Two-bounce optical arbitrary permutation network

Marc P. Christensen and Michael W. Haney

The two-bounce free-space arbitrary interconnection architecture is presented. It results from a series of three-dimensional topological transformations to the Benes network, the minimum rearrangeable nonblocking network. Although functionally equivalent to the Benes network, it requires only two stages of global (spanning multiple chips) optical interconnections. The remaining stages of the modified Benes interconnection network are local and are implemented electronically (on individual chips). The two-bounce network is optimal in the sense that it retains the Benes minimum number of electronic switching resources yet also minimizes the number of optical links needed for global interconnection. Despite the use of higher-order k-shuffle ($k > 2$) global optical interconnects, the number of 2×2 switching elements is identical to the two-shuffle Benes network: there is no need for $k \times k$ crossbar switches for local interconnection at each stage. An experimental validation of the two-bounce architecture is presented. © 1998 Optical Society of America

OCIS codes: 200.4650, 250.7260.

1. Introduction

Free-space optical interconnections (FSOI's) have been shown to overcome communications limitations in large, globally interconnected multiprocessor architectures by scaling well for the multiterabit bisection bandwidth regime.[1,2] Several macro-optical approaches to shuffle interconnection networks have been proposed and demonstrated.[3–10] There appears, however, to be a significant trade-off between the fundamental scaling benefits of three-dimensional (3-D) free-space macro-optical approaches and the inherent arbitrary interconnection flexibility of space-variant micro-optical interconnection approaches. While multichip macro-optical interconnection approaches, such as the one shown in Fig. 1, have been shown to scale effectively to high bisection bandwidth problems, they are limited, by their high degree of space invariance, to implementing only regular shuffle link patterns. A macro-optical interconnection approach is desired that provides arbitrary interconnections yet retains the beneficial scaling properties of macro-optics.

It is commonly assumed that using higher-order k-shuffle-based optical interconnections will require the use of $k \times k$ crossbar switches for the local switching elements to achieve arbitrary link patterns. However, as shown in this paper, 3-D topological transformations make it possible to avoid the use of $k \times k$ crossbar switches entirely, while requiring only the minimum number of 2×2 switching elements. The two-bounce architecture achieves, without changing lens positions or attributes, a completely arbitrary interconnection pattern through changing only local 2×2 switch electronic interconnections. Furthermore, the optical system can be implemented with a symmetric macro-optical multichip arrangement, thereby allowing the interconnection to be folded back onto itself in a reflective single-plane architecture that achieves the required high degree of optomechanical alignment.[10]

Section 2 of this paper provides background on the topological transformations that facilitate efficient optoelectronic packaging in 3-D FSOI systems. Section 3 reviews the Benes network and the transformations of a perfect-shuffle-based[11] Benes network that lead to the higher-order two-bounce architecture.[12] Example applications of the two-bounce architecture are illustrated in Section 4 for two distinctly different types of global permutation interconnection patterns. Section 5 is a discussion of how the two-bounce architecture can be generalized to various network sizes. An experimental module, which implements the two-bounce architecture with simulated smart-pixel arrays, is presented in Section 6. The paper is summarized in Section 7.

The authors are with the Department of Electrical and Computer Engineering, George Mason University, Fairfax, Virginia 22030-4444.

Received 28 August 1997; revised manuscript received 16 January 1998.

0003-6935/98/142879-07$15.00/0

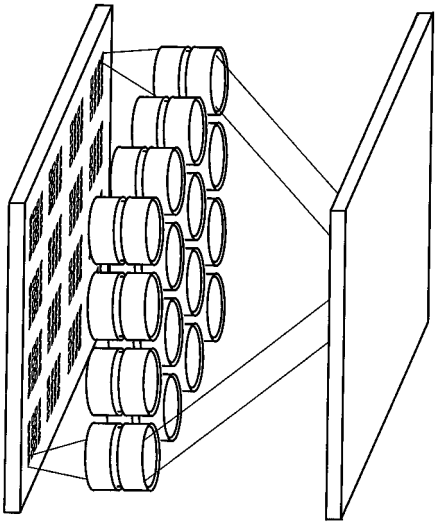

Fig. 1. Schematic depiction of a reflective macro-optical multichip interconnection module.

