done using the Optical Network Interface Unit which is capable of performing high-speed packet switching for static routing schemes.

The ONIU is an adaptation of our Optical Interface Message Processor (OPTIMP) [6]. It performs high speed packet switching by fully utilizing the inherent parallelism and multidimensional capability offered by optics. The hypercube-based ring network is particularly well suited for optical switching because it is a

ONIU 0

S	D	Channel 0	1	2	Rcv
X	0	0	0	0	1
0	1	1	0	0	0
0	2	0	1	0	0
0	3	1	0	0	0
0	4	0	0	1	0
0	5	1	0	0	0
0	6	0	0	1	0
0	7	1	0	0	0
2	1	0	1	0	0
3	1	0	1	0	0
3	2	1	0	0	0
4	1	1	0	0	0
4	2	0	0	1	0
4	3	0	0	1	0
⋮			⋮		
7	4	1	0	0	0
7	5	0	1	0	0
7	6	1	0	0	0

ONIU 1

S	D	Channel 0	1	2	Rcv
X	1	0	0	0	1
0	2	0	1	0	0
0	3	0	1	0	0
0	4	0	0	1	0
0	5	0	0	1	0
0	6	0	1	0	0
0	7	0	1	0	0
1	0	1	0	0	0
1	2	1	0	0	0
1	3	0	1	0	0
1	4	1	0	0	0
1	5	0	0	1	0
1	6	1	0	0	0
1	7	0	1	0	0
⋮			⋮		
7	4	1	0	0	0
7	5	0	1	0	0
7	6	1	0	0	0

ONIU 2

S	D	Channel 0	1	2	Rcv
X	2	0	0	0	1
0	3	0	1	0	0
0	4	0	0	1	0
0	5	0	0	1	0
0	6	0	0	1	0
0	7	0	1	0	0
1	0	1	0	0	0
1	3	0	1	0	0
1	4	1	0	0	0
1	5	0	0	1	0
1	6	1	0	0	0
1	7	0	1	0	0
2	0	0	1	0	0
2	1	1	0	0	0
⋮			⋮		
7	4	1	0	0	0
7	5	0	1	0	0
7	6	1	0	0	0

ONIU 7

S	D	Channel 0	1	2	Rcv
X	7	0	0	0	1
1	0	1	0	0	0
2	0	0	1	0	0
2	1	0	1	0	0
3	0	0	1	0	0
3	1	0	1	0	0
3	2	1	0	0	0
4	0	0	0	1	0
4	1	0	0	1	0
4	2	0	0	1	0
4	3	0	0	1	0
5	0	0	0	1	0
5	1	0	0	1	0
5	2	0	0	1	0
⋮			⋮		
7	4	1	0	0	0
7	5	0	1	0	0
7	6	1	0	0	0

X = 'Don't Care'

Figure 4: Sample Routing Tables

static structure requiring a simple deterministic routing scheme and nontunable optical devices. Therefore optical switching and routing is realizable using currently available optical technology. Figure 5 shows a block diagram of the proposed ONIU, for a hypercube-based ring network of dimension $r = 3$, which requires three wavelength division multiplexed channels (Figure 3). Extension to higher dimensions is obvious.

A signal enters the demultiplexor (DEMUX) where a beam splitter (BS) divides the incoming optical power among three optical filters, tuned to the wavelengths λ_0, λ_1 and λ_2. Each filter passes one wavelength which is sent to a Space-Time Conversion Unit (STCU) and to a delay line. The STCU extracts the source/destination (SD) information from the packet header and stores it and its complement ($\overline{SD}$) onto optical shift registers. The delay line delays the propagation of the message signal until the ONIU can configure the switching network according to the SD information extracted from the header. The signals from each STCU are sent to the Wavelength Select Unit (WSU) which selects one input at a time in a round-robin manner.

The selected SD and $\overline{SD}$ are compared against all the columns of the routing table (RT) and its complement ($\overline{RT}$) in the Optical Routing Unit (ORU). The RT consists of all SD pairs for that ONIU and, for each pair, indicates the outgoing wavelength to be used (see Figure 4). The contents of this routing table can be stored on a two-dimensional optical array (e.g., holograms or spatial light modulators) and can be read in a single step parallel search to select the column corresponding to the incoming SD pair and, thus, generate the control signals required to configure the Optical Switching Unit (OSU). Upon completing the configuration of the OSU for an incoming packet, the packet is taken from the dalay line and is either routed onto the proper outgoing wavelength or received, if the current ONIU is the destination. The next input is then selected by the WSU and that packet is similarly processed. All outgoing packets are then wavelength-division multiplexed (WDM) and transmitted onto the ring.

