

Highly efficient interconnection for use with a multistage optical switching network with orthogonally polarized data and address information

Dane C. Butzer, Bradley D. Clymer, and Betty Lise Anderson

A novel optical interconnection is introduced for a multistage optical switching network that uses orthogonally polarized data and address information. The network is unique in that the data information is never regenerated and remains in optical form throughout (i.e., it is never converted into electrical information). This has two main consequences: (1) the bandwidth of the data is not restricted by electrical circuit considerations, and (2) the optical interconnections from one stage of the network to the next must be highly efficient. The interconnection meets several goals: high efficiency, preservation of cross polarization of data and address, low cross talk between polarizations, good manufacturability, resistance to misalignment caused by thermal expansion, and absence of significant aberrations. In addition, synchronization of the signals is maintained, as the optical path lengths for all routes through the system are equal.

1. Introduction

A. Background

Multistage permutation networks consisting of multiple stages of 2×2 switches connected in various ways include Batcher, omega, indirect binary n -cube, and R networks.¹ These networks have many applications, including fast-Fourier transforms, polynomial evaluation, sorting, and matrix transposition.² Research in the area of optical implementations of these networks has focused on either the connections between the stages or on the 2×2 switch implementation. The data-addressing mechanism has generally been omitted. An exception is a system that has been proposed with orthogonally polarized data and address information.³

The advantage of having orthogonally polarized data and address information is that the data need never be converted into electrical information throughout the network. When the address information needs to be converted into an electrical signal for use by the switching elements, a portion of it is

stripped off the input beam by a partial polarizing beam splitter (PPBS) and then converted into an electrical form. The data signal is not affected by the beam splitter, permitting it to stay in optical form throughout the system. As a result, neither data nor address signals are regenerated. This has two main implications: (1) the bandwidth of the data is not restricted by electrical circuit considerations; and (2) the optical interconnections from one stage of the network to the next must be highly efficient.

The specific system to be examined is a switching network that routes up to eight channels to specific outputs. This network consists of an eight-element Batcher network in conjunction with an omega network. The Batcher network sorts the inputs into an ascending address order. Unused inputs are sorted to the bottom of the network. The omega network then directs the data to the appropriate outputs. The graph of the total network is given in Fig. 1. The fundamental concepts developed for this interconnection are equally applicable to other multistage switching topologies.

B. Objectives

The objectives for this interconnection system are as follows:

- (1) High efficiency, preferably over 90% (power) for each stage.

The authors are with the Department of Electrical Engineering, Ohio State University, 2015 Neil Avenue, Columbus, Ohio 43210.

Received 7 March 1994; revised manuscript received 18 October 1994.

0003-6935/95/111788-13\$06.00/0.

© 1995 Optical Society of America.

- (2) Preservation of cross polarization of data and address signals.
- (3) Low cross talk between polarizations.
- (4) Good manufacturability.
- (5) Resistance to misalignment due to thermal expansion.
- (6) Absence of significant aberrations.

The motivation for the first goal has been briefly discussed (cf. Subsection 1.A). The value of 90% power efficiency is arbitrarily chosen. This results in approximately 57% power attenuation through the nine stages (eight interconnections) of the network.

The second and the third goals are necessary in order to prevent corruption of the address and the data information. Because the address signal is separated from the data signal by PPBS's, any part of the data signal that crosses to the address polarization and vice versa is noise. A nonguided implementation is chosen because construction and use of a waveguide that maintains separation between copropagating cross polarizations is rather difficult.

Good manufacturability and resistance to misalignment due to thermal expansion help to ensure that a practical system is developed. This implies that aligning the interconnection elements should be as simple as possible, the size of the system should be reasonably small, and the number of elements in the system should be kept to a minimum.

Aberrations must be avoided because the switches tend to rely on the incoming signal having undistorted and relatively planar wave fronts. This is true for the two switches being considered here, an acousto-optic switch³ and a ferroelectric liquid-crystal switch.⁴

C. Organization

The following sections present the theory and experimental results pertaining to the optical interconnection introduced. The network under consideration is reviewed in Section 2. Then, the utility of various types of interconnections are examined in Section 3. A new interconnection is proposed in Section 4. Layout, efficiency, polarization, and wave-front curvature considerations are examined. Conclusions are presented in Section 6.

The scope of this paper is limited in two areas. First, only two-dimensional network topologies are considered. Three-dimensional networks are left for future research. Second, the size of the switching

elements is assumed to be macroscopic. Therefore miniature systems are not considered.

2. Switching Network

A. Topology

The graph of the particular switching network to be examined is given in Fig. 1. This network has eight inputs. Each box corresponds to a 2×2 switch. Each switching element either passes or swaps its inputs (see Fig. 2).

The first six stages of this network correspond to a Batcher network.⁵ The last three stages correspond to an omega network.⁶ This combination of a Batcher and an omega network, introduced in Ref. 7, creates a nonblocking network that routes an arbitrary input consisting of $N = 2^n$ differently addressed signals to the proper outputs. This network has a total of $(n^2 + n)/2 + n$ stages of 2^{n-1} switches.

In order for the omega-network part of the system to cope correctly with unused inputs, the Batcher network must sort them to its bottom (or top) outputs. For this reason, inactive signals are treated as having a higher (or lower) address than all other signals in the Batcher network. Thus the omega network and the system as a whole are nonblocking for inputs with up to N unique addresses.

The overall network does not cope with multiple inputs that have the same address. For example, the input [4, 2, 7, 3, 3, unused, unused, unused] would be sorted by the Batcher network to [2, 3, 3, 4, 7, unused, unused, unused]. Then, the two signals with address 3 would block each other in the omega network. The result would be routing of one of the inputs with address 3 to the incorrect output.

The switching decisions for these networks are completely local to each 2×2 switch. In other words, each switch can determine how to switch the two incoming signals (pass or exchange) without any other information from the rest of the network or any other signals. Only the address information for the two inputs is required. For the Batcher network the decision is based on a comparison of the values of the addresses. For the omega network the decision is based on the binary representation of each of the addresses of each of the inputs.

The representation of the switching elements in Fig. 1 can be misleading. In actuality the input and the output beams are not parallel unless extra optics are incorporated. For example, the acousto-optic switch in Ref. 3 and the ferroelectric liquid-crystal




Fig. 1. Graph of the sorting network.




Fig. 2. Configuration of a typical 2×2 switch. The inputs are either passed or swapped.




Fig. 3. Inputs and outputs for the 2×2 switches are not parallel unless extra optics are incorporated. For the acousto-optic switch, $\theta = 45^\circ$. For the ferroelectric liquid-crystal switch, $\theta \approx 68^\circ$.

switch in Ref. 4 have input-output configurations as shown in Fig. 3. Extra optics, such as mirrors or prisms, are necessary to redirect the input and the output beams into parallel paths; however, the optics involved are considered as part of the interconnections. In other words the presence of these extra optics increases the number of elements in the system. In the case in which the extra optics are not present, the input (and output) beams are nonparallel.

