

An Optical Interconnect Model for k-ary
n-cube Wormhole Networks

Mongkol Raksapatcharawong and
Timothy Mark Pinkston

CENG Technical Report 95-22

Department of Electrical Engineering - Systems
University of Southern
Los Angeles, California 90089-2562
(213)740-4482

An Optical Interconnect Model for k -ary n -cube Wormhole Networks

Mongkol Raksapatcharawong

Timothy Mark Pinkston

SMART Interconnects Group
EE-Systems Dept., University of Southern California
Los Angeles, CA 90089-2562
{<http://www.usc.edu/dept/ceng/pinkston/SMART.html>}

Abstract

This paper presents an optical interconnect model for k -ary n -cube network topologies based on free-space analysis. This model integrates relevant parameters inherent to optics with traditional network parameters to make it meaningful for performance evaluation of optical network designs. We apply this model to a free-space diffractive-reflective optical interconnect design and compare our results with electronic-based networks.

Keywords: free-space optical interconnects, k -ary n -cube networks, optical interconnect model, performance evaluation, wormhole switching.

I. Introduction

Computer systems capable of teraflop computation are becoming a near-term reality with the emergence of high-speed, powerful processors that can run cooperatively in a *multiprocessor system*. The underlying fabric that supports communication among processors in the system is the *interconnection network*. It must be able to supply the required bandwidth for it not to be a bottleneck to system performance. Therefore, the design of low latency interconnection networks is important and is an active research topic in parallel computer architecture [1-3].

Realizing the bandwidth limitations imposed by electrical interconnects, network designers are exploring alternatives, both architectural and technological, that can overcome these limitations. Optics is one such alternative. Optical technology has successfully penetrated the wide-area machine-to-machine interconnect regime. Because of its inherent parallelism, high interconnectivity, and large bandwidth, optics may also be applicable to chip-to-chip and board-to-board level interconnects [4-10]. Studies comparing optics with electronics at these interconnect levels conclude that optics may have a speed/power advantage [5,7,8]. However, few analytical models exist which characterize performance of optically-based multiprocessor networks.

In this paper, an analytical model is developed for three-dimensional free-space optical k -ary n -cube interconnection networks that extends Dally's planar VLSI model [11]. We then apply our model to a diffractive-reflective optical interconnect (DROI) design [12]. The rest of this paper is organized as follows. Section II gives some background on k -ary n -cube topologies, wormhole switching, and free-space optical interconnects. Section III presents our analytical model for optical interconnects. Section IV discusses the performance of optical and electrical k -ary n -cube networks when applying our model. Section V discusses some practical considerations in free-space optical systems. We give some conclusions in Section VI.


II. Background

A. k -ary n -cube topologies


The k -ary n -cube class of networks has a number of favorable characteristics as given by Table 1 and, hence, are among the more popular. They employ static, direct point-to-point connections between nodes and support locality of communication to reduce delay of messages in the network. Topologies in this class of networks have channels which span n dimensions and have k nodes connected in each dimension. Examples of a 2-ary 3-cube network with no wrap-around (mesh-connected) and a 3-ary 2-cube network with wrap-around (torus-connected) are shown in Figure 1 (a) and (b).

Table 1. k -ary n -cube network characteristics (for unidirectional links).

- | |
|--|
| <ul style="list-style-type: none"> • node and edge symmetric • regular topology • connectivity = $2n$ • maximum degree = n • maximum diameter = $n(k-1)$ • average distance $D_{avg} = n(k-1)/2$
(for uniform message distribution) • channels = $nN = nk^n$ • bisection width (channels) = $2N/k = 2nk^{n-1}$ • nodes in system = $N = k^n$ |
|--|



(a) 2-ary 3-cube network.



(b) 3-ary 2-cube network.

Figure 1. Examples of k -ary n -cube networks.

B. Wormhole Switching

The switching technique also largely influences the delay of messages in the network. Our model assumes wormhole switching [11] which pipelines the transfer of flits¹ along the path from source to destination. Once the header flit of a message (which contains all the relevant routing information) is received by a node, it is routed to an appropriate output channel. If that channel is free, the header is transferred to the next node; all other flits follow sequentially. If the required channel is busy, all flits are blocked behind the header

¹ A flit or flow control unit is the unit of message transfer on which flow control is performed.

and wait until the channel becomes available. Therefore, the latency resulting from wormhole switching is expressed simply as

$$T_{lat} = T_c \left(D + L_F \cdot \frac{F}{W} \right) + T_{contention} , \quad (1)$$

where T_c is the channel cycle time for transceiving and routing flits, D is the number of network hops required from source node to destination node, L_F is the message length in flits, F is the flit size in bits/flits, and W is the channel width in bits. The congestion along the path from source to destination due to messages contending for the same channel is parameterized by the $T_{contention}$ variable. The channel cycle time, T_c , takes into account all the latencies that are incurred internal and external to the network router. Note that by pipelining logic functions in the network router, the external propagation delay (i.e., signal propagation time in the interconnect medium and signal conversion/regeneration, if applicable) of signals can become the critical path which determines the channel cycle time.

C. Free-Space Optical Interconnects

Conceptually, free-space optical interconnects can be comprised of a transmitter plane, a receiver plane, and an optical imaging system in between. Optical beams are transmitted by transmitters (e.g., light sources or modulators) in the transmitter plane, deflected and/or split by the optical imaging system (e.g., holograms, lenses, etc.), and detected by detectors (e.g., photodiodes, modulators etc.) at the receiver plane. Free-space optical interconnects make use of the third dimension to route signals. This is a major distinction between optical and conventional electrical interconnects.

Due to two-dimensional (planar) VLSI implementation, electrical interconnects have traditionally been evaluated in terms of *bisection width*—which is the number of wires crossing an imaginary plane that divides the system into two equal halves. This notion is no longer valid in developing expressions for network latency in optical interconnects because of optics' three-dimensional routing characteristics which surpass those of electronics. In optics, interconnections are established over a volume. This directly influences the network latency due to a wider channel width for each topology as confirmed by this study. We therefore introduce the notion of *connection capacity* to evaluate optical interconnects as opposed to bisection width.

III. The Model

An analytical model for k -ary n -cube optical networks with wormhole switching is developed in this section. This model is an extension of Dally's analysis applied to optical interconnects [11]. Our analysis is primarily based on the notion of connection capacity as opposed to bisection width. Any free-space optical interconnect system can be described in terms of its connection capacity. Below, we show that the connection capacity has a significant bearing on the types of topologies that are efficiently supported.

A. Connection-Efficient Topologies

Let C be the connection capacity of an optical imaging system (C is constant). The number and width of channels supported for various topological configurations can be described in terms of this connection capacity:

$$C = lW = nNW = nk^n W, \quad (2)$$

where the number of unidirectional channels or links required by a k -ary n -cube network is given in Table 1.

Comparisons between various topological configurations are more comprehensible when we normalize connection capacity to that of the hypercube (binary n -cube) topology with unity channel width. Normalized connection capacity results in $C = N \log N$. The channel width $W(k, n)$ of a k -ary n -cube with normalized connection capacity is therefore given by²

$$W(k, n) = \frac{N \log N}{nN} = \frac{\log N}{n} = \log k. \quad (3)$$

This expression for channel width with normalized connection capacity is different from that for normalized bisection width derived by Dally [11], where a *constant bisection width* was assumed. Under that assumption, channel width was shown to grow linearly with increasing k (i.e., $W(k, n) = k/2$). Intuitively, when k increases the dimension, n , decreases accordingly, increasing the number of links and, thus, decreasing the channel width, W . Clearly, our model shows that optically implemented topologies are less sensitive to the ary-ness, k , as the logarithmic function ($\log k$) changes less rapidly than the linear function ($k/2$). Hence, the advantages expected from lower dimension networks (namely, wider channels) should not be as pronounced for optical (connection capacity limited) networks as they are for electrical (bisection bandwidth limited) networks. Next, we show how this conclusion impacts the latency of optical interconnects according to our model.