2. Background: Topological Transformations

The application of FSOI techniques to a multiprocessor interconnection problem can be viewed as a mapping of the network's functional interconnection pattern onto a 3-D optical interconnection architecture.[13] Such a mapping amounts to a topological transformation, which preserves the interconnection pattern and functionality of the architecture's configuration, but achieves performance advantages owing to the use of 3-D space and smart-pixel capabilities. In fact, the architecture can be represented as a series of topological transformations that each exploit a performance advantage of photonic interconnects. The cumulative performance advantage of a FSOI implementation of a network architecture is therefore derived from the aggregate advantages of several distinct topological transformations of the link interconnection pattern.

Examples of topological transformations that apply to FSOI banyan-based networks and the motivation for using them are shown in Fig. 2. Figure 2(a) shows the isomorphism between banyans consisting of butterflies and shuffles. Using an optical-shuffle link pattern between stages of the banyan simplifies the optical design and facilitates further transformations, as described below. Figure 2(b) shows the formatting of the shuffle as a two-dimensional (2-D) shuffle, rather than a one-dimensional shuffle, to take better advantage of optical and multichip module (MCM) packaging techniques. Arraying the smart pixel on self-similar grids [Fig. 2(c)] rather than on rectilinear grids increases the multichip pixel density and optical efficiency.[14] Figure 2(d) shows the spatial interleaving of multiple stages to cluster nodes, thereby reducing the amount of required electronic resources in the smart pixel.[15,16] Furthermore, if every stage is a shuffle, then this topological transformation enables the use of a single reflective optical system. Figure 2(e) shows this common

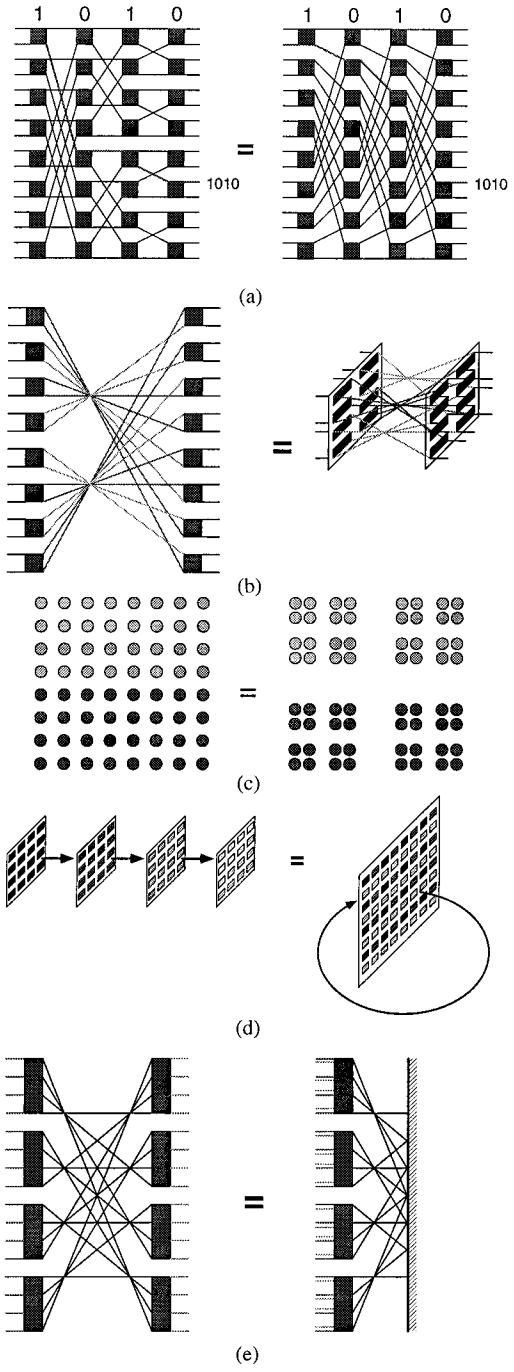

Fig. 2. Example topological transformations of multistage FSOI architectures.

plane reflective approach: to distribute the smart pixels across a single backplane, to simplify optical alignment, and to reduce the number of output drivers required.[10] Each of these FSOI topological transformations is motivated by a packaging advantage that leads to a performance enhancement or packaging simplification. The performance enhancements achieved by these topological transformations are made practical only through the use of a 3-D FSOI.

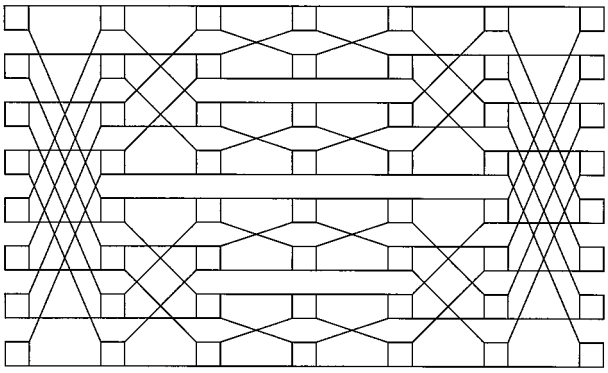

Fig. 3. Butterfly-based Benes network for $N = 16$.