B. Other Topologies

Any multistage network consisting of 2×2 switches can utilize the interconnection developed here. Networks with more than eight inputs, such as a Batcher-omega network with 16 inputs (and hence 14 stages of 8 switches), can also be implemented. Therefore when issues such as loss, signal cross talk, and wave-front curvature are considered, the case of an arbitrary number of stages is discussed. Different networks such as the crossbar network⁸ can also be implemented.

C. Orthogonally Polarized Address and Data

Eight beams of light are used as the inputs to the network, each beam containing both address and data information. The exact method by which this information is coded is discussed in Ref. 3. The method by which the address information is separated from the data information at each switch is pertinent.

If the address information is simply coded as a header to the data information, it would be necessary to convert some portion of each beam to electrical form at each switch. This results in attenuation of both the data and address information throughout the network. As the data information can be assumed to have a much higher bandwidth than the address information, this loss could have a detrimental effect on its signal-to-noise ratio.

Alternatively, the address information can be coded on the orthogonal polarization of a beam for each data signal.³ Then, a portion of the address information can be stripped from the beam with a PPBS at each switch with minimal impact on the data information. The address information can then be converted to electrical form and decoded. In addition, the ad-

dress information can be present during the entire period when the data information is valid. This simplifies the switching decision at each 2×2 switch because no storage of the current input addresses or parsing of the input signal (into address and data) is then necessary. With this arrangement the data remain in optical form throughout.


D. 2×2 Switches

Two types of switches are considered for this network, an acousto-optic switch³ and a ferroelectric liquid-crystal switch.⁴ The second switch has a higher switching speed, but it must be slightly modified to work with this network. Other types of switches may also work for this architecture.

The configuration of the acousto-optic switch is shown in Fig. 4. If the beams are to be swapped, they simply pass through the acousto-optic modulator (AOM). If they are to be passed without swapping, an acoustic wave is generated in the AOM, causing the beams to diffract.

The ferroelectric liquid-crystal switch is shown in Fig. 5. Again, a pair of partially polarizing beam splitters (PPBS's) strip off a small portion of the address information. Then, both inputs pass through a single full polarizing beam splitter (PBS). This beam splitter is constructed of a nematic liquid-crystal (NLC) cell sandwiched between two glass prisms. It separates the polarizations of the beams so that one path carries the address from one signal and, orthogonally polarized, the data from the other signal. The other path carries the remaining address and data information. To understand this, one needs to realize that the data polarizations in both beams are in the same direction (say, p polarized) as the address polarizations (say, s polarized). The beam splitter transmits one polarization and deflects the other. Assume that the s polarization is transmitted. Then, the data information (p polarized) from the top beam is reflected, whereas the address information (s polarized) from the bottom beam is transmitted. These two follow the top path after the beam splitter. The other data and address signals follow the bottom path.

After the beams pass through the polarizing splitter, they propagate through a pair of surface-stabilized ferroelectric liquid-crystal (SSFLC) cells. If the SSFLC's do not rotate the polarizations of the beams, the data and the address signals recombine and propagate in the direction that they entered the switch. This is an exchange operation, as the top



Part of Address Polarizations

Fig. 4. Configuration of the acousto-optic 2×2 switch.




Fig. 5. Configuration of the ferroelectric liquid-crystal 2×2 switch.

input is routed to the bottom output and vice versa. If the SSFLC's do rotate the beams, the signals again recombine properly. The directions of the beams are switched, resulting in a pass operation.

In the exchange operation the output beams have the same polarizations as the input beams. This is not the case with the pass operation. If these switches were used, any subsequent switches could not distinguish which polarization to use for address information. Therefore an additional pair of SSFLC's must be added to the output of the switch. These SSFLC's rotate the polarizations back to their original directions if necessary. This extra pair of SSFLC's is not included in the switch described in Ref. 4. They are necessitated by the orthogonal polarization of the data and the address information for this network.

A pair of Fabry-Perot cavities exist within the ferroelectric switch⁴ because each beam splitter is made of a pair of glass prisms sandwiching a passive nematic liquid-crystal (NLC) cell. The slight index mismatch between the glass and the liquid crystal forms the cavity. The effect of the cavity is to reflect part of the transmitted polarization onto the reflected path. This leads to some cross talk between the signals. This effect can be minimized by proper choice of the angle of incidence for the beams.

3. Survey of Interconnection Elements

A preliminary layout for implementing the network described in Section 2 and the utility of diffractive, refractive, and reflective routing elements for this layout are presented in Ref. 9. Routing elements are those passive elements that connect the active 2×2 switches. The important results from this discussion are summarized here.

A. Preliminary Layout

An intuitive approach to constructing the network would be to align the stages of 2×2 switches in columns or planes and to place routing optics between each stage. A two-dimensional version of this approach, in which the switches are aligned in columns, forms the basis for the preliminary layout. With this approach the smallest worst-case deflections for the optical beams are always greater than 45° .

The size of the overall network is determined by the spacing between the switches. The ratio of the

vertical spacing between switches within a column (VS) to the horizontal spacing between columns (HS) gives a measure of the network size that scales easily with component size.

B. Diffractive Elements

Thin holograms, especially computer-generated ones, are relatively easy to mass produce. This is especially true in light of recent research involving the use of phototypesetters such as the Lintronic 300 Imagesetter.¹⁰ Unfortunately, diffraction efficiencies of the order of 20% to 30% are considered good for these elements.¹¹ In addition, as the deflection angle increases, the deflection efficiency for the first-order beam decreases greatly. Because of their low efficiency, thin holograms by themselves are not suitable for the routing elements in this system.

Thick holograms and diffraction gratings, on the other hand, are capable of extremely high diffraction efficiencies. With a proper choice of films (such as a dichromated gelatin) or grating profiles (triangular) the deflection efficiencies have a theoretical limit of 99.0% for the first-order beam.¹²

C. Refractive Elements

Refractive elements include lens arrays and prisms. If lenses are used by themselves to deflect the beams in the preliminary layout, significant aberrations such as coma are likely to occur because of the large deflections needed. If the VS:HS spacing ratio for the switching elements is decreased, these angles are reduced; however, the system size increases disadvantageously. For example, going from a 90° worst-case deflection to a 50° worst-case deflection causes the system width to increase by a factor of 3.73 over the 1:1 case.

Another solution involves permuting the order of the switches in each column of the system. By doing this, one can implement the network with low aberration and high efficiency.¹³ Again, the VS:HS ratio is decreased, but not as severely. Unfortunately, additional optics are still required in order to align the inputs to each switch appropriately.