Latency is the time required to deliver a message from source to destination. The average latency for a k -ary n -cube network can be found as follows. If we randomly select, with equal probability, the source and destination nodes P_s and P_d , the average number of hops between them is given by

$$D = \left(\frac{k-1}{2} \right) \cdot n, \quad (4)$$

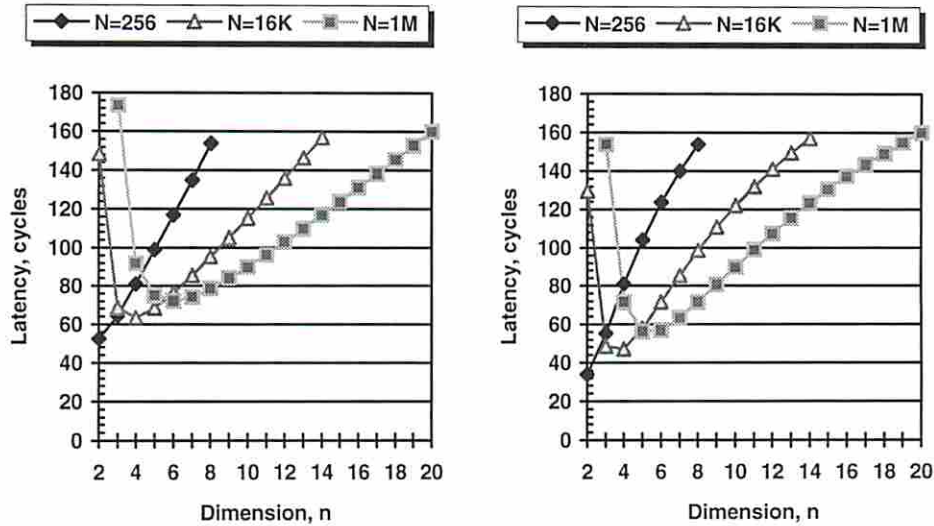
where $(k-1)/2$ is the average number of hops the message travels in each dimension given that links are unidirectional, and n is the number of dimensions.

Substituting this expression and the channel width from Eq.[3] into Eq.[1], we express the average latency of an optical k -ary n -cube network as follows:

$$T_{lat} = T_C \left(\left(\frac{k-1}{2} \right) \cdot n + L_F \cdot \frac{F}{\log k} \right) + T_{contention}. \quad (5)$$

² Throughout this paper, $\log x$ stands for $\log_2 x$.

The second term in Eq.[5] is the only difference between our optical model and Dally's electrical model; it is less sensitive to the ary-ness of the network than the electrical model. Therefore, we should expect to see latency increase more slowly with dimension as compared to Dally's model. Latency characteristics for both models are illustrated in Figure 2.



(a) Optical model. (b) Electrical model (Dally, 1990).

Figure 2. Latency versus dimension with unit channel cycle time ($T_{contention}$ excluded).

Figure 2 depicts the average network latency as a function of dimension for k -ary n -cube networks with $N=256$, 16K and 1M nodes. Unit channel cycle time is assumed, and the message length $L=F \cdot L_F$ (in bits) is assumed to be 150 bits (flit size and channel width are assumed to be equal henceforth). Thus, the above figure represents the latency for *constant delay* of both optical and electrical signals regardless of the physical distance between source and destination nodes. It should be noted that Figure 2 does not intend to compare latency between optics and electronics; rather, it shows how dimension affects latency for each design space (optics or electronics).

For each curve, the rightmost data point corresponds to a hypercube and the leftmost data point corresponds to a 2-D torus. In low-dimensional networks, messages travel a greater number of hops. Latency is dominated by this hop distance even with wormhole routing (network congestion would further degrade performance). In contrast, messages suffer increased transfer time (in flits) between nodes for high-dimensional networks due to the smaller channel width offered by the topology. Here, latency is dominated by message length. Hence, our results agree with [11] in that low-dimensional networks outperform high-dimensional networks in terms of latency for both design spaces. However, lower dimensional networks are not as advantageous for optics, especially given that it is less difficult to implement higher dimensional networks with free-space optics than with wire-limited electronics.

B. The Channel Cycle Time (T_C)

The previous analysis assumed constant channel cycle time. In what follows, the model is developed in more detail to include the effects of optical signal delay assuming that the path external to the router defines channel cycle time (e.g., $T_C = \max[\text{external delay, internal delay}]$). The time to convert and propagate an optical signal between a pair of nodes is given by the following:

$$T_C = T_{E/O} + T_{O/E} + T_{PROP}. \quad (6)$$

The first term, $T_{E/O}$, is the electro-optical delay time for the optical source (or modulator) circuit. The second term, $T_{O/E}$, is the opto-electronic delay time of the receiver circuit. The last term, T_{PROP} , is the light propagation time which is approximately 1 ns per foot in free-space. Our optical signal delay model is depicted in Figure 3.

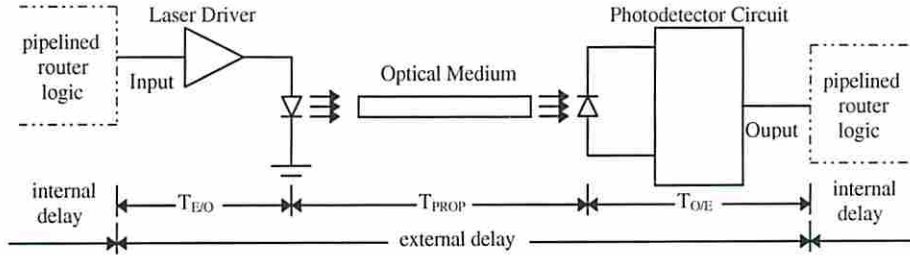


Figure 3. The optical signal equivalent propagation path.




Figure 4. Schematic of a transmitter circuit [13].

We model transmitter delay assuming the transmitter circuit shown in Figure 4. The output of the logic gate is amplified and drives a transistor current buffer which supplies sufficient current to drive the *Vertical-Cavity Surface-Emitting Laser* (VCSEL). When the current flow through the VCSEL is above its threshold current, the VCSEL starts emitting light, resulting in an electrical to optical modulation. The electrical to optical conversion delay is described by [14]

$$T_{E/O} = (2 \cdot N/P + 1)R_n(C_{out} + C_{in}) + T_{laser}, \quad (7)$$

where C_{in} is the input capacitance of the output and current driver transistors, and C_{out} is the output capacitance of the amplifier transistor, N/P is the ratio of n-MOS to p-MOS, R_n is the n-MOS linear resistance, and T_{laser} is the laser response time.

Commercially available VCSELs require 7 mA of current for 1 mW output power. Thus, the current driver transistor must be sized to supply that amount of current. Assuming a 0.8- μm CMOS process and an N/P ratio of 3, we get $C_{out} = 3.08$ fF, $C_{in} = 88.5$ fF for current driver transistor, and $R_n \sim 882 \Omega$ for amplifier transistor. This results in a laser

delay, T_{laser} , of 0.1 ns for 1 mW output. These values yield ~ 0.67 ns electrical to optical conversion delay according to Eq.[7].




Figure 5. Schematic diagram of a P-I-N photodetector circuit [5].