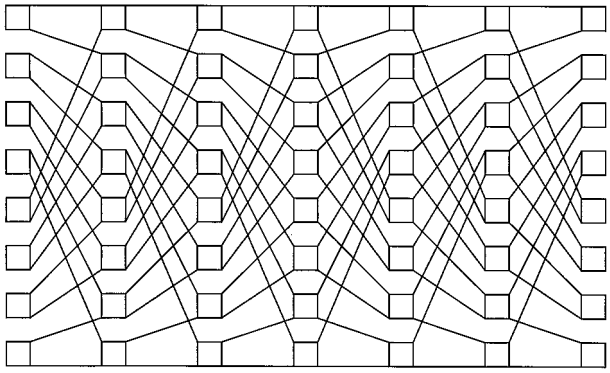

Fig. 4. Perfect-shuffle-based Benes network for $N = 16$, which is isomorphic to the network shown in Fig. 3.

3. Two-Bounce Architecture

Section 2 described the topological transformations that map regular shuffle-interconnected multistage interconnection networks (MIN's) onto optical interconnection modules like that of Fig. 1. The Benes network is a regular modulo-2 MIN-based network that achieves arbitrary rearrangeable nonblocking interconnections with the minimum number of switching resources.[17] But can the Benes network be implemented with higher-order k-shuffle optical modules without paying the increased switching penalty associated with higher-order $k \times k$ crossbars? As discussed below, the two-bounce architecture achieves exactly this result through the judicious application of topological transformations that can be implemented with the reflective k-shuffle FSOI module. The topological transformations rearrange the interconnections required for the Benes network, resulting in two stages of global interconnections, performed optically, and multiple stages of local electronic interconnections. While the two-bounce architecture retains the Benes minimum number of switching resources for an arbitrary permutation network, it also minimizes the global interconnection requirements, thereby minimizing the FSOI interconnect resource requirement. The Benes network and the topological transformations applied to it are discussed in Subsections 3.A and 3.B below.

A. Benes Architecture

The Benes network, shown in Fig. 3, consists of back-to-back butterfly networks. The resulting network consists of $2[\log_2(N)] - 1$ switching stages and $2[\log_2(N)] - 2$ interconnection stages, where N is the number of nodes. As shown in Fig. 3, the first butterfly interconnection is oriented in a forward direction, whereas the second butterfly interconnection is oriented in the reverse direction. This network has been shown to require the minimum number of 2×2 switching elements for effecting a rearrangeable nonblocking permutation network[17]; any permutation of inputs to outputs can be realized with this relatively simple switching network. The simplicity of the network has its price: the routing algorithm for the Benes network requires global information of the permutation and is iterative and therefore does not readily lend itself directly to low-latency packet-switching applications. However, the Benes network is useful for networks that can use out-of-band reconfigurations or that can store a precompiled set of interconnection patterns. For example, fast Fourier transforms (FFT's) used in digital signal processing can be implemented on multiprocessor architectures in which the processors are linked by the butterfly patterns required by the FFT algorithm. Two-dimensional FFT implementations, used, for example, in synthetic aperture radar, also require a memory corner-turn interconnection that amounts to a transpose of the data. These types of interconnections are notoriously difficult because of their high bisection bandwidth. The two-bounce architecture is particularly well suited to these types of interconnections because it can prestore the required switch settings for each stage of the FFT's butterfly as well as the corner-turn settings.

As described in Section 2, the butterfly interconnection network is isomorphic to $\log_2 N$ shuffle interconnections, as shown in Fig. 2(a). The application of this topological transformation results is a new shuffle-based Benes network, shown in Fig. 4, which comprises identical shuffle interconnections between switching elements. The identical plane-to-plane interconnection patterns make possible another topological transformation in which the interconnection module is interleaved and folded back onto itself. However, at this point, the Benes network still comprises $2[\log_2(N)] - 2$ stages of shuffles, each requiring global interconnections resources. The scaling benefits of macro-optics are best utilized when the optical interconnection pattern is global between multiple chips, i.e., higher-order shuffles corresponding to the number of optoelectronic integrated circuits (IC's) interconnected in the architecture.[1,2] This motivates the transformation of the Benes network into an architecture that uses higher-order shuffles.