Prisms appear to be well suited to deflecting optical beams by large angles. However, the Fresnel losses associated with large deflections rule out these elements.

D. Reflective Elements

The final type of interconnection element is reflective. These elements include metallic mirrors, such as gold-coated aluminum, and dielectric mirrors. Reflective elements have been utilized in different architectures. For example, highly efficient film-based reflective elements can be embedded in the system.¹⁴

With the preliminary layout a pair of routing elements is again required for each connection. Therefore the reflection efficiency for a single mirror should be at least 95%. The highest efficiency for

reflective elements is achieved by use of total internal reflection (TIR). Using TIR yields an efficiency of 100% and leads to a novel system layout. The new layout is considered in Section 4.

4. Proposed Interconnection

A novel layout for implementing the switching network is presented here. Efficiency, polarization, and wave-front considerations for the layout are examined. The utility of the layout for other topologies is also considered.

While two specific types of 2×2 switches are envisioned for this architecture, other types of 2×2 switches may also work. An attempt has been made to include some of the pertinent considerations for the acousto-optic and the ferroelectric liquid-crystal switch in the following discussion. Considerations that are especially significant because of the nature of the interconnection are examined particularly.

A. Layout

A new layout for the system is shown in Fig. 6. The most important characteristic of this layout is that it contains no discrete passive routing elements. Instead, the system consists of a substrate that contains the switching elements, and the light beams are routed by TIR off the edges of the substrate. The placement of the switching elements forms the sorting network.

This network is equivalent to the one in Fig. 1, as verified by tracing of the corresponding paths through each. The switches are labeled in the new layout to facilitate this. Switches labeled with an a correspond to those in the top row of Fig. 1. Those with a b correspond to those in the second, etc. The stages are distinguished by shading. For example, switch a in the first stage of the new layout is connected to switches a and b in the second stage. This matches the graph of the network, in which the top switch in the first column is connected to the top two switches in the next column. Furthermore, the arrow on switch a in the first stage indicates that the higher-address output is routed to switch b in the next stage. This also matches the graph. Thus the first six stages form a Batcher network, and the last three




Fig. 6. Novel layout for the sorting network. Each successive stage is shaded differently than the ones before and after it. For example, the switches forming the first stage are all shaded dark gray. The switches forming the second stage are all shaded white. The first six stages form a Batcher network, and the last three form an omega network. The letters are for comparing the network with Fig. 1, and the arrows indicate the same operations as in that figure.

form an omega network, with the corresponding directions for the outputs indicated by the arrows on the switches. The entire network is a signal router based on the input addresses.

The substrate must have a sufficiently high index of refraction to ensure TIR of the beams at the edges of the network. If the system uses acousto-optic switches, the angle of incidence for the beams at the substrate-air boundary at the edges of the network is 45° . If ferroelectric liquid-crystal switches are used, the angle of incidence is $\sim 68^\circ$. Given that $\theta_c = \arcsin(1/n)$, the substrate must have an index of refraction greater than ~ 1.414 for the acousto-optic switch and greater than ~ 1.078 for the liquid-crystal switch. The substrate material should also have good thermal properties in order to satisfy goal (5) in Section 1.

Amorphous alumina is a good candidate for the substrate. It is transparent to many infrared frequencies and has an adequately high index of refraction for both switch types ($n_{\text{Al}_2\text{O}_3} \approx 1.76$). It also has excellent thermal properties.¹⁵

B. Efficiency

Because this configuration is based on TIR, the routing efficiency is 100%. In practice there may be some minor losses from surface imperfections at the substrate-air boundaries. The only major losses through the system are those caused by the 2×2 switches. These include losses from coupling into and out of the switches and losses internal to the switches.

C. Polarization Considerations

If the polarizations are imperfectly aligned, cross talk results. The geometry of a light beam in the system incident at a substrate-air boundary and undergoing TIR is shown in Fig. 7. The angle ϕ is the misalignment of the data and the address polarizations with respect to E_{\parallel} and E_{\perp} . In this subsection the data are assumed to be (approximately) in the E_{\parallel} direction, and the address is assumed to be (approximately) in the E_{\perp} direction. For $\phi = 0$ the data and the address polarizations reflect at the boundary without any cross talk. However, for $\phi \neq 0$ there is cross talk between the data and the address components of the light beam.




Fig. 7. Geometry of a total internal reflection at the substrate-air interface. The angle of incidence is θ , and the angle of polarization misalignment is ϕ .

The cross talk that arises in the case of nonzero ϕ is the result of the phase delay in E_{\parallel} with reference to E_{\perp} at TIR. The value of this delay, δ , is given by¹⁶

$$\delta = 2 \arctan \left\{ \frac{(\eta_2/\eta_1)[(v_2/v_1)^2 - 1] \sin^2 \theta}{[(\eta_2/\eta_1)^2 - 1] \cos \theta [(v_2/v_1)^2 \sin^2 \theta - 1]^{1/2}} \right\}, \quad (1)$$

where $\eta = (\mu_0/\epsilon)^{1/2} = (\mu_0)^{1/2}/n$ and $v = c/n$. The variable θ refers to the angle of incidence of the light at the boundary. Variables with subscript 1 refer to the substrate; those with subscript 2 refer to the air at the TIR boundary. Replacing v with c/n yields

$$\delta = 2 \arctan \left\{ \frac{(\eta_2/\eta_1)[(n_1/n_2)^2 - 1] \sin^2 \theta}{[(\eta_2/\eta_1)^2 - 1] \cos \theta [(n_1/n_2)^2 \sin^2 \theta - 1]^{1/2}} \right\}. \quad (2)$$

The relation between the data and the address polarizations and E_{\parallel} and E_{\perp} is given by

$$\begin{aligned} \begin{bmatrix} E_{\parallel} \\ E_{\perp} \end{bmatrix} &= \begin{bmatrix} \cos(\phi) & \sin(\phi) \\ -\sin(\phi) & \cos(\phi) \end{bmatrix} \begin{bmatrix} E_d \\ E_a \end{bmatrix} \\ &= \begin{bmatrix} E_d \cos(\phi) + E_a \sin(\phi) \\ E_d [-\sin(\phi)] + E_a \cos(\phi) \end{bmatrix}, \end{aligned} \quad (3)$$

where E_d and E_a are the fields in the data and the address polarizations. After the TIR there is a phase delay of δ in E_{\parallel} with respect to E_{\perp} , as given in Eq. (2). Thus the fields become

$$\begin{bmatrix} E'_{\parallel} \\ E'_{\perp} \end{bmatrix} = \begin{bmatrix} E_d \cos(\phi) \exp(i\delta) + E_a \sin(\phi) \exp(i\delta) \\ E_d [-\sin(\phi)] + E_a \cos(\phi) \end{bmatrix}. \quad (4)$$