By proper synchronization beween the ORU, the delay lines used, and the OSU, conflicts at the wavelength converters in OSU can be avoided. That is, given the packet size used, the delay lines can be designed [6] such that they delay a packet for a period of time long enough for the previously routed packet to be transmitted. Also, the ORU and OSU are synchronized so that the control signals needed to configure the switch for a packet are available after the previous packet has been transmitted and no conflict can occur.

In the following, we describe in detail each of the

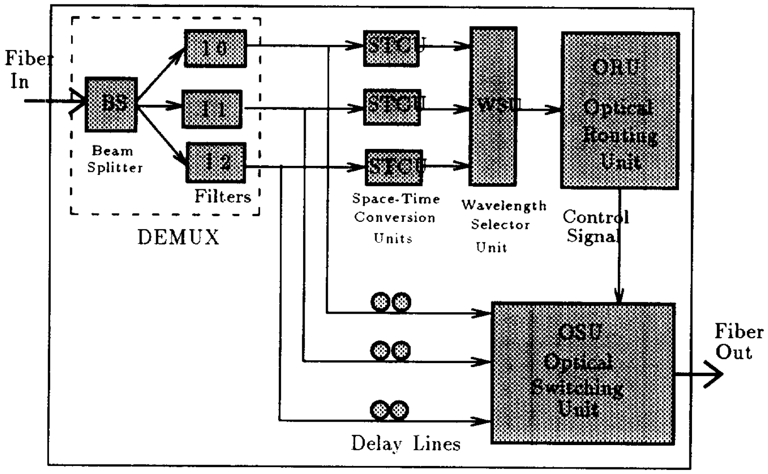

Figure 5: BlockDiagram of ONIU

functional units of the ONIU shown in Figure 5. For simplicity, most of the light guiding elements, such as mirrors, cylindrical lenses, and beam splitters are omitted from the figures.

3.1 Space-Time Conversion Unit

Figure 6 shows the architecture of the Space-Time Conversion Unit (STCU). The SD code from the header of the incoming packet is first converted from time domain to spatial domain. The purpose of this conversion is that we can then use the parallel processing advantage of optics to achieve a high-speed optical routing table search. The SD code is separated from the input data packet by a time domain AND gate and stored in an S-SEED (Symmetric Self Electro-optic Effect Device) based optical shift register [9]. The size of the shift register is $2r$ bits since each source and destination in an r dimensional hypercube requires r address bits.

SEEDs are essentially photodetectors combined with multiple quantum wells (MQWs) to give an optically controlled device with optical outputs. If the input beam is above a threshold, the device absorbs the entire incident light, thereby inverting the input signal. Hence, we first send the incoming signal through an inverter so that an incoming high bit (binary 1) will cause a SEED to be transparent and a low bit (binary 0) will be absorbed. The 1/2 reference in the other input port of the shift register (also inverted) provides the ground value for the device to determine whether the input bit is a one or zero. With the synchronization of the clock pulses, the SD code propagates and is stored on the shift register.

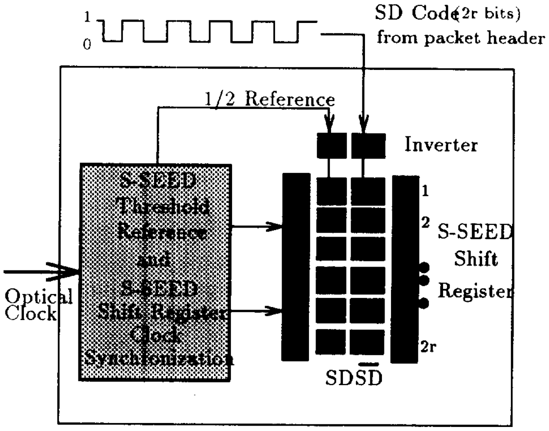

Figure 6: Space-Time Conversion Unit

Since each S-SEED produces dual outputs, both SD and its complement code $\overline{SD}$ are stored on the shift register simultaneously. Next, an imaging system reads out the SD and $\overline{SD}$ codes in parallel and displays them separately in the spatial domain, thus completing the time domain to spatial domain conversion. Since the S-SEED can operate at a rate above Gigabits/s [10], the domain conversion for each bit of the SD and $\overline{SD}$ code can also be in the range of Gigabits/s. The output of the STCUs is then sent to the Wavelength Selector Unit (WSU), a shift register which selects only one of the STCU outputs and sends it to the ORU.