B. Topological Transformation of Two-Shuffles into Higher-Order Shuffles

A perfect shuffle is a global interconnection pattern that amounts to a 1-bit rotation of an address.[11] A shuffle-exchange stage consists of a perfect shuffle

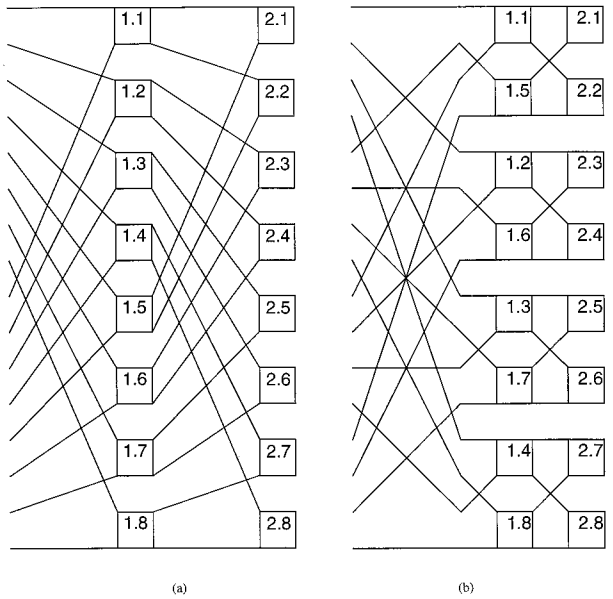

Fig. 5. (a) Two perfect-shuffle-exchange stages for $N = 16$. (b) Topological transformation of (a) to a single global four-shuffle followed by four banyans of four elements each.

followed by a set of $N/2$ 2×2 exchange–bypass switches, where N is the number of nodes. Therefore a series of M shuffle-exchange stages performs a sequence of M rotations, after each of which the locally connected bypass–exchange switch causes the least-significant bit to remain unchanged or to switch to its complement. This network can be topologically transformed into a single global $k = 2^M$ shuffle followed by routing and switching among the M least-significant bits. This makes sense because an M-stage two-shuffle MIN performs the same function as a $k = 2^M$ shuffle-based MIN that performs M left rotations (in one step) followed by a single set of $N/2^M$ banyans of M stages to set the M least-significant bits.[18] This transformation is shown in Fig. 5. Figure 5(a) shows two two-shuffle stages of 16 nodes; the switching elements are labeled for reference. Figure 5(b) shows a single four-shuffle on the same 16 nodes, with the resultant node labeling. The transformation from Fig. 5(a) to Fig. 5(b) moved only the switching elements, retaining the original interconnections between them. In this fashion, any M two-shuffle stages can be transformed into a single 2^M global shuffle followed by local routing and switching (amounting to a banyan).

This transformation of a two-shuffle-based architecture into a higher-order-shuffle-based architecture facilitates the mapping of the Benes network onto a k-shuffle optical module, where k is a power of 2. In fact, $k = N^{1/2}$ is the optimum choice for implementing the network on reflective folded modules,[19,20] such as that shown in Fig. 1, because the resulting k shuffle is symmetric[9] (i.e., the shuffle rotates half the bits). Since the equivalence of two-shuffle mappings to higher-order mappings requires an initial shuffle pattern [as shown in Fig. 5(a)], the shuffle-based Benes

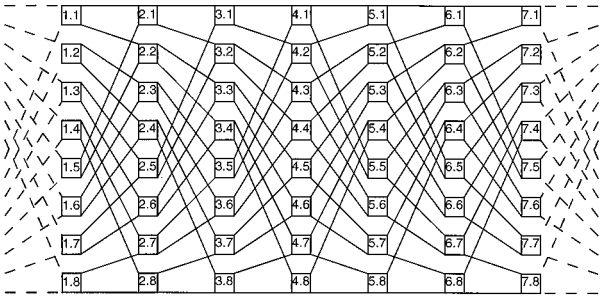

Fig. 6. Perfect-shuffle-based Benes network modified to include preshuffle and postshuffle stages.