Converting from E_{\parallel} and E_{\perp} back to E_d and E_a , we obtain an expression for E'_d and E'_a after the TIR in terms of E_d and E_a before the reflection, given by

$$\begin{bmatrix} E'_d \\ E'_a \end{bmatrix} = \begin{bmatrix} [E_d \cos(\phi) \exp(i\delta) + E_a \sin(\phi) \exp(i\delta)] \cos(\phi) \\ + [E_d [-\sin(\phi)] + E_a \cos(\phi)] [-\sin(\phi)] \\ [E_d \cos(\phi) \exp(i\delta) + E_a \sin(\phi) \exp(i\delta)] \sin(\phi) \\ + [E_d [-\sin(\phi)] + E_a \cos(\phi)] \cos(\phi) \end{bmatrix}. \quad (5)$$

For the case of perfect polarization alignment with the data information entirely in E_{\parallel} and the address information entirely in E_{\perp} , $\phi = 0$. There is no cross talk in this case. However, for $\phi \neq 0$ the term E'_d has a component of E_a and the term E'_a has a component of E_d . This indicates cross talk between the address and the data signals. This condition is referred to as polarization misalignment in the rest of this paper.

A computer simulation of the above equations was

performed to predict the level of cross talk versus polarization misalignment. This simulation was also verified experimentally. Table 1 shows the results of this analysis for six TIR's and misalignments of up to 5°. This is the worst-case number of reflections for the proposed eight-input system. Also, the last reflection does not always exhibit the worst cross talk. Thus the worst cross talk of the six reflections for each misalignment is shown. These results indicate that the cross-talk considerations are minimal for even significant polarization misalignments ϕ , of the order of a few degrees.

D. Wave-Front Considerations

Because TIR at a flat boundary does not distort the shape of the wave fronts of the beams, wave-front considerations are minimized. The main source for concern in this area is that the light beams travel in a substrate of limited thickness. If the edges of a beam approach the top or the bottom of the substrate, the optical situation changes from an unguided to a guided one.

In a guided-wave situation the wave-front becomes distorted by the waveguide. In addition, in the extreme case one polarization may propagate through the system with a lower efficiency than the other. Therefore one must ensure that the beam size is sufficiently small throughout the system while maintaining relatively planar wave fronts. The main factor in determining this is the divergence of a beam as it propagates through the system.

One method of ensuring small spot sizes and flat wave fronts through the system is to place a lens or lenses between all of the switches. This multitude of lenses could correct for the beam divergence as the light propagates through the system, keeping the spot size small throughout. Such an approach is sufficient from this standpoint. Unfortunately, it tends to make the system less manufacturable by greatly increasing the number of elements that must be incorporated and aligned. Alternatively, it may be possible simply to place a single lens at each input to the overall system. Whether or not this is possible is determined by the thickness and the material of the substrate and the size of the system and its elements.

Table 1. Results of Computer Simulation of Cross Talk between Polarizations for the Proposed Eight-Input Interconnection for $n = 1.8$, $\theta = 45^\circ$, and $\phi = [0^\circ, 1^\circ, 2^\circ, 3^\circ, 4^\circ, 5^\circ]^a$

ϕ (degrees)	Power in p (mW)	Power in s (mW)	Power of Powers (p/s)	Ratio of Powers (s/p)	Ratio of Powers (p/s in dB)
0	1.0000	0.0000	∞	0.000	∞
1	0.9989	0.0011	944.7	0.0011	29.75
2	0.9958	0.0042	235.7	0.0042	23.72
3	0.9905	0.0095	104.4	0.0096	20.19
4	0.9832	0.0168	58.47	0.0171	17.67
5	0.9738	0.0262	37.20	0.0269	15.71

^aThe worst-case cross talk for six consecutive reflections at each misalignment is shown.

First, the optical distance for a single beam propagating through the network must be determined. All input signals travel the same distance before reaching the outputs, independent of any switching in the system. A general expression for optical distance from any input to any output is

$$OD = kn_1 \frac{1}{\sin \theta} \frac{VS}{VS:HS} + mn_3d. \quad (6)$$

The value k is the number of interconnections for each path, and the value m is the number of stages. For the eight-input system in Fig. 6, $k = 8$ to connect $m = 9$ stages. The term n_3 is a weighted average of the indices of refraction in a switch. Letting d be the physical distance for a light beam to propagate through a switch, we choose n_3 so that n_3d gives the effective optical distance through a switch. The ratio VS:HS is the average vertical spacing between switching elements divided by the average horizontal spacing. Thus, taking VS:HS as its equivalent fraction, we obtain the term VS/(VS:HS), which gives the average horizontal spacing between switches. The VS:HS ratio for the configuration in Fig. 6 with $\theta = 45^\circ$ is 1:2, or, equivalently, 1/2. The angle of incidence of the light beam on the substrate-air boundaries (top and bottom of the system in Fig. 6) is θ . Thus $1/\sin(\theta)$ is a scaling factor from the overall horizontal interconnection distance to the overall optical interconnection distance. Therefore OD is the sum of the optical distances through all of the interconnections and through all of the switches, giving the total optical distance desired.

Now that the optical distance is known, the beam divergence through the system can be determined. Because the spot size is symmetric about $\tilde{z} = 0$, placing the beam waist in the center of the system results in the smallest possible sizes at the ends of the system. The magnitudes of the spot sizes and the radii of curvature throughout the system are $w(\tilde{z})$ and $R(\tilde{z})$ for $0 \leq |\tilde{z}| \leq OD/2$.

The maximum value of the spot size is of particular importance. Because a significant fraction of the beam energy exists outside the defined spot size, a spot size that is roughly equal to the thickness of the substrate is still unacceptable. A maximum value for the spot size is arbitrarily set at one sixth the substrate thickness. Assuming Gaussian beams, this ensures that the fields for the beams at the top and the bottom edges of the substrate are $1/(e^9) \approx 0.0001234$ of their maximum values.

The wave-front curvature throughout the system is also important. The radius of curvature alone is not a good measure of the flatness of the wave front. For example, two wave fronts with the same radius of curvature but different spot sizes have a different degree of flatness, as illustrated in Fig. 8. A better measure of flatness is the angle γ between the normal to the direction of propagation of the beam and the tangent to the wave fronts at the edge of the beam. The arctangent of $w(\tilde{z})$ over $R(\tilde{z})$ is approximately the




Fig. 8. Radius of curvature of the wave front is not an adequate measure of its flatness. Instead, the angle γ between the normal to the direction of propagation of the beam and the tangent to the wave fronts at the edge of the beam is used.

angle subtended by the arc from the center of the beam to an edge. Simple geometry can be used to show that this is equal to γ ; thus

$$\gamma(\tilde{z}) \approx \arctan \frac{w(\tilde{z})}{R(\tilde{z})}. \quad (7)$$

This is also illustrated in Fig. 8.