We model receiver delay assuming the receiver circuit shown in Figure 5. The receiver consists of PIN diode and an output driver. An expression for the delay in charging up voltage at the receiver is shown to be [5]

$$T_{OIE} = \frac{V}{S\eta P_{laser}} (C_{PD} + C_{in}) F_{fan-out}, \quad (8)$$

where S is the P-I-N detector sensitivity, P_{laser} is the optical power emitted by the VCSEL, V is the power supply voltage, C_{PD} is the P-I-N detector capacitance, C_{in} is the input capacitance of the output driver, $F_{fan-out}$ is the fan-out (which is one in this study), and η is the optical link efficiency (which includes that of the hologram and the microlens).


Assuming the same CMOS process and an input capacitance of a unit sized inverter of 5.31 fF, we get $C_{PD} = 53$ fF for a detector area of $15 \mu\text{m} \times 15 \mu\text{m}$. Its sensitivity is 0.5 A/W at 15V reverse-bias. In a DROI design with a 1 mW VCSEL and an optical link efficiency of 63% (81% hologram efficiency for 4-level diffractive optical element (DOE) [15] and 99.5% microlens efficiency), we get an optical to electrical conversion delay of 0.9 ns (which can be reduced to 0.3 ns if a 3-mW VCSEL is used). This number is not an under-estimate as detectors that operate beyond 700 Mb/s with 800 μW optical power are reported in the literature [16].

The last major component of channel cycle time is the propagation delay. This delay is dependent on the medium and its length. The most efficient way to implement a network topology in a volume (where nodes are to reside in a plane) is to map the connections as symmetrically as possible so as to minimize connection length. Previous attempts to map optical k-ary n-cube topologies in a volume did not consider wrap around connections (see Drabik [17]). We define the *maximum connection path* to be the longest connection between two nodes in the system. Figure 6 shows a suggested layout of nodes and the mapping of connections in a volume (3-D) for various 4-ary n -cube topologies.



(a) Layout and mapping of 1-D network.

(b) Layout and mapping of 2-D network.



(c) Layout and mapping of 3-D network.

Figure 6. Embedding of 4-ary n -cubes in a volume (nodes in 2-D plane) for $n=1,2,3$. Only the connections of nodes along edges are shown for clarity. Moreover, the mirror plane shown here does not correspond to the real system where it would be above the transmitter-receiver plane.

With this layout, the maximum connection path, R_{max} , is given by

$$R_{max} = \begin{cases} \frac{2pk^{(n/2-1)}}{\sin \theta}, k = 2^i, i > 1 \\ \frac{2pk^{(n/2-2)}}{\sin \theta}, k = 2^i, i = 1 \end{cases}, \quad (9)$$

where $p = \sqrt{A/N}$ is the minimum connection length (lateral distance) between adjacent nodes, N is the number of nodes, A is the square area of the node plane, θ is the maximum hologram deflection angle, n and k are dimension and ary-ness of the network. It should be noted that Eq.[9] applies for any configuration of n and k (which are integers) that fit perfectly in a square area. Therefore, the light propagation time is

$$T_{PROP} = \frac{R_{max}}{c}, \quad (10)$$

where c is the speed of light.

We can now estimate the channel cycle time, T_C . According to the values calculated previously, the channel cycle time of an optical interconnect over a distance of $R_{max} = 1$ foot is 2.57 ns.

C. Network Latency with Linear Optical Signal Delay

The latency figures shown in Section III-A do not reflect the more realistic situation where the channel cycle time in optical networks is not constant but depends on interconnect distance. Assuming that the efficiency of a free-space optical system does not depend on distance, $T_{E/O}$ and $T_{O/E}$ in Eq.[6] remain constant. In this case, *linear optical signal delay* results where $T_C \propto T_{PROP} \propto R_{max}$.

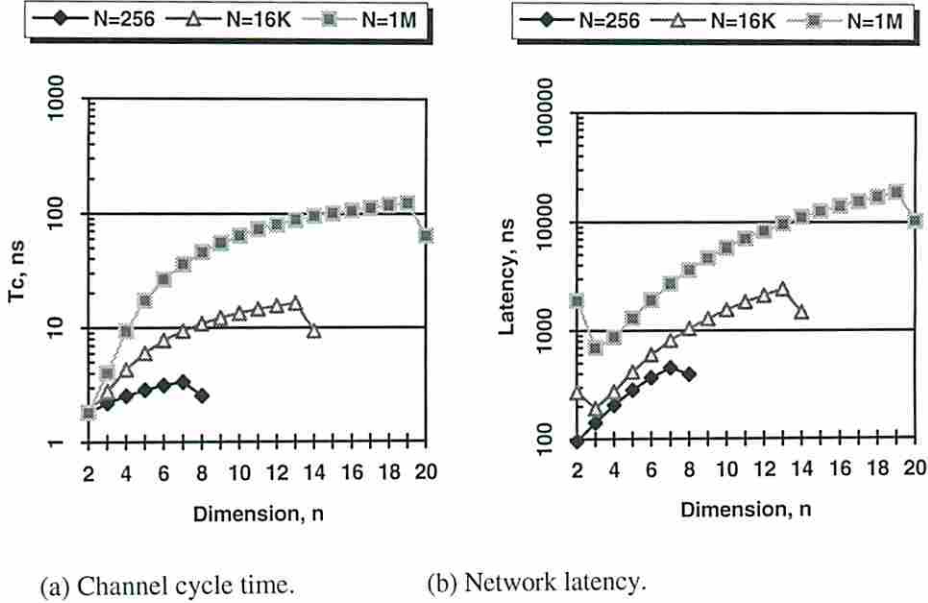


Figure 7. The effect of maximum connection path on channel cycle time and network latency with linear optical signal delay ($T_{contention}$ excluded).

Figure 7 (a) and (b) depict channel cycle time and network latency for systems with $N=256$, $16K$, and $1M^3$ nodes when normalized connection capacity and linear optical signal delay is assumed. (Message length is 150 bits and the minimum connection length is assumed to be 1.5 cm.) As expected, channel cycle time increases with dimension.⁴ Low-dimensional networks feature shorter maximum connection path, thus yielding lower channel cycle time. This is seen simply if we rewrite Eq.[9] as $R_{max} \propto 1/\sqrt[k]{N^2}$ for $k=2$, and $R_{max} \propto 1/\sqrt[k]{N}$ for $k > 2$.

Clearly, the channel cycle time of high-dimensional networks is comparatively higher. The smaller channel width further accentuates the latency difference between low- and high-dimensional networks. Hence, with the linear delay assumption, low-dimensional networks still outperform high-dimensional networks in terms of network latency for a broad range of system sizes.

³ For the sake of comparison, we allow k to be non-integer values.

⁴ There is a special case when $k=2$ (hypercube network) where the channel cycle time tends to be lower. This discrepancy is due to the topologies that do not actually exist in our mapping scheme (i.e., k is not an integer).

D. Connection Capacity (C)

Channel width of each topology in k -ary n -cube optical networks is determined by connection capacity. In general, the connection capacity for an optical imaging system is expressed as

$$C = \frac{A_{system}}{A_{spot}}, \quad (11)$$

where A_{system} is the area over which interconnects can be established and A_{spot} is the maximum light beam area along the propagation path. Assuming diffractive-reflective optical interconnects (DROI) and Gaussian beam propagation [Appendix A], these two parameters are shown to be functions of other system parameters [Appendix B]:

$$\begin{aligned} A_{system} &= F(\theta, h, p, n, k) \\ A_{spot} &= F(f, w, \lambda, \theta, h) \end{aligned} \quad (12)$$

where θ is the hologram deflection angle, h is the separation between mirror and microlens planes, f is the microlens focal length, w is the transmitted beam radius, λ is the wavelength, p is the minimum connection length, n is the dimension, and k is the ary-ness. The hologram deflection angle itself is also a function of other system parameters [Appendix B]:

$$\theta = F(\lambda, n_x, L_b, w_f), \quad (13)$$

where n_x is the index of refraction of the material through which optical signals propagate, L_b is the number of hologram levels, and w_f is the minimum feature size of each hologram. Figure 8 illustrates the DROI geometry. Shown is one optical signal connection path.