shown in Fig. 4 must be modified to include this initial shuffle pattern to the first and the last stages of the Benes. Figure 6 shows the modified two-shuffle Benes network, with the initial and the final shuffles shown as dashed lines. Figure 7 shows the result of transforming Fig. 6 to utilize higher-order-shuffle interconnections. Note that the resultant architecture also has initial and final k shuffles ($k = 4$ in this example), again shown in dashed lines. Figure 7 is completely equivalent to Fig. 6: it contains the same number of switching elements, and they are all interconnected in the same pattern. When a module is built to realize the architecture in Fig. 7, the initial and the final global interconnections (dashed lines) are not required. The dashed lines define only a mapping between the inputs of Figs. 6 and 7, used to determine the switch settings. To implement a mapping of permutation **A** to permutation **B** in the interconnection module shown in Fig. 7, the two-shuffle Benes is solved for the mapping of **A*** to **B***, where **A*** and **B*** are defined as

$$\mathbf{A}^* = [(A)^{-4}]^2, \tag{1}$$

$$\mathbf{B}^* = [(B)^{-4}]^2, \tag{2}$$

where the -4 exponent represents an inverse four-shuffle of the pattern and the 2 exponent represents a two-shuffle of the pattern. After the switch settings are determined with **A*** and **B*** for the two-shuffle implementation, they are applied directly to the higher-order implementation and the dashed lines are not needed and are therefore not implemented.

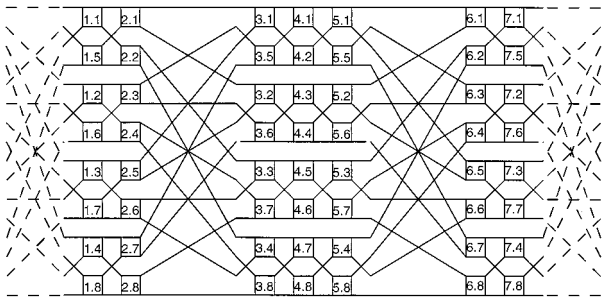

Fig. 7. Higher-order-shuffle Benes network topologically equivalent to that of Fig. 6.

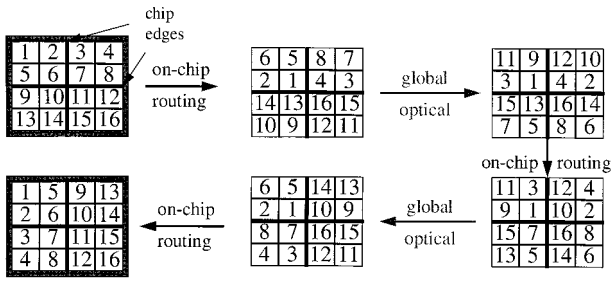

Fig. 8. Example of a two-bounce interconnection pattern: transpose.

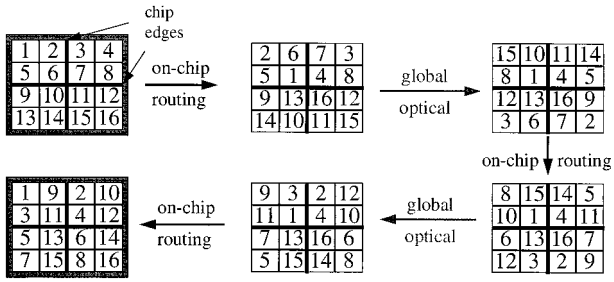

Fig. 9. Example of a two-bounce interconnection pattern: folded perfect shuffle.

Even though the global interconnection pattern is implemented with a higher-order k shuffle, the Benes network remains, logically, a two-shuffle Benes implementation. There are still $2[\log_2(N)] - 1$ switching stages and $2[\log_2(N)] - 2$ interconnection stages; only now all but two of the interconnection stages are local interconnections. The two global interconnections are symmetric optical shuffles with shuffle order (k) equal to $N^{1/2}$. Note that the local electronic routing and switching in the middle switching plane is identical to an N/k Benes network, each containing k elements. The first and the last electronic switching and routing planes each comprise simply N/k banyans with k nodes each, because they are not required for performing all permutations within the Benes structure. The result of mapping the two-shuffle Benes network onto a higher-order shuffle, while retaining the 2×2 switching, results in *fewer* switching resources than had the Benes network been constructed of higher-order shuffles, which would have required higher-order $k \times k$ crossbars at each of the three switching stages. Again, the resultant architecture comprises symmetric shuffles that facilitate the folding of the optical systems and the interleaving of resources into a module such as the one depicted in Fig. 1.