3.2 Optical Routing Unit

The SD and $\overline{SD}$ selected by the WSU are displayed on two optical masks (also call them SD and $\overline{SD}$) and sent into the Optical Routing Unit (ORU). Figure 7 outlines the ORU operation for the case of $S = 0, D = 2$ extracted from the packet header. The function of ORU can be divided into two operations, parallel routing table search and control signal generation. The parallel routing table search decides which SD code is present and then based on this information, a control signal is generated in order to configure the switching unit properly.

Because of the inverter used in the STCU, the SD and $\overline{SD}$ masks use positive logic (i.e., a binary one is represented by high intensity and a binary zero is represented by low intensity). The routing table and its compliment are displayed by two programmable SLMs, RT and $\overline{RT}$. A pixel on a SLM is transparent if it is a binary one or opaque if it is a binary zero. Each column of pixels on the table corresponds to one SD pair.

In general for a hypercube structure of size $N = 2^r$, we would require N^2 SD pairs. However, due to the physical ring structure of our hypercube-based network, the number of SD pairs that need to be supported by each ONIU is significantly reduced. The total number of pairs necessary, and hence the total number or columns needed in RT and $\overline{RT}$, for the ring of N ONIUs is reduced from N^2 to $D = (N-1)(N+2)/2$; N pairs are eliminated since an ONIU will never send a packet to itself, and another $(N-1)(N-2)/2$ are eliminated due to the ring topology because a packet may reach its destination enroute to an intermediate ONIU, as was discussed in Section 2.2. Therefore, the total size of the RT array in an ONIU is $2r \times D$ pixels. For $N = 8$, this gives a 6×35 array (6 bits needed for an SD pair, 35 SD pairs).

A reduction in the size of the routing table corresponds to a reduction in the energy requirements of the system. The RT is searched by expanding each bit in the SD code over every column of the RT. Therefore the energy efficiency can, at best reach $1/D$ and so by reducing D we increase this efficiency. A smaller RT also results in shorter reconfiguration times when the network changes. The achievable frame rate, to reconfigure the routing table, of SLM is in the order of microseconds, however during normal routing operations, the patterns displayed on the SLMs (routing table entries) do not need to be changed. To read these patterns by light beams takes less than a nanosecond (10^{-9} sec. for light to travel a few centimeters). Therefore, the relative slow writing speed of the SLM does not limit the performance of the system in normal operation.

The parallel routing table look-up is based on a vector-matrix multiplication scheme. The cylindrical optical system (not shown in the figure) spreads each bit of the spatial SD and $\overline{SD}$ into a horizontal row that illuminates a corresponding row in one of the SLMs RT and $\overline{RT}$, respectively. In this manner, the input SD and its complement code, represented by two columns of intensity distributions, are *simultaneously* multiplied by each column in the routing tables. A beam combiner (symbolically shown in the figure) combines the results of the two sets of bit multiplications into one $2r \times D$ pixel array. The following cylindrical lens (not shown in the figure) focuses each column of this array onto a pixel in the one-dimensional ($1 \times D$ array) thresholding device TH. The intensity on each pixel of the one-dimensional TH array can be expressed as the Boolean Equivalence function: $\sum_{i=1}^{k}(SD)_i \cdot (RT)_{ij} + (\overline{SD})_i \cdot (\overline{RT})_{ij} = (Equivalence)_j$ where $j = 1, 2 \ldots n$, $(SD)_i$ is the ith bit of the SD code, and $(RT)_{ij}$ is the ith bit in the jth column in the routing table RT.