The spot size for a Gaussian beam at an optical distance \tilde{z} from the location of $|R| = \infty$ is

$$w^2(\tilde{z}) = w_0^2 [1 + (\tilde{z}/\tilde{z}_0)^2]. \quad (8)$$

The radius of curvature for the wave fronts of the beam at an optical distance \tilde{z} from the location of planar wave fronts is

$$R(\tilde{z}) = \tilde{z} [1 + (\tilde{z}_0/\tilde{z})^2]. \quad (9)$$

For both of these equations,

$$\tilde{z}_0 = \frac{\pi w_0^2}{\lambda}, \quad (10)$$

where λ is the wavelength of the light, n is the index of refraction for the material through which the beam is propagating, and w_0 is the minimum spot size. The value \tilde{z}_0 is approximately the optical distance at which the value for the radius of curvature, $R(\tilde{z})$, is minimum. This is also known as the (optical) Rayleigh range.

Smaller values of γ correspond to flatter wave fronts. Therefore the maximum value for γ gives the worst case. To find where this occurs, we take the derivative of approximation (7) with respect to \tilde{z} .

$$\begin{aligned} & \frac{d}{d\tilde{z}} \arctan \frac{w(\tilde{z})}{R(\tilde{z})} \\ &= \frac{1}{1 + [w(\tilde{z})/R(\tilde{z})]^2} \frac{R(\tilde{z})[dw(\tilde{z})/d\tilde{z}] - w(\tilde{z})[dR(\tilde{z})/d\tilde{z}]}{R^2(\tilde{z})}, \end{aligned} \quad (11)$$

The local maxima, if any exist, occur when Eq. (11) is

equal to zero. This can only happen when

$$R(\tilde{z}) \frac{dw(\tilde{z})}{d\tilde{z}} - w(\tilde{z}) \frac{dR(\tilde{z})}{d\tilde{z}} = 0. \quad (12)$$

With $dw/d\tilde{z} = w_0^2 2\tilde{z}/(\tilde{z}_0^2)$, $dR/d\tilde{z} = 1 - (\tilde{z}_0^2)/(\tilde{z}^2)$, and some algebraic manipulation, Eq. (12) becomes

$$w_0^2 \left(\frac{\tilde{z}^2}{\tilde{z}_0^2} + \frac{\tilde{z}_0^2}{\tilde{z}^2} + 2 \right) = 0. \quad (13)$$

This cannot occur for real \tilde{z} . Furthermore, because the left side of Eq. (13) is always greater than zero, the maximum value of γ occurs at the maximum value of $|\tilde{z}|$. Thus the beam has the least flat wave front at $|\tilde{z}| = OD/2$.

This may seem to be an obvious result, but an examination of Eq. (9) shows why the above analysis is necessary. The radius of curvature $R(\tilde{z})$ decreases with \tilde{z} up to $\tilde{z} \approx \tilde{z}_0$, then increases for large \tilde{z} . Thus if $OD/2 > \tilde{z}_0$, the wave fronts with the smallest radii of curvature would not occur at the ends of the system. In other words, with $R(\tilde{z})$ as the measurement of flatness, the most curved wave fronts [smallest $R(\tilde{z})$] would not occur at $OD/2$. As shown above, however, the most curved waves always occur at $OD/2$ when γ is used as the measure of curvature.

The upper allowable limit for γ is determined by the internal structure of the 2×2 switches used. For the acousto-optic switch the wave-front curvature must be small enough to maintain the Bragg condition for efficient diffraction. With the ferroelectric liquid-crystal switch the curvature must be small enough to minimize unwanted reflection of the through polarizations by the PBS's. This reflection is caused by the presence of a slight Fabry-Perot effect in the splitters. As different phenomena are involved in each case, the two switches are treated separately below.

An alternative and perhaps more direct method of approaching the question of wave-front curvature is to set the maximum allowable optical path length to twice the Rayleigh range for the light beam. With this approach the minimum radii of curvature will exist at the ends of the network (with a flat wave front at the center). Hence if the wave fronts are sufficiently flat at the Rayleigh range, they will be sufficiently flat throughout the network. However, as just mentioned, the critical measure for wave-front flatness for the 2×2 switches is γ , not $R(\tilde{z})$. With γ as the controlling design factor the minimum $R(\tilde{z})$ may occur well within the network. Toward the ends of the system it is possible that $R(\tilde{z})$ may be increasing, whereas γ is still sufficiently limited. Therefore limiting the network parameters so that the ends of the system correspond to the Rayleigh range is a sufficient method for limiting wave-front curvature, but it may be overly restrictive in some cases.

E. Polishing Considerations

The quality of the reflective surfaces also has an impact on the polarization and the wave-front curvature considerations. If the reflecting surfaces have significant imperfections or aberrations, cross talk could occur between the polarizations, and the wave fronts could become distorted. Polishing the reflecting surfaces to a satisfactory degree of flatness is well within the scope of current technology. The polishing is simplified because all reflections occur on two flat surfaces. Any polishing technique that results in $\lambda/8$ flatness should be adequate.

F. Acousto-Optic Switch

As mentioned above, one switch implementation uses an acousto-optic modulator. This device either permits the input beams to pass unchanged or redirects them. The process of redirection is equivalent to diffraction of the light by a moving, periodic grating. The maximum diffraction efficiency possible is given in Ref. 3 by

$$\eta = \frac{\kappa^2}{\kappa^2 + (\frac{1}{2} K \Delta\theta)^2}, \quad (14)$$

where κ is a coupling constant, $K = 2\pi/\Lambda$ is the wave number for the sound wave, $\Delta\theta$ is the Bragg-angle error, and Λ is the acoustic wavelength. The coupling constant is given by

$$\kappa = \frac{\pi}{\lambda} \left(\frac{M I_{\text{acoustic}}}{2} \right)^{1/2}. \quad (15)$$

Here, λ is the wavelength of the light and I_{acoustic} is the acoustic intensity.

When $\Delta\theta = 0$, the diffraction efficiency of the AOM is maximum. This corresponds to a perfectly flat wave front striking the modulator at exactly the Bragg angle. For the efficiency to be reduced by a significant amount, $\frac{1}{2} K \Delta\theta$ must be significant compared with κ . Because $\kappa \gg K$, $\Delta\theta$ can be fairly large. From Fig. 8, $\Delta\theta$ is seen to be the same as γ . A somewhat arbitrary but conservative upper limit for γ is set at 0.1° .