4. Interconnection Pattern Examples

To illustrate the steps involved for the two-bounce arbitrary permutation architecture, we present two example interconnection patterns with different requirements. The two interconnection patterns are a matrix transpose and a folded perfect shuffle.[7] For illustrative purposes the two-bounce interconnection is shown in six steps: (1) original data positions, (2) data positions after local electronic routing and switching, (3) data after the first global optical interconnection, (4) data after the second stage of local electronic routing and switching, (5) data after the second global optical interconnection, and finally, (6) final data positions. To make the example easy to follow, we utilize a simple 2×2 chip array with 2×2 data positions within each chip, corresponding to a data set of 16 nodes. Figure 8 is the two-bounce interconnection's effecting a transpose of the original data in a matrix fashion. Note that data remain within chip boundaries during local routing operations (between stages 1 and 2, 3 and 4, and 5 and 6). The global optical interconnections take place between stages 2 and 3 and stages 4 and 5, and they are fixed for this and all interconnection patterns when the two-bounce architecture is used.

Figure 9 shows a two-bounce interconnection effecting the perfect shuffle, in this case folded, of the data set. The global optical interconnection stages of Fig. 9 are identical to those of Fig. 8. This is a key feature of the two-bounce architecture: The optical interconnection module is fixed. No modification is required for changing the interconnection pattern. Only the local electronic routing is changed to modify a transpose interconnection to a perfect-shuffle interconnection. While the resulting optical interconnection of the two-bounce architecture is a folded perfect shuffle, the optical interconnection module is physically different. It contains two lens planes arranged in a symmetric fashion, which facilitates the folding of the optical system into the single-plane two-bounce module. Additionally, the two-bounce architecture requires one lens per chip, so a two-bounce module performing a folded perfect shuffle on a 4×4 chip array would utilize 16 lenses and perform unity magnification.

5. Discussion

The two-bounce network generalizes directly to any permutation network of size $N = 2^M$, where M is even. For example, if $M = 10$ and $N = 1024$, then an optimum choice for k is $2^{M/2} = 32$. Therefore at least 32 lenses are required for the reflective optical-shuffle module. These lenses could be arranged in a 4×8 pattern or more lenses could be utilized to make the array square. For networks of arbitrary sizes, two approaches can be considered. The network can be mapped onto the next largest readily packaged size array (2^M, where M is even), or, if the interconnection can be partitioned into a number of separate smaller arbitrary permutations, then each of these can be interleaved and implemented in parallel with a single optical system.

It has been pointed out that a symmetric k-shuffle network ($k = N^{1/2}$) allows any node to communicate with any other node with a single pass through the optical system and two stages of switching.[21] This is a k-shuffle-based banyan and therefore suffers from internal blocking (for a full permutation this amounts to $\sim 2/3$ of the data[22]), i.e., not every node can simul-

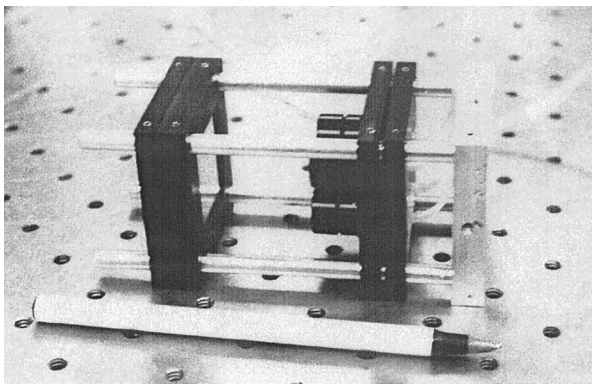

Fig. 10. Close-up view of the optical interconnection module.

taneously communicate with another node; not all permutations are possible. Since the two-bounce architecture implements the full Benes network, it truly achieves an arbitrary rearrangeably nonblocking performance with just two optical passes. These two optical passes provide the necessary global interconnection for the entire Benes network: all other interconnections are local and therefore are contained within each chip.

6. Experimental Results: Arbitrary Interconnection Prototype

Implementation of the two-bounce architecture relies on interleaving multiple k shuffles within the same optical system. A reflective optical interconnection module that interleaved multiple k shuffles on a 4×4 simulated MCM smart-pixel array[10] was designed, fabricated, and tested. This demonstration utilized a photolithographically etched mask to mimic the placement of active optical emitters and detectors. This evaluation demonstrated a registration accuracy of $\sim$10 μm over the 10-cm backplane, which is consistent with the anticipated required smart-pixel IC alignment dictated by sources such as vertical-cavity surface-emitting lasers (VCSEL's) that will be 10–20 μm in diameter. In this section the first reflective k-shuffle optical interconnection module evaluated with active sources and detectors is described.