The Equivalence function determines the best match between the input SD code and one of the columns on the routing table. For the example case shown in Figure 7 input SD code matches the second column in the routing table and the total (combined) energy passing the second column of RT and the second column of $\overline{RT}$ will be a maximum. That is, after the beam combiner/cylindrical lens combination, the total light energy passing the second pixel of the 1×35 TH array will be a maximum (i.e., j=2).

The threshold device TH is properly set such that a pixel transmits light only if the total intensity impinging on it is above this predetermined maximum value, so there will be only one transparent pixel on the device (pixel 2 in the above example). To insure an accurate table search, each column in the routing table must be able to transmit equal amount of light intensity, and the added results on the one-dimensional thresholding device TH must be sufficiently different so that the TH can make a correct decision as to which pixel is to be illuminated. This pixel indicates the result of the parallel table search.

The cylindrical optical system behind TH, also not shown, expands the light transmitted through the pixel in column direction to illuminate a corresponding column in the control signal table CST. The con-

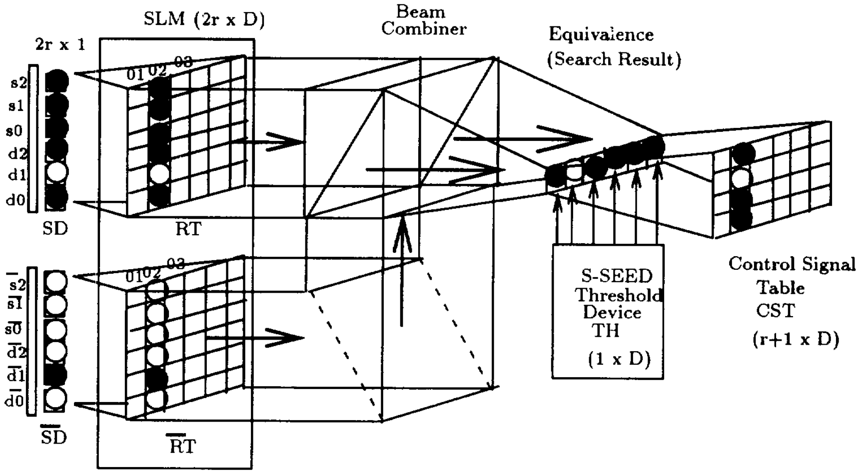

Figure 7: Optical Routing Unit

trol pattern will consist of an array of $r + 1$ pixels (4 in this case), r corresponding to the possible outgoing wavelengths and one to indicate that the packet has reached its destination (i.e. receive the packet). The selected column is imaged onto the output plane to provide an optical binary control pattern for the Optical Switching Unit (OSU).

3.3 Optical Switching Uint

The control pattern is used in the Optical Switching Unit (OSU), shown in Figure 8, to configure the optical switches to route the packet onto the proper wavelength or receive the packet (receiver omitted from figure). The OSU also can resolve conflicts by synchronization with the WSU and delay lines. The delay lines can be adjusted such that the time needed for the ORU to process the packet header of a new incoming packet and generate a control signal for that packet, is larger than the time needed for OSU to configure the switch and transmit the previous packet. Then, when the control signal for the new packet is sent to the OSU, the previous packet will have completed transmission and no conflict occurs.

With no conflicts, the control signal generated by ORU is sent to a shift register which is synchronized with the WSU. This selector passes the control signal to set one of r S-SEED arrays (in this example $r = 3$). Each of these arrays corresponds to a switch for one of the incoming packets and only the array corresponding to the packet for which the control signal was generated is set. Each array is of size $r + 1$ and has input from one of the delay lines. According to the incident control signal, one SEED in the array is transparent, passing the optical packet signal from the delay line to the appropriate wavelength converter, and the others are opaque. Thus, the packet is transmitted on the proper wavelength as indicated by the routing table in the ONIU. All routed packets are then wavelength division multiplexed (MUX) and transmitted to the next ONIU in the ring.