G. Ferroelectric Liquid-Crystal Switch

A slight index mismatch exists between the glass prisms and the NLC cells forming the full PBS's in the ferroelectric switch. This forms a parasitic Fabry-Perot cavity whose reflectivity is a function of its electrical length.¹⁷ This length is the number of wavelengths it takes to cross the cavity. It is given by

$$L_e = 2\pi l/\lambda, \quad (16)$$

where l is the distance the light travels across the cavity and λ is the wavelength of the light in the material.

The reflectivity of the cavity is minimum when the electrical length is an integer multiple of π . The light from the edges of a curved, contracting wave front travels a distance across the cavity different than the light in the center of the beam. Recall that the wave front is most curved at the ends of the system; thus the distances traveled by the light at the edges of the beams can be approximated as the lengths of the normals to the tangents of the wave front at its edges through the cavity. This is illustrated in Fig. 9.

In order to find the maximum variation in the electrical length, we can replace l in Eq. (16) by $l + \Delta l_1$ and $l - \Delta l_2$. Some minor trigonometric manipulation yields

$$\Delta l_1 = l \left[\frac{\cos \theta}{\cos(\theta + \gamma)} - 1 \right]. \quad (17)$$

$$\Delta l_2 = l \left[1 - \frac{\cos \theta}{\cos(\theta - \gamma)} \right]. \quad (18)$$

Because

$$\left| \frac{\cos \theta}{\cos(\theta + \gamma)} - 1 \right| > \left| 1 - \frac{\cos \theta}{\cos(\theta - \gamma)} \right|, \quad (19)$$

Δl_1 gives the maximum change in distance and is always used. The value of Δl_1 should be small enough that it does not change the electrical length by a significant amount in comparison with π . Because the cavity is parasitic, it has a low finesse, and a small change in the electrical length does not cause a significant change in the reflectivity across the wave front.

As an example, for $\lambda = 1300$ nm, Δl_1 should be much less than $6.5 \mu\text{m}$ to keep the electrical distance




Fig. 9. Distance traveled by light at the edges of a curved wave front through a Fabry-Perot cavity differs from that at the center of the beam. Assuming a spherical wave front, Δl_1 and Δl_2 give the changes in the electrical distance for a given γ .

much less than π . Solving Eq. (17) for γ gives

$$\gamma = \arccos \left(\frac{1}{1 + \Delta l_1} \cos \theta \right) - \theta. \quad (20)$$

For an l of the order of $10 \mu\text{m}$ (Ref. 18) and $\theta = 45^\circ$ this implies that γ should be much less than $\sim 3^\circ$. For $\theta = 68^\circ$, γ should be much less than $\sim 1.6^\circ$. In both cases, 0.1° is chosen as a somewhat arbitrary but sufficient limit for γ .

H. Alignment between the 2×2 Switches

The polarization misalignment discussed so far is misalignment of the data and address polarizations with respect to E_{\parallel} and E_{\perp} for the TIR's. The issue of polarization misalignment due to any tilt of one switch with respect to another has not been addressed. Misalignment considerations between switches are mitigated by the fact that all of the switches are placed in a single substrate. The rigidity of the overall system should greatly limit any possible tilts. For example, if the switches themselves come in rigid packages, then the substrate could be constructed with guides or grooves that force correct placement and alignment of the switches with respect to each other. If one of the switches is misaligned because of internal misalignments, then all of the switches would be similarly misaligned. Thus the most pertinent alignment consideration becomes the polarization misalignment at the TIR's. The precision with which the switches could be aligned with one another in an actual network and the construction of such guides and grooves are areas for further research.

I. Using a Single Lens

It may be possible to focus a beam to a beam waist at the center of the system with a single lens at the input in such a way that the beam has a sufficiently small spot size and a sufficiently flat wave front throughout the network. As an example with the proposed network, let the substrate material be a 0.5-cm-thick piece of amorphous alumina (Al_2O_3). This substrate has an index of refraction of approximately 1.76. Let θ be 45° , VS be 1.5 cm, and the optical distance through a switch $n_3 d$ be 1 cm. Recall that there are eight interconnections ($k = 8$) between nine switches ($m = 9$) for each path through the system. Also, the VS:HS ratio for the proposed system is 1:2. Then, from Eq. (6) the total optical distance through the network is 0.687 m. Using a wavelength of 1300 nm, choosing 0.05 cm as the minimum beam waist, and applying eq. (10), we obtain $\bar{z}_0 = 0.604$ m. From Eq. (8) the maximum spot size is $w(0.687/2) = 0.0575$ cm. This is well below the limit of $0.5 \text{ cm}/6 = 0.0833$ cm. From Eq. (9) the value of the $R/OD/2$ is 1.406 m. Using Eq. (7), we obtain $\gamma(0.687/2) = 0.0234^\circ$. This is also well within the set limit. Therefore a lens that focuses an input beam to a waist of 0.05 cm at the center of the system is sufficient to ensure that the beam meets the criteria for acceptable spot size and wave-front curvature. In this case,

only a single lens is needed at each input to the network, as opposed to lenses between all of these stages.

Table 2 shows the results of this analysis for various values of total system optical distance OD, substrate thickness, wavelength of light λ , and minimum spot size w_0 . Certain entries are italicized to help indicate which of these values change between cases. A Yes in the One Lens? column indicates that a single lens at an input to the system, focusing the beam to the specified beam waist at the center of the system, results in a beam with sufficiently small maximum spot sizes and sufficiently flat wave fronts. A No indicates that a single lens at the input does not result in adequately small maximum spot sizes, flat wave fronts, or both. The second-from-the-last combination in the table is of particular interest; reducing the minimum spot size to 0.02 cm at the center of the system actually increases the spot size at the ends of the system. This is not surprising because the beam divergence angle is larger for smaller beam waists.

J. Synchronous Data

The analysis of the wave-front-curvature and spot-size considerations above was simplified by the fact that all of the possible paths through the new layout have the same optical length. The uniformity of the path lengths has another important consequence. It ensures that synchronous signals that enter the system at different inputs remain synchronous at the outputs. In addition, because the optical path lengths for any two signals to any single switch are identical, the signals are synchronous at each switch. This synchronization may not hold between different switches in a single stage. However, because the 2×2 switches in any given stage do not communicate with each other in any way (all switching decisions are local to each switch), the lack of interswitch synchronization within a stage does not have any implications for the operation of the system.