Figure 8. A DROI geometry.

In our study, we maintain A_{system} constant and allow h to vary according to the topology (no volume constraint). We also make the simplifying assumption that the spot size area, A_{spot} , is equal to the microlens area, M_D^2 . Therefore, the connection capacity for our assumed DROI simplifies to

$$C = \frac{A_{system}}{2M_D^2}. \quad (14)$$

The factor of two in the denominator takes into account the fact that both transmitters and receivers are in the same plane. For example, assuming the DROI optical imaging system

supports interconnection of nodes over an area of $A_{system}=64 \text{ cm}^2$ and the lens diameter of each interconnection is $M_D=125 \text{ }\mu\text{m}$, the connection capacity is 204,800 connections for all topologies.

IV. Application of the Model: Optics vs Electronics

It is interesting to determine whether optics performs better than electronics and, if so, by how much. Therefore, we compare the latency given by our model with that given by [11] for k -ary n -cube networks. Our optical model has as its constraint the connection capacity whereas Dally's electrical model has bisection width as its constraint. These constraints, although different, are actually related by the fact that they are used to determine channel width of the various topologies.

A. Electrical Interconnect Delay Model

The latency model of an electrical interconnect is based on distributed RC effects of a transmission line using microstrip conductor type with no transmission line effect [18]. The channel cycle time in electrical interconnects is given by

$$T_{C-elec} = T_{PROP-elec} + T_{RC} . \quad (15)$$

Here, T_{PROP} is the propagation delay in an electrical medium which is 0.148 ns/in [18]. The RC delay, T_{RC} , takes into account the distributed RC effect of the transmission line [14] and delays associated with driver and receiver circuits as shown in Figure 9.




Figure 9. Simple model for electrical interconnect delay.

Here, C_l is the input capacitance of a receiver, C_o is the output inverter capacitance, C_b is the bonding pad capacitance, τ is the signal delay on a transmission line, C_r and R_r are the total lumped capacitance and resistance of the transmission line. The signal delay on the transmission line is given by [14]

$$\tau = \frac{rcl^2}{2} , \quad (16)$$

where r and c is the unit length resistance and capacitance which are 45.4 m Ω /in and 1.0 pF/in for 5-mil-wide, 5-mil-apart, and 2.7-mil-thick conductor [19], and l is the transmission line length.

The total RC delay, T_{RC} , is expressed as follows:

$$T_{RC} = \left(\frac{C_l + C_b + C_o}{V} \right) \left[\frac{1}{\beta_n} + \frac{1}{\beta_p} \right] + \tau + R_t(C_l + C_b) . \quad (17)$$

The driver is assumed to provide an output current of 7 mA. Thus, the output inverter capacitance based on 0.8- μm CMOS process is 64 fF. The input capacitance of the receiver is 3.54 fF. The n- and p-MOS transistor gains, β_n and β_p are equal to 2213.5 $\mu\text{A}/\text{V}^2$. The bonding pad capacitance is 0.4 pF for the $100\mu\text{m}^2$ pad area. Given a maximum connection path, $R_{max-elec}$, we can find the channel cycle time of an electrical interconnect by using Eqs.[15-17]. For 1-foot interconnection length, the channel cycle time of the above interconnects is 4.03 ns (2.25 ns for RC delay and 1.78 ns for propagation delay).

The channel cycle time for both optical and electrical interconnects vary differently with the maximum connection path R_{max} . We plot the relation between the two according to the parameters we have been assuming to observe their difference in Figure 10.

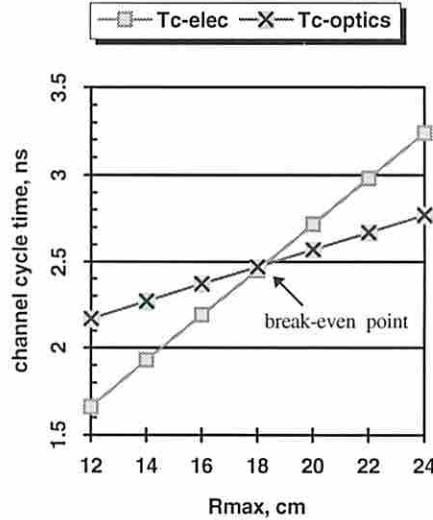


Figure 10. T_C and break-even point.

We observe that, under the realistic parameters assumed, optics gives a channel cycle time that grows about 2.6 times slower than electronics. This means optics is better for longer R_{max} . The *break-even point* where both schemes yield comparable channel cycle time indicates at what point optics is better. Figure 10 shows the break-even point is around 18 cm. Since both signal conversion times in optics depend entirely on optoelectronic and micro-optics technologies, we envision that this break-even point will become smaller with further development, resulting in shorter distance for optical interconnects to be superior. In contrast, the major contribution to channel cycle time in electronics is the line capacitance. It is worth mentioning that free-space optics is inherently more compact than electronics, so it generally operates at a smaller R_{max} and, thus, at a smaller channel cycle time. For instance, given a system size of 144 in^2 for electronics and 144 cm^2 for optics

(assuming a deflection angle of 24°), we get $R_{max-elec}=18.53$ cm and $R_{max-optics}=7.41$ cm for a 2-D torus network.

B. Channel Width

Substituting Eq.[11] into Eq.[2], we express channel width for optical networks as a function of topological and implementation parameters:

$$W_{optics}(k, n) = \left(\frac{A}{2M_D^2 \cdot N \log N} \right) \cdot \log k. \quad (18)$$

Applying the bisection width notion to electrical networks, we express its channel width as a function of topological and implementation parameters:

$$W_{elec}(k, n) = \left(\frac{L\sqrt{A}}{NT_w} \right) \cdot \frac{k}{2}, \quad (19)$$

where A is the printed circuit board (PCB) area (or area of the microlens plane in optical networks), N is the system size, T_w is the electrical wire pitch, L is the number of PCB layers that can be routed in same direction, and M_D is the microlens diameter.

C. Latency Comparison

We now compare latency between electrical interconnects with aggressive PCB technology and optical interconnects with available optoelectronic and micro-optics technology. Parameters assumed are listed in Table 2 and 3 for electrical and optical interconnects, respectively.

Table 2. Parameters for electrical system.