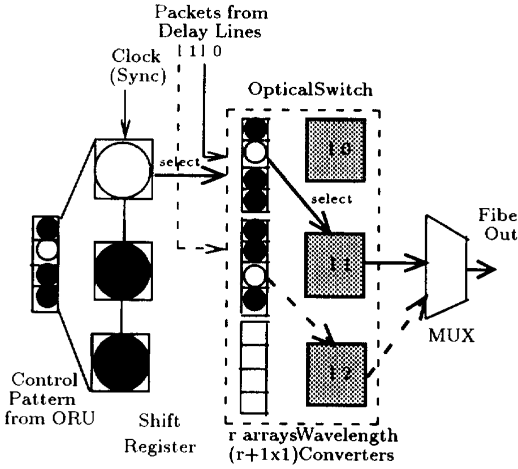

Figure 8: Optical Switching Unit

Figure 8 shows an example where the current packet being routed entered the ONIU on wavelength λ_0 and is being transmitted on wavelength λ_1 (solid lines). It also shows one previous packet being routed (dashed lines) which entered on λ_1 and is being transmitted on wavelength λ_2.

4 Conclusions and Future Work

We have presented an optical Hypercube-Based Multichannel Ring Network on which the hypercube routing functions are implemented using wavelength division multiplexed channels on a single fiber. The number of required channels is shown to be less than that needed in previously proposed multichannel systems. Hypercube routing is achieved using the d-routing scheme, a variation of e-routing which exploits the multiple shared channels and physical ring topology to distribute network traffic more evenly and reduce the size of the routing tables needed.

We have also designed an Optical Network Interface Unit (ONIU) to perform high-speed packet switching using a static routing scheme. The ONIU takes advantage of the high-bandwidth, parallelism, multi-dimensional capability, and high storage density offered by optics. The most time consuming operations in communication networks, such as switching and routing, are performed in optical domain in the proposed system. Thus ONIU does not suffer from the optical/electrical conversion bottlenecks present in hybrid switching systems.

The proposed ONIU is based on current SEED device technology. The high speed (less than 1 ps) and low energy (on the order of fJ) of the self electro-optical device (SEED) [10] has drawn a great attention because of its inherent advantages suited for optical computing. Also, since SEED fabrication is based on molecular beam epitaxy (MBE), the number of SEED elements on a chip can be, in theory, very large. Hence, the proposed system can be implemented using current optical technology and is also easily scalable to build large networks as the technology improves.

Future work for this research includes a performance analysis of ONIU in the Hypercube-Based Multichannel Ring Network. Furthermore, we are considering implementation of the main functions of ONIU, particularly the parallel routing table search using programmable SLM, in our optics laboratory. We are also investigating the use of ONIU for circuit switching as well as other all-optical hypercube-based topologies.

References

[1] A.S. Acampora and M.J. Karol, "An Overview of Lightwave Packet Networks," *IEEE Network*, pp. 29-40, Jan. 1989.

[2] A.S. Acampora, "A Multichannel Multihop Local Lightwave Network," *GLOBECOM '87 Conf. Rec.*, pp. 1495-1467, Nov. 1987.

[3] M.J. Karol, "Optical Interconnection Using ShuffleNet Multihop Networks in Multi-Connected Ring Topologies," *Proc. ACM SIGCOMM '88 Symposium*, pp. 25-34, Aug. 1988.

[4] P. Dowd, "High Performance Interprocessor Communication Through Optical Wavelength Division Multiple Access Channels," *18th Int. Symp. on Computer Arch.*, May 1991.

[5] F.Reichmeyer, S.Hariri and K.Jabbour, "Hypercube-Based Local Lightwave Networks," *35th Midwest Symposium on Circuits and Systems*, paper Tp2,6, Washington, D.C., Aug. 1992.

[6] Wang Song, Salim Hariri and Alok Choudhary, "An Optical Interface Message Processor for Fiber Communication Networks," *Optics Communications* 91 (1992), pp. 304-311.

[7] H. Scott Hinton, "Architectural Considerations for Photonic Switching Networks," *IEEE Journal on Selected Areas in Communications* Vol. 6, No. 7, August 1988, pp. 1209-1226.

[8] K.Hwang and F.A.Briggs, *Computer Architecture and Parallel Processing*. New York, NY: McGraw-Hill Book Company, 1984.

[9] F. B. McCormick et al., "All-optical Shift Register Using Symmetric Self Electro-Optical Effect Devices," Topical Meeting on Photonic Switching, paper Thc5, Washington DC, 1989.

[10] D. Miller, "Optoelectronic Applications of Quantum Wells," *Optics and Electronics News*, February 1990.