The optical interconnection module is shown in Fig. 10. The module consists of a fiber-coupled active plane, a 2×2 lens plane, and a mirror. The purpose of this module was to apply the two-bounce architecture to implement an optical Viterbi decoding trellis,[23] which requires simultaneous forward and backward perfect shuffles.[24] Since optoelectronic IC's with the required functionality are not currently available, a fiber-coupled array was utilized for optical input–output (I/O). This two-bounce prototype has 64 fibers mounted in a face plate. The capacity of the two-bounce module greatly exceeded the number of fibers utilized for the architecture verification. In this experiment the usable (to avoid vignetting and cross talk) I/O real estate under each lens is approximately 8 mm^2. Therefore, if the four lenses were to use optical emitters and detectors spaced at 100 μm, 1600 optical links could be implemented simultaneously, leading to 100 different simultaneous arbitrary interconnections of the 16 nodes. A laptop PC performs the requisite smart-pixel functionality and is interfaced through a data-acquisition system to an array of fiber-coupled emitters and detectors. The fiber ends are situated in the optical module's backplane to mimic the placement of VCSEL's and detectors. If integrated smart pixels were utilized, only the small optical module of Fig. 10 would be present in this system.

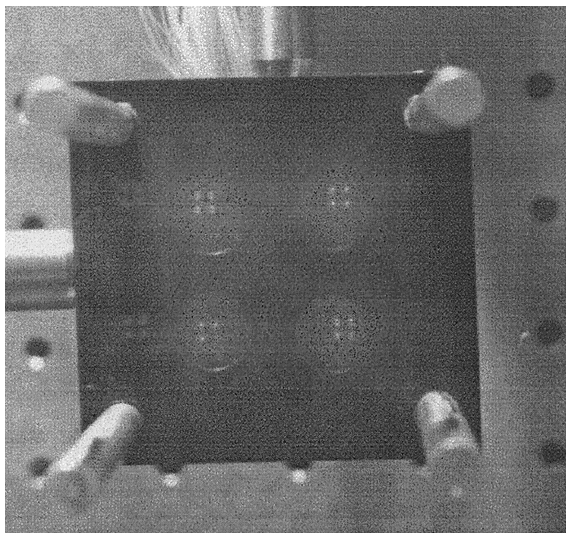

Fig. 11. Photograph of the fiber-coupled simulated smart-pixel I/O plane in the experimental setup (mirror removed).

A close-up of the active plane (mirror removed) with all the backplane emitters turned on is shown in Fig. 11. The dark fibers correspond to the detectors. There are two emitting fibers and two detecting fibers, spaced 500 μm apart, in each sixteenth of the array. In operation, a 2×2 lens array is placed a focal length above this plane, with one quadrant behind each lens. The quadrant behind each lens is a self-similar grid representation[14] of the active plane, i.e., behind each of the lenses in the 2×2 array there are 2×2 groups of fibers.

This low-density fiber-based module implements two interleaved four-shuffles on 16 nodes. This provides a single arbitrary permutation on 16 nodes. As described above, this module is capable of implementing 100 simultaneous permutations of the 16 nodes. Furthermore, this reflective module has been demonstrated to scale to larger arrays[2,10] of IC's (4×4, etc.), thereby providing arbitrary interconnections to larger numbers of nodes.

7. Summary and Conclusion

The two-bounce architecture provides an optimal mix of global optical and local electronic interconnections to achieve a rearrangeable nonblocking arbitrary interconnection network. It requires the minimum number of 2×2 switching elements and the mini-

mum number (two) of global interconnection stages. In general, the arbitrary interconnection pattern can be fixed on integration (routing wires) or circuit switched for added flexibility. Experimental results showed that this architecture could be implemented in a scalable multichip, single-plane reflective architecture. As part of the FAST-Net (Free-Space Accelerator for Switching Terabit Networks) prototype program, efforts are now focused on the implementation of the two-bounce architecture with MCM's of smart pixels consisting of high-density VCSEL/detector arrays integrated with high-performance complementary metal-oxide semiconductor logic.[25]

Aspects of this study have been sponsored by the U.S. Defense Advanced Research Projects Agency and the Ballistic Missile Defense Organization through contracts monitored by the U.S. Air Force Office of Scientific Research.