Chip area	1 in ²
PCB size	12 in x 12 in
# of layers	20
min. connection length (p)	1.5 in

Table 3. Parameters for optical system.

laser wavelength (λ)	850 nm
VCSEL beam radius	5 μm
VCSEL output power	1 mW
P-I-N detector size	15 μm x 15 μm
microlens diameter	125 μm
chip area	1 cm^2
interconnection area	12 cm x 12 cm
usable microlens area (A)	64 cm^2
min. connection path (p)	1.5 cm
max. deflection angle (θ_{max})	$\sim 24^\circ$

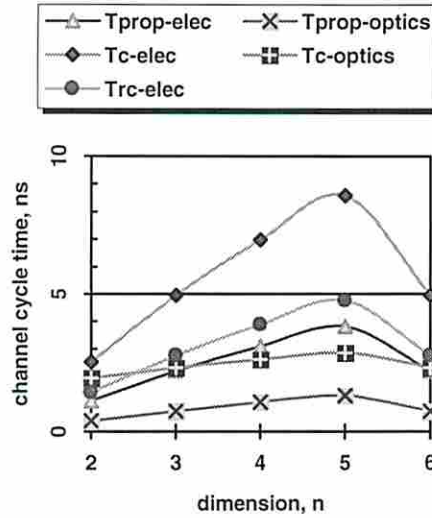
The electrical interconnects are implemented using a 20-layer PCB (12 in x 12 in) in which 10 layers can be used to route signals in the same direction. We assume the router node die size to be that of the CHAOS Router chip which is 10 mm x 10 mm [20]. Due to die packaging, each node occupies a square area of 1 in². All nodes are placed 0.5 in apart, thus, the minimum connection length is 1.5 in. The conductor and spacing are 10 mils. This number is reported and implemented by Hewlett Packard [21]. Substituting these values yields a bisection width for the electrical interconnects of 12,000 connections. Other parameters for latency calculations are assumed to be the same as in Section IV-A.

Our optical interconnects are assumed to be implemented with a 12 cm x 12 cm transmitter-receiver plane. Each node occupies a square area of 1 cm² and is separated by 0.5 cm from its neighbors for a minimum connection path of 1.5 cm (die only). Recent studies confirm that packaging of multiprocessor free-space optical interconnects at this level of compaction is feasible [22-23]. The VCSEL and P-I-N detector arrays are integrated on top of the CMOS circuits via flip chip bonding [24]. Therefore, only 64 cm² is available for microlens and hologram fabrication. Minimum feature sizes of 0.35- μm for holograms⁵ and 0.8- μm for CMOS circuits are assumed. Every plane is packaged together within glass with a refraction index of 1.5. Under the Gaussian-beam propagation assumption, we find that microlens with 125 μm diameter is sufficient to collect light with 99.5% efficiency and, hence, resulting in the connection capacity of 204,800 connections. All other parameters for $T_{E/O}$ and $T_{O/E}$ are consistent with Section III-B.


In comparing both types of interconnects, we can immediately see the great disparity in connectivity. Optics provides ~ 17 times more connectivity. The maximum volume needed to sustain this connectivity is a modest $\sim 980 \text{ cm}^3$ (12 cm x 12 cm x 6.78 cm). (We chose a 64-node system due to the limited PCB area and transmitter-receiver plane area for the nodes.) The channel cycle time and network latency of the 64-node system with a message length $L=1024$ bits (or 128 byte packets) are plotted in Figure 11. We vary the

⁵ Note that 0.35- μm technology has been employed for AMD microprocessor fabrication [25].

dimension n and observe the channel cycle time and network latency as given by both models. We allow k to be non-integer values to see results trending. The results corresponding to non-integer values of k were found based on the curve-fitting method.



(a) T_c and its components (electronics and optics).



(b) Channel width.

(c) Network latency.

Figure 11. Latency and channel width of the 64-node system.

We show channel cycle time and its components for both optical and electrical interconnects in Figure 11 (a). Every component in T_{c-elec} grows with dimension whereas only the propagation time does so in $T_{c-optics}$. This makes channel cycle time in optics grow much slower. Although routing in the third dimension increases propagation distance, the effect is negligible.

Channel width for both schemes is depicted in Figure 11 (b) assuming 10% of the signal lines are used for data (due to practical considerations such as power, ground, and control

lines). Our results confirm that optical interconnects can be implemented more compactly with lower latency for all topologies shown here. Optics' wider channel width makes network latency less dependent on message length even for higher dimensions. Together with wormhole switching, which makes hop distance even less of an impact on network latency, optical interconnects are closer to achieving constant minimal network latency for various k -ary n -cube configurations as shown in Figure 11 (c).

V. Other Considerations

Two ~~well-known~~ ^{other important} considerations in implementing free-space optical ~~systems~~ ^{interconnects} are power dissipation and packaging tolerance. We discuss both issues from a performance perspective in how they affect the network latency of optical interconnects. In our DROI study, power dissipation and cooling capability for current technology put limits on channel width. Misalignment in system packaging leads to either larger transmitter or receiver microlenses. Given that power dissipation, cooling capability, and interconnection area are known, we can determine the channel width, the network latency, and the packaging tolerance of the system.

A. Power Dissipation

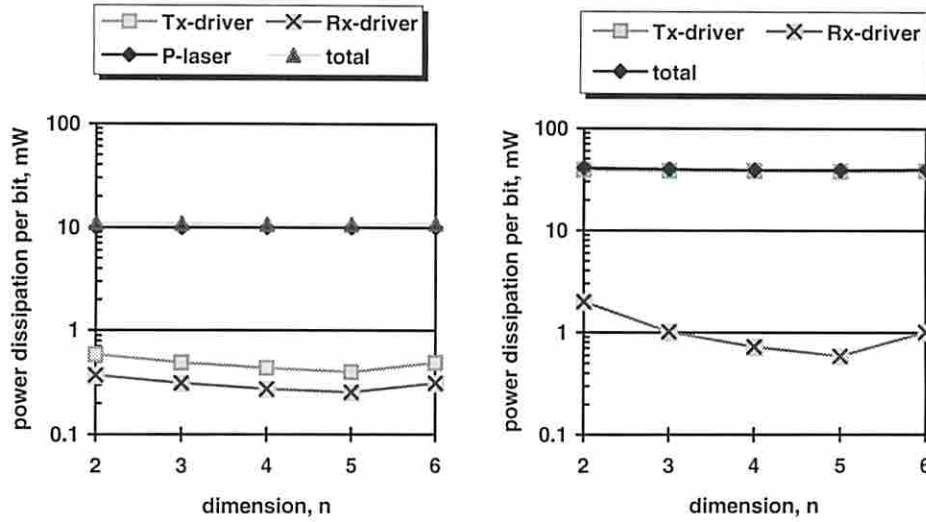
Power dissipated as heat generated by VCSEL is

$$P_{laser} = P_{th} + P_o \left(\frac{1 - \eta_{VCSEL}}{\eta_{VCSEL}} \right), \quad (20)$$

where P_{th} is the threshold power, P_o is the optical output power, and η_{VCSEL} is the slope efficiency (mW/mA). Similarly, the power dissipation of an electronics circuit is given by

$$P_{elec} = \frac{1}{2T_c} CV^2, \quad (21)$$

where C is the total load capacitance, V is the supply voltage, and T_c is the channel cycle time. Therefore, the overall power dissipation per optical channel takes into account the heat generated by VCSEL, laser driver, and receiver circuit. Each electrical channel includes only the heat generated by transmitter (or line driver) and receiver circuits.