Another network that shares the property of maintaining synchronous signals through the system was

Table 2. Wave-Front Analysis for Various Values of Total System Optical Distance OD, Substrate Thickness, Frequency Light λ , and Minimum Spot Size w_0

OD	Substrate Thickness (cm)	λ (nm)	w_0 (cm)	$w(\text{OD}/2)$ (cm)	$\gamma(\text{OD}/2)$	One Lens? ^a
0.687 m	0.50	1300	0.05	0.058	0.023°	Yes
0.5 m	0.50	1300	0.05	0.054	0.018°	Yes
<i>1.0 m</i>	0.50	1300	0.05	0.065	0.030°	Yes
<i>1.5 m</i>	0.50	1300	0.05	0.078	0.037°	Yes
<i>2.0 m</i>	0.50	1300	0.05	0.097	0.041°	No (<i>w</i> bad)
2.0 m	<i>0.75</i>	1300	0.05	0.097	0.041°	Yes
2.0 m	0.50	<i>632.8</i>	0.05	0.064	0.014°	Yes
2.0 m	0.50	1300	<i>0.02</i>	0.207	0.118°	No (<i>w</i> and γ bad)
2.0 m	<i>0.60</i>	1300	<i>0.06</i>	0.091	0.030°	Yes

^aIndicates whether a single lens at each input suffices for the system.

introduced by Wise.¹⁸ In fact, Wise's layout is similar to the new interconnection in that the optical paths through the network form a grid, and the paths are routed to form a switching network through the use of reflection. However, Wise's layout has the switches aligned in columns, and there are discrete mirrors within the system.

5. Versatility and Scalability of the Interconnection

The method for connecting the 2×2 switches for the new layout is not limited to eight-input Batcher-omega-type networks. It can also be applied to sorting networks with more inputs and stages and to networks with different topologies. For example, a $N = 16$ Batcher-omega network can be implemented. Alternatively, a pure perfect-shuffle network or an optical crossbar switch is possible. In order to construct these systems, we alter the positioning and/or the number of the switches. In addition, the direction of the outputs for each switch is modified. The method for determining the new positions is discussed here, as is the impact of increasing the number of inputs.

A. Creating the Layout

The positioning of switches for a new layout begins with a standard graph of the desired system. In this graph, switches within a stage are in columns, such as in Fig. 1. The switches are labeled [$a_i, b_i, c_i, d_i, \dots$] within each column. Then, an image of the TIR network with only one row of switches is drawn. These switches are labeled in any desired order. Next, the intersections of the optical paths from all of the switches are noted. Finally, the next stage of switches is placed at the appropriate crossings of the optical path according to the graph of the system, and the process is repeated.

For perfect-shuffle and banyan-type networks (including the Batcher and omega networks) there are always two possible positions for the switches for any stage. This is shown in Fig. 10 for a perfect-shuffle connection with [$a_0 + c_0 \rightarrow a_1, a_0 + c_0 \rightarrow b_1, b_0 + d_0 \rightarrow c_1, b_0 + d_0 \rightarrow d_1$]. The switches for the second stage could either be placed at the intersections marked with the black circles or at the intersections marked with the black boxes. The functionality of the connection is the same in each case.

The choice of switch positions for any one stage affects the possible choices for all subsequent stages. Two overall goals are considered when one chooses the switch placements: assurance that the spacing between successive stages is kept to a minimum and assurance that the input and output switches are aligned in columns. The method for picking placements to meet these two goals is as yet trial and error.

B. VS:HS Ratio

The VS:HS ratio, which determines the overall system size, has a slightly different meaning in the proposed architecture than in the preliminary layout. In the preliminary layout (stages in columns), the




Fig. 10. Two possible placements for a perfect-shuffle connection with $[a_0 + c_0 \rightarrow a_1, a_0 + c_0 \rightarrow b_1, b_0 + d_0 \rightarrow c_1, b_0 + d_0 \rightarrow d_1]$. The next stage can occupy the intersections covered by either the squares or the circles. The horizontal centers of both possible placements occur at the same place. VS and HS for either of these placements are also shown.

horizontal distance between stages (HS) is self-evident. However, in the proposed network the switches forming a given stage are not necessarily all in a column. In addition, some stages actually overlap one another. The value of the horizontal distance between stages must be redefined. In the proposed layout the horizontal distance between stages is the horizontal distance between the geometric centers of the stages. This is shown in Fig. 10. The geometric center of a stage is the point equidistant between all of the switches composing that stage. The vertical spacing, VS, for a network is not redefined. It is taken as the distance between two adjacent switches from any stage in which the switches are aligned in a column. This is also shown in Fig. 10.

For any type of interconnection between stages there is a minimum spacing between the horizontal centers of the two stages. For many interconnections the horizontal center of a stage is fixed regardless of the choice of switch placements. This is true for the example shown in Fig. 10. Other interconnections create geometries in which there are switch placements that miss the minimum spacing. These placements are to be avoided, as they result in an increase in the overall size of the network. The overall HS value for the entire system is the average of all of the HS values for each interconnection between stages.

The minimum spacing between the centers of stages depends on the network to be implemented. The minimum stage spacing for perfect-shuffle interconnections is equal to that for banyan interconnections because these connections are topologically equivalent. In these cases, with $\theta = 45^\circ$ the VS:HS ratio representing the minimum spacing is

$$VS:HS = 1:(N/4), \quad (21)$$

where N is the number of inputs to the system. This agrees with the layout in Fig. 6, where $N = 8$ and $VS:HS = 1:2$. For arbitrary θ the VS:HS ratio becomes

$$VS:HS = 1:(N/4)\tan \theta. \quad (22)$$

We can determine the above expressions by noting that, in this topology, two switches from any one stage are interconnected to two switches in the next stage. This forms a loop, as shown in Fig. 11. Two of these loops exist between each stage. The geometry of the loops constrains the locations of the centers of the stages. If one switch moves away from the center of its stage, the loops force another switch to move away. If one switch moves closer, the loops constrain another switch to move closer.

Other networks with the same number of inputs may have a different minimum HS. The optical crossbar switch is an example. The crossbar can utilize a switch at each and every intersection between light beams. Hence its VS:HS ratio is simply $1:\tan \theta$.

As HS in the VS:HS ratio increases, the size of the network becomes unwieldy. $VS:HS = 1:4$ is approaching the limit of what might be considered efficient. However, the Batcher and the omega networks can both be broken down into smaller networks.⁷ By implementing these smaller networks and connecting them, we can limit the total number of inputs to each network and thus the VS:HS ratio. However, the method for connecting the smaller networks must be considered. It must also be highly efficient and practical.

C. Input and Output Stages in Columns

As there is a minimum spacing between the centers of successive stages, there is a minimum spacing between the centers of the first and the last stages. Now, for minimization of the size of the overall system, the ends of the network need to occur as close to these centers as possible. Therefore the input and the output stages of the network should be in columns. Achieving this is, as mentioned, still a trial-and-error process.




Fig. 11. Presence of a pair of loops between stages constrains the locations of the centers of the stages for the perfect-shuffle (banyan) interconnections. The switch placement on the left corresponds to the circles in Fig. 10; the placement on the right corresponds to the boxes.