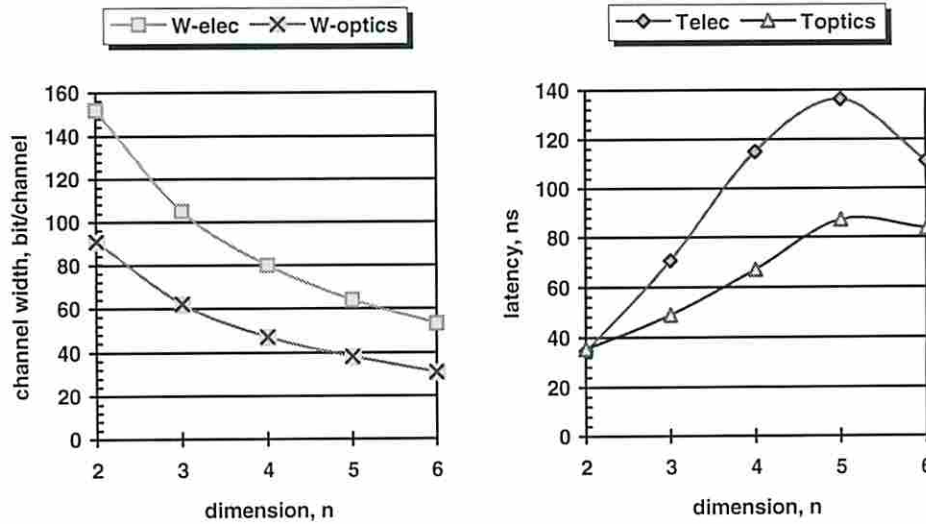


(a) Optical interconnects.

(b) Electrical interconnects.

Figure 12. Power dissipation of the 64-node optical and electrical interconnects.

Figure 12 illustrates power dissipation per bit and its components calculated using Eq.[20-21] and parameters from Section III-B and IV-A are assumed. The threshold power, P_{th} , for the 1-mW output VCSEL is assumed to be 8.5 mW ($I_{th} = 5$ mA, $V_{th} = 1.7$ V).



(a) Channel width.

(b) Network latency.

Figure 13. Latency and channel width of the 64-node system with 2 W/cm² cooling capability. Only chip areas are taken into account for cooling which are 64 cm² for optics and 400 cm² for electronics.

Electrical interconnect dissipates most power to drive the transmission line while optical interconnect dissipates most power on the VCSEL. This dissipation is so large that it dominates other components. Figure 12 describes that power dissipation per bit for both interconnects virtually independent of topology. It is obvious for optical interconnect to

maintain constant power dissipation according to Eq.[20]. In contrast, Eq.[21] shows that power dissipation in electrical interconnect depends on capacitance (mostly line capacitance) and channel cycle time. Coincidentally, they change at almost the same magnitude and, hence, a virtually constant power dissipation is achieved. Although both electrical and optical interconnects share the same power dissipation trend, electrical interconnect still generates heat about four times larger than optical interconnect. This gap is likely to be larger in a near-future due to the emergence of low-threshold, high-efficiency VCSELs which thereby lower a significant source of power dissipation in optics.

To comprehend the effect of power dissipation on network latency we assume the cooling capability of 2 W/cm^2 of interconnection area. The number of I/O channels is therefore more strictly limited than that imposed by connection capacity or bisection width in the optical or electrical interconnects. Finally, we would expect smaller channel width and larger network latency as a result. We show this effect in Figure 13.

Electronics provides wider channel width due to its larger chip area. Although this is the case, network latency in optics still surpasses that of electronics. The latency difference is more pronounced for higher dimension because of optics' lower channel cycle time. Thus, when power dissipation and cooling capability are being considered, optical interconnect does not yield as much performance advantage as when connection capacity is the sole consideration. However, it must be emphasized again that progress in optoelectronic and micro-optic technology will enhance optics' performance advantage over electronics.

B. Packaging Tolerance

The analyses presented here are based on DROI scheme. They can, however, be applied to other free-space optical systems with only slight modifications. From Section V-A we realize that power dissipation places more stringent limitation on channel width than connection capacity of an optical interconnect. Therefore, assumption in Section V-A still holds in this section and we add the effect from each type of misalignment individually to this limitation. By doing so, we can derive the individual tolerance corresponding to each type of misalignments allowed in DROI packaging.




Figure 14. Three types of misalignment in DROI.

We are interested in three common misalignments exhibited in free-space optical systems. They are shown in Figure 14 for DROI scheme. The first is lateral misalignment, which is the horizontal misalignment between transmitter-receiver plane and microlens-hologram plane. The second one is longitudinal misalignment which is the vertical misalignment

between the microlens-hologram plane and mirror plane. The last is angular misalignment of the mirror plane to horizontal plane. The effect of each type is somewhat similar to each other.

- 1) *Lateral Misalignment*: To maintain the microlens efficiency at 99.5%, the diameter must be four times larger than spot radius at the microlens. Since this misalignment shifts both transmitters and receivers, it is necessary to fabricate larger transmitter and receiver microlenses to maintain this condition. We can express the lateral misalignment as

$$\Delta x = \Delta_{Lat}, \quad (22)$$

where Δx is the lateral shift of the optical beam respect to the microlens-hologram plane and is Δ_{Lat} the lateral misalignment. Referring to the assumption in Section V-A, we can fabricate a maximum of 1,936 microlenses/cm² with diameter of 227 μm each. Compared to the correctly aligned microlens with 125 μm diameter, we can tolerate 102 μm of lateral misalignment.

- 2) *Longitudinal Misalignment*: Misalignment either upward (+ sign) or downward (- sign) with respect to the microlens-hologram plane results in a lateral shift of optical beams at the receiver microlenses and can be written as

$$\Delta x = 2\Delta_{Long} \tan \theta, \quad (23)$$

where Δx is the lateral shift of the optical beams, Δ_{Long} is the longitudinal misalignment, and θ is the maximum hologram deflection angle. The effect of this misalignment is somewhat different from the previous one because it requires only larger receiver microlenses. By keeping transmitter microlenses at 125 μm diameter, the receiver microlenses can be enlarged to 330 μm . From Eq.[23], we can tolerate 230 mm of longitudinal misalignment.

- 3) *Angular Misalignment*: The angular misalignment of the mirror plane results in a lateral shift of optical beams at the receiver microlenses similar to 2) and can be written as

$$\Delta x = 2h\Delta\theta, \quad (24)$$

for small angle approximation where Δx is the lateral shift of the optical beams, h is the mirror and microlense-hologram planes separation, and $\Delta\theta$ is the angular misalignment. For h equals to 6.77 cm (the 64-node hypercube network) and the receiver microlenses with 230 μm diameter, we can tolerate the angular misalignment of 7.75×10^{-4} radian or only 0.044°!

C. Wavelength Variation

Uniformity of optoelectronic devices is not easily achieved, particularly the VCSEL array. Although the uniformity of the VCSELs also includes some other parameters such as threshold current, threshold voltage, its wavelength variation causes the most severe performance deterioration. This issue is analyzed in this section.

Wavelength variation affects two things in DROI. This variation changes the spot radius at both transmitter and receiver microlenses due to Gaussian beam propagation (see Appendix A). It also results in a lateral shift of optical beams with respect to the receiver microlenses. To accommodate both changes, we need to fabricate larger microlenses to maintain their efficiency. Since the variation has little effect on the spot radius, we will address in this paper, only the lateral shift of optical beams.

The hologram deflection angle changes proportionally to the wavelength (see Appendix B) and can be expressed as (for small angle approximation)

$$\Delta\theta = \frac{\Delta\lambda}{n_x T}, \quad (25)$$

where $\Delta\theta$ is the deflection angle sensitivity, $\Delta\lambda$ is the wavelength variation, n_x is the refractive index of material, and T is the hologram period. This deflection angle sensitivity resembles the angular misalignment (for small angle) previously discussed and causes the lateral shift between the optical beams and the receiver microlenses which is

$$\Delta x = \frac{2h\Delta\lambda}{n_x T}. \quad (26)$$

Assume the receiving microlens with 230 μm diameter and the 64-node hypercube network, a very small wavelength variation of 0.8 nm is allowed!

These results clearly show that immature optical interconnects have only a few (but important) obstacles to overcome before they can be widely employed. Optics' (free-space) small feature makes not only heat removal but also packaging much harder than electronics. Together with device nonuniformity such as wavelength variation, they considerably reduce the performance advantage of optical interconnects. This is the reason why free-space optics is not likely to be implemented in a large single volume. Once these issues have been substantially improved, free-space optics will certainly be practical for interconnects.

VI. Discussion and Conclusions

In this paper we present an optical interconnect model for k -ary n -cube networks. The notion of connection capacity of an optical imaging system is the key concept that makes our model useful for performance evaluation of networks without neglecting critical design details. One of the merits of our model is that it quantitatively approximates performance of optical interconnects for direct comparison with electrical interconnects for a given topology, assuming the connection capacity and the physical mapping of the topology are known.

According to our results, channel cycle time grows comparatively faster in electronics than in optics because of line capacitance. In contrast, the channel cycle time (and network latency) in optics depends largely on signal propagation and conversion. Optical-to-electrical conversion delay (T_{OE}) is crucial to our latency analysis since it is very sensitive

to the optical link efficiency of an imaging system. For instance, this delay for a 2-level DOE hologram (with 41% efficiency) is 3.5 ns rather than 0.9 ns for a hologram with 81% efficiency. Thus, efficient optoelectronic and micro-optics devices are needed for optical interconnects to achieve high-speed operation.

We quantitatively show that much wider channel widths can be supported in optics as compared to electronics for the same network topology due to the massive connectivity afforded by free-space implementations. Optical implementations are also seen to be more compact. Further, when wormhole routing is used, network latency is seen to be lower for low-dimensional networks, but not overwhelmingly so. This suggests that implementing high-dimensional optical networks does not degrade performance as much as electrical networks do.

Our study shows that optics is very promising even when using available technology. It features lower latency, higher bandwidth, a more compact system, and more topological design flexibility as compared to electronics. Nevertheless, a caveat with free-space optics is the practical need for a compact volume which limits system scalability. This limitation, along with other practical considerations including power dissipation and packaging tolerance must be taken into account in the design and evaluation of optical interconnects [26]. We finally show some design calculations that address the power dissipation, packaging tolerance, and devices' variation for a DROI system. In fact, their effects on the number of optical channels in a system are more stringent than the connection capacity. Hence, these considerations reduce a benefit of optical interconnects but do not overwhelm the overall performance offered by optical interconnects.

Acknowledgment

We wish to thank Professor Alexander A. Sawchuk for his helpful suggestions and relevant sources. In particular we acknowledge Dr. Charles Kuznia who provided useful information on VCSELs. Jen-Ming Wu and Chih-Hao Chen reviewed earlier drafts of this work.

Appendix

A. Guassian Beam Propagation Through a Lens

We assume an optical beam with Gaussian irradiance profile in our study which is represented by [27]

$$I = I_0 e^{(-2r^2/w^2)}, \quad (\text{a.1})$$

where I_0 is the intensity at the center of the beam, e is the base of the natural logarithm ≈ 2.718 , r is the distance from the beam center, and w is the *spot radius* of the beam where its intensity drops to $1/e^2$ of its peak value, I_0 .

As travel in free-space, the spot size of the Gaussian beam increases as [27]:

$$w(z) = w_0 \left[1 + \left(\frac{\lambda z}{\pi w_0^2} \right)^2 \right]^{1/2} \quad (\text{a.2})$$

where $w(z)$, w_0 and z represent the spot radius, travels at wavelength λ , at the distance z along propagation axis from the beam waist w_0 at $z=0$ (where its wavefront was flat, e.g., at the source windows).

The Gaussian beam propagation through a microlens is shown in Figure A-1.




Figure A-1. The Gaussian beam propagation through a microlens.

This propagation complies the lens law [27] in which the distances and sizes of object and image are shown to be

$$d_1 = f + \frac{w_1}{w_2} \sqrt{f^2 - \left(\frac{\pi}{\lambda} w_1 w_2 \right)^2}, \quad (\text{a.3})$$

$$d_2 = f + \frac{w_2}{w_1} \sqrt{f^2 - \left(\frac{\pi}{\lambda} w_1 w_2 \right)^2}, \quad (\text{a.4})$$

where d_1 and d_2 is the object and image distances, f is the focal length of microlens, w_1 and w_2 are the beam radii of the object and image.

B. Finding Connection Capacity in DROI

Each interconnect in DROI is simply described by an imaging system with two microlenses and two subholograms as shown in Figure B-1.




Figure B-1. Gaussian beam propagation in DROI.

Recall that, from Section III-D, connection capacity is a function of an area over which interconnects can be established and the maximum light beam area along the propagation path—Eq.[11]. This maximum light beam area can be observed at the microlenses as shown in Figure B-1. Due to symmetry of the system, both transmitter and receiver microlenses happen to be the same (e.g., size, focal length, and f-number).

Theoretically, for a system with no volume constraint, the interconnects area, A_{system} , will not be limited (as shown below). The only limitation in this case is processing technologies (e.g., free-space system packaging, VLSI fabrication process, etc.). In contrast, in a volume-limited system, the maximum interconnects area is represented by

$$A_{system} = \frac{2V}{R_{max} \cos\theta}, \quad (b.1)$$

where V is the system volume, R_{max} is the maximum connection path (Eq.[9]), and θ is the maximum hologram deflection angle.

Each optical beam in DROI is deflected by assuming a linear blazed grating [15] as shown below.




Figure B-2. Linear blazed grating structure.

Suppose that the linear blazed grating is implemented by binary optics in L_b levels with a feature size equal to w_f and a grating period T . After passing through the grating, if the optical signal propagates through the material with refractive index n_x then the angle of deflection, θ , can be written as [15]

$$\theta = \sin^{-1} \left[\frac{\lambda / n_x}{L_b w_f} \right]. \quad (b.2)$$

Due to its discrete feature, several diffraction orders are generated once the light beam propagates through the blazed grating. Only the first diffraction order is usable for interconnection. The hologram efficiency is therefore defined as the ratio of the first diffraction order power the total input power and is given by [15]

$$\eta_{+1} = \sin^2 \left(1 / L_b \right) = \left[\frac{\sin(\pi / L_b)}{\pi / L_b} \right]^2. \quad (b.3)$$

We show in this study that the optical link efficiency (a ratio of power at the detector and at the transmitter) plays a great role in our results. This hologram efficiency is a major contribution to that efficiency and we show here how the hologram structure affects

system performance. As a feature size in VLSI fabrication gets smaller we expect a better hologram efficiency and, thus, improving overall system performance.




Figure B-3. Optical link efficiency versus $T_{O/E}$.

Figure B-3 shows the relationship between optical link efficiency and opto-electronic conversion time at the receiver circuit according to Eq.[8]. Obviously, higher link efficiency yields lower conversion time and lower channel cycle time. However, the advantage of high link efficiency is only marginal for link efficiency greater than 0.5. Realize that high efficiency optical imaging system is not easily achieved, this figure clearly suggest that the system with 0.5 link efficiency is considerably good enough in terms of cost/performance. Besides, this efficiency is technically sound for current technology.

To find a spot size at a microlense we employ Eq.[a.2-a.4]. Parameters for Eq.[a.2-a.3] correspond to those shown in Figure B-1. First we need to find the object distance d_1 where the image distance d_2 is half of the maximum connection path determined by Eq.[9]. Our calculation shows the object distance is approximately $555 \sim 560 \mu\text{m}$. Once the object distance is found, we use Eq.[a.2] to find the spot size at the microlense. However, to collect 99.5% of beam power the microlense diameter, M_D , must be about four times larger than the beam radius impinging onto it (using Eq.[a.1]). Note that we also assume off-axis hologram does not change the beam radii; rather it elongates the image or object distances.