D. Efficiency, Scalability, and Examples

The efficiency of the proposed eight-input (nine-stage) network is equal to the efficiency of the switching elements multiplied by the efficiency of the interconnection. The performance of the proposed network with hypothetical individual 2×2 switching efficiencies of 90% and 95% are examined below.

For the nine-stage network and 2×2 switching efficiencies of 90% the efficiency of the switching part of the system is $0.90^9 = 38.7\%$. For nine stages and 2×2 switching efficiencies of 95%, the efficiency of the switching part of the system is $0.95^9 = 63.0\%$. There are at most six internal reflections for any signal through the nine-stage network. Thus for even a worst-case-20-dB cross talk due to interconnection polarization misalignment of just over 3° (99.05% of each signal correctly routed by each reflection), the worst-case interconnection efficiency is $0.9905^6 = 94.4\%$. Therefore with the 90%-efficient switches, at least roughly 36.4% of the input signals reach the correct outputs. For increased switching efficiencies of 95%, the overall worst-case efficiency for the proposed eight-input network is 59.5%.

For a 16-input, 14-stage perfect-shuffle network, as shown in Fig. 12(a), 2×2 switching efficiencies of 90%, and -20 dB cross talk due to interconnection polarization misalignment, roughly 21.3% of the input signals reach the correct outputs in the worst case. There are at most 12 internal reflections from the start of the network to the end for any signal. For increased switching efficiencies of 95% the overall efficiency for the 16-input network is 45.4%.

Perfect-shuffle architectures are particularly useful because they are equivalent to banyan networks (also known as binary n -cube networks), data manipulators, and flip networks.¹⁹ However, while the per-


fect shuffle of Fig. 12(a) is similar to a 14-stage Batcher network, it is not equivalent.

Scalability of perfect-shuffle and Batcher architectures beyond 16 inputs is an area for further research. We believe that any given 16-input network can be designed with VS:HS = 1:4. A 32-input network should exhibit VS:HS = 1:8. At this point the actual sizes of the switching elements, wave-front considerations, and efficiency considerations begin to limit the practicality of the interconnection. However, as switch sizes decrease and switch efficiencies increase, even larger networks may become practical.


An example of an entirely different architecture using the proposed interconnection is the 16-input crossbar switch, shown in Fig. 12(b). This is also a useful network, especially in the area of parallel computing.⁸

6. Conclusions

In conclusion, the novel interconnection meets the stated objectives. The use of total internal reflection ensures high routing efficiency. The cross polarization of the data and the address information is maintained with little cross talk. The system is simple, consisting of a single substrate in which the switches are placed. A single lens at each input and output to the system is sufficient to control beam divergence. There is a significant amount of tolerance in aligning the address and data polarizations. These three facts should result in good manufacturability. Because the interconnection consists of relatively few components, the system should be robust with respect to thermal expansion. This is especially true if a material that has excellent thermal properties, such as Al_2O_3 , is used for the substrate. Finally, because the input beams pass through only one lens and encounter reflections only at flat surfaces, optical aberrations should be minimal.



(a)



(b)

Fig. 12. (a) 16-input, 14-stage perfect-shuffle network using the new interconnection. Successive stages are indicated by shading. (b) A 16-input crossbar network using the new interconnection.

References

1. J. D. S. Parker, "Notes on shuffle/exchange-type switching networks," *IEEE Trans. Comput.* **C-29**, 213-222 (1980).
2. H. S. Stone, "Parallel processing with the perfect shuffle," *IEEE Trans. Comput.* **C-20**, 153 (1971).
3. R. B. Jenkins and B. D. Clymer, "An acousto-optic comparison switch for optical switching networks with analog addressing techniques," *Appl. Opt.* **31**, 5433-5463 (1991).
4. L. McAdams, R. McRuer, and J. Goodman, "Liquid crystal optical routing switch," *Appl. Opt.* **29**, 1304-1307 (1990).
5. K. E. Batcher, "Sorting networks and their applications," in *Proceedings AFIPS 1968 Spring Joint Computer Conference* (American Federation of Information Processing, Washington, D.C., 1968), Vol. 32, pp. 307-314.
6. D. H. Lawrie, "Access and alignment of data in an array processor," *IEEE Trans. Comput.* **C-25**, 1145-1155 (1975).
7. S. Knauer, J. H. O'Neill, and A. Huang, "Self-routing switching networks," in *Principles of CMOS VLSI Design*, N. H. E. Weste and K. Eshraghian, eds. (Addison-Wesley, Reading, Mass., 1985), pp. 428-437.

8. A. A. Sawchuk, B. K. Jenkins, and C. S. Raghavendra, "Optical crossbar networks," *Computer* (June 1987), p. 51.
9. D. Butzer and B. Clymer, "A highly efficient interconnect for use with a multistage optical switching network," in *Optoelectronic Interconnects*, R. T. Chen, ed., Proc. Soc. Photo-Opt. Instrum. Eng. **1849**, 153–158 (1993).
10. "MacBEEP offers desktop system for binary optics applications," in *Optoelectronic Reports*, November 1992 (Society of Photo-Optical Instrumentation Engineers, Bellingham, Wash., 1992), p. 11.
11. M. Feldman and C. Guest, "Computer-generated holographic optical elements for optical interconnection of very-large-scale-integrated circuits," *Appl. Opt.* **26**, 4377–4382 (1987).
12. T. Gaylord and G. Moharam, "Analysis and applications of optical diffraction by gratings," in Proc. IEEE **73**, 894–905 (1985).
13. S. Kawai, "Free-space multistage optical interconnection networks using micro lens arrays," *J. Lightwave Technol.* **9**, 1774–1777 (1991).
14. J. Skinner and C. H. R. Lane, "A low-cross-talk micro-optic liquid-crystal cell," *IEEE J. Select. Areas Commun.* **6**, 1178–1185 (1988).
15. Goodfellow Corporation, *Metals and Materials For Research and Industry* (Goodfellow Corporation, Malvern, Pa., 1992).
16. M. Born and E. Wolf, *Principles of Optics*, 6th ed. (Pergamon, Oxford, 1980).
17. J. T. Verdeyen, *Laser Electronics*, 2nd ed. (Prentice-Hall, Englewood Cliffs, N.J., 1989).
18. D. S. Wise, "Compact layout of banyan/FFT networks," in *VLSI Systems and Computations*, H. T. Kung, B. Sproull, and G. Steele, eds. (Computer Science Press, Rockville, Md., 1981), pp. 186–195.
19. C.-L. Wu, "On a class of multistage interconnect networks," *IEEE Trans. Comput.* **C-29**, 694–702 (1980).