Without power dissipation consideration, connection capacity would be limited only by the microlens diameter and the maximum interconnection area. Hence, the maximum connection capacity of a system is simply

$$C = \frac{A_{system}}{2M_D^2}, \quad (b.4)$$

where C is the maximum connection capacity, A_{system} is the maximum interconnection area, and M_D is the microlens diameter. Factor two in the equation takes into account that both transmitter and receiver are on the same plane. Nonetheless, a practical value of connection capacity might be different from Eq.[b.4] due to not all of A_{system} is used.

Furthermore, the transmitter/receiver circuit area could be larger than the microlense itself and, hence, determine the system connection capacity.

By employing the above procedures and parameters assumed in Table 3, we find the microlense with the diameter of 125 μm , the focal length of about 460 ~ 467 μm , and the f-number of about 3.8 is required for our hypothetical system with no volume constraint. These values are practical which confirm the validity of our results.

References

- [1] Andreas G. Nowatzky and Paul R. Prucnal, "Are Crossbars Really Dead? The Case for Optical Multiprocessor Interconnect Systems," In Proceedings of The 22nd International Symposium on Computer Architecture, Santa Margherita Ligure, Italy, June 1995, pp. 106-115.
- [2] Anjan K. V. and Timothy M. Pinkston, "An Efficient, Fully Adaptive Deadlock Recovery Scheme: *DISHA*," In Proceedings of The 22nd International Symposium on Computer Architecture, Santa Margherita Ligure, Italy, June 1995, pp. 201-210.
- [3] William J. Dally, "Virtual-Channel Flow Control," IEEE Transactions on Parallel and Distributed Systems, Vol. 3, No. 2, March 1992, pp. 194-205.
- [4] Ashok V. Krishnamoorthy, Philippe J. Marchand, Fouad E. Kiamilev, and Sadik C. Esener, "Grain-size considerations for optoelectronic multistage interconnection networks," Applied Optics, Vol. 31, No. 26, September 1992, pp. 5480-5507.
- [5] Michael R. Feldman, Sadik C. Esener, Clark C. Guest, and Sing H. Lee, "Comparison between optical and electrical interconnects based on power and speed considerations," Applied Optics, Vol. 27, No. 9, May 1988, pp. 1742-1751.
- [6] Timothy M. Pinkston and Joseph W. Goodman, "Design of an optical reconfigurable shared-bus hypercube interconnect," Applied Optics, Vol. 33, No. 8, pp. 1434-1443.
- [7] T.J. Cloonan, "Architectural Considerations for Optical Computing and Photonic Switching," Optical Computing Hardware, Academic Press 1994, pp. 1-43.
- [8] Ronald A. Nordin, A. F. J. Levi, Richard N. Nottenburg, J. O'Gorman, T. Tanbun-Ek, and Ralph A. Logan, "A Systems Perspective on Digital Interconnection Technology," Journal of Lightwave Technology, Vol. 10, No. 6, June 1992, pp. 811-827.
- [9] Ronald A. Nordin, William R. Hollowd, and Muhammed A. Shahid, "Advanced Optical Interconnection Technology in Switching Equipment," Journal of Lightwave Technology, Vol. 13, No. 6, June 1995, pp. 987-994.
- [10] A. Guha, J. Bristow, C. Sullivan, and A. Husain, "Optical interconnections for massively parallel architectures," Applied Optics, Vol. 29, No. 8, March 1990, pp. 1077-1093.
- [11] William J. Dally, "Performance Analysis of k -ary n -cube Interconnection Networks," IEEE Transaction on Computers, Vol. 39, No. 6, June 1990, pp. 775-785.

- [12] Karl-Heinz Brenner and Frank, "Diffractive-reflective optical interconnects," *Applied Optics*, Vol. 27, No. 20, October 1988, pp. 4251-4254.
- [13] S. Matsuo et al., LEOS Annual Meeting, 1994.
- [14] Neil H. E. Weste and Kamran Eshraghian, "Principles of CMOS VLSI Design," Addison-Wesley 1988, pp.132-134.
- [15] J. Jahns, "Diffractive Optical Elements for Optical Computers," *Optical Computing Hardware*, Academic Press 1994, pp. 137-167.
- [16] A. L. Lentine, K. W. Goossen, J. A. Walker, L. M. F. Chirovsky, L. A. D'Asaro, S. P. Hui, B. T. Tseng, R. E. Leibenguth, D. P. Kossives, D. W. Dahringer, D. D. Bacon, T. K. Woodward, and D. A. B. Miller, "700 Mb/s operation of optoelectronic switching nodes comprised of flip-chip-bonded GaAs/AlGaAs MQW modulators and detectors on silicon CMOS circuitry." Postdeadline papers, Conference on Lasers and Electro-optics (CLEO) (1995), paper CPD-11.
- [17] Timothy J. Drabik, "Optoelectronic Integrated Systems Based on Free-Space Interconnects with an Arbitrary Degree of Space Variance," *Proceeding of the IEEE*, Vol. 82, No. 11, November 1994, pp. 1595-1622.
- [18] James Buchanan, "CMOS/TTL Digital Systems Design," McGraw-Hill 1990, pp. 184-185.
- [19] Clyde F. Coombs Jr., "Printed Circuits Handbook," 3rd edition, McGraw-Hill 1988, pp. 4.24-4.25.
- [20] Kevin Bolding, Sen-Ching Cheung, Sung-Eun Choi, Carl Ebeling, Soha Hassoun, Ton Ngo, Robert Wille, "The Chaos Router Chip: Design and Implementation of an Adaptive Router," *VLSI '93, IFIP*, Sept. 1993.
- [21] T. B. Alexander, K. G. Robertson, D. T. Lindsay, D. L. Rogers, J. R. Obermeyer, J. R. Keller, K. Y. Oka, and M. M. Jones, "Corporate Business Servers: An Alternative to Mainframes for Business Computing," *HP Journal*, June 1994, pp. 8-33.
- [22] Philippe J. Marchand, Ashoj V. Krishnamoorthy, Sadik C. Esener, and Uzi Efron, "Optically Augmented 3-D Computer: Technology and Architecture," in *Proceedings of the First International Workshop on Massively Parallel Processing using Optical Interconnects*, April 1994, pp. 133-139.
- [23] W. Stephen Lacy, Christophe Camperi-Ginestet, Brent Buchanan, D. Scott Wills, Nan Marie Jokerst, and Martin Brooke, "A Fine-Grain, High-Throughput Architecture Using Through-Wafer Optical Interconnect," in *Proceedings of the First International Workshop on Massively Parallel Processing using Optical Interconnects*, April 1994, pp. 27-36.
- [24] K. W. Goossen, A. L. Lentine, J. A. Walker, L. A. D'Asaro, S. P. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. Dahringer, L. M. F. Chirovsky, and D. A. B. Miller, "Demonstration of a dense, high-speed optoelectronic technology integrated with silicon CMOS via flip-chip bonding and substrate removal," 1995 Spring Tropical Meeting—Optical Computing Section, Salt Lake City UT, March 1995, pp. 142-144.

[25] Information from AMD Word-Wide-Web home page at <http://www.amd.com/>.

[26] Timothy M. Pinkston, "Design Considerations for Optical Interconnects in Parallel Computers," in Proceedings of the First International Workshop on Massively Parallel Processing using Optical Interconnects, April 1994, pp. 306-322.

[27] Miles V. Klein and Thomas E. Furtak, "Optics," 2nd edition, John Wiley & Sons, 1986, pp. 474-478.