References

1. M. W. Haney and M. P. Christensen, "Fundamental geometric advantages of free-space optical interconnects," in *Proceedings of the Third International Conference on MPPOI*, A. Gottlieb, Y. Li, and E. Schenfeld, eds. (IEEE Computer Society, New York, 1996), pp. 16–23.
2. M. W. Haney and M. P. Christensen, "Performance scaling comparison for free-space optical and electrical interconnection approaches," Appl. Opt. **37,** 2886–2894 (1998).
3. A. W. Lohmann, "What classical optics can do for the digital optical computer," Appl. Opt. **25,** 1543–1549 (1986).
4. G. Eichmann and Y. Li, "Compact optical generalized perfect shuffle," Appl. Opt. **26,** 1167–1169 (1987).
5. S.-H. Lin, T. F. Krile, and J. F. Walkup, "2-D optical multistage interconnection networks," in *Digital Optical Computing*, R. Arrathoon, ed., Proc. SPIE **752,** 209–216 (1987).
6. K.-H. Brenner and A. Huang, "Optical implementations of the perfect shuffle interconnection," Appl. Opt. **27,** 135–137 (1988).
7. C. W. Stirk, R. A. Athale, and M. W. Haney, "Folded perfect shuffle optical processor," Appl. Opt. **27,** 202–203 (1988).
8. A. A. Sawchuk and I. Glaser, "Geometries for optical implementations of the perfect shuffle," in *Optical Computing '88*, P. Chaval, J. W. Goodman, and G. Roblin, eds., Proc. SPIE **963,** 270–282 (1988).
9. G. C. Marsden, P. J. Marchand, P. Harvey, and S. C. Esener, "Optical transpose interconnection system architecture," Opt. Lett. **18,** 1083–1085 (1993).
10. R. R. Michael, M. P. Christensen, and M. W. Haney, "Experimental evaluation of the 3-D optical shuffle interconnection module of the sliding banyan network," J. Lightwave Technol. **14,** 1970–1978 (1996).
11. H. S. Stone, "Parallel processing with the perfect shuffle," IEEE Trans Comput. **C-20,** 81–89 (1971).
12. M. P. Christensen and M. W. Haney, "Two-bounce free-space arbitrary interconnection architecture," in *Proceedings of the Fourth International Conference on MPPOI*, J. Goodman, S. Hinton, T. Pinkston, and E. Schenfeld, eds. (IEEE Computer Society, New York, 1997), pp. 61–67.
13. M. W. Haney, M. P. Christensen, K. Raj, and P. Milojkovic, "Packaging advantages of macro-optical free-space interconnections over micro-optical and electrical interconnections," in *Advances in Electronic Packaging 1997*, Vol. 19-1 of EEP Series, E. Suhir, M. Shiratori, and Y. C. Lee, eds. (American Society of Mechanical Engineers, New York, 1997), pp. 811–817.
14. M. W. Haney, "Self-similar grid patterns in free-space shuffle/exchange networks," Opt. Lett. **18,** 2047–2049 (1993).
15. M. W. Haney, "Pipelined optoelectronic free-space permutation network," Opt. Lett. **17,** 283–285 (1992).
16. M. W. Haney and M. P. Christensen, "Sliding-banyan network performance analysis," Appl. Opt. **36,** 2334–2342 (1997).
17. F. T. Leighton, *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes* (Morgan-Kaufmann, San Mateo, Calif., 1992).
18. A. V. Krishnamoorthy, P. J. Marchand, F. E. Kiamilev, and S. C. Esener, "Grain sized considerations for optoelectronic multistage interconnection networks," Appl. Opt. **31,** 5480–5507 (1992).
19. M. W. Haney and M. P. Christensen, "Optical freespace sliding tandem banyan architecture for self-routing switching networks," in *Digest of the International Conference on Optical Computing* (Heriot-Watt U., Edinburgh, UK, 1994), pp. 249–250.
20. M. W. Haney and M. P. Christensen, "Optoelectronic sliding banyan network," U.S. patent 5,467,211 (14 November 1995).
21. W. L. Hendrich, P. J. Marchand, F. B. McCormick, I. Cokgur, and S. C. Esener, "Optical transpose interconnection system: system design and component development," in *Optical Computing*, Vol. 10 of 1995 OSA Technical Digest Series (Optical Society of America, Washington, D.C., 1995), pp. 283–285.
22. C. P. Kruskal and M. Snir, "The performance of multistage interconnection networks for multiprocessors," IEEE Trans. Comput. **C-32,** 1091–1098 (1983).
23. M. W. Haney and M. P. Christensen, "Smart pixel based Viterbi decoder," in *Optical Computing*, Vol. 10 of 1995 OSA Technical Digest Series (Optical Society of America, Washington, D.C., 1995), pp. 99–101.
24. G. D. Forney, "The Viterbi algorithm," Proc. IEEE **61,** 268–278 (1973).
25. M. W. Haney, M. P. Christensen, P. Milojkovic, J. Eckman, P. Chandramani, R. Rozier, F. Kiamilev, Y. Liu, M. Hibbs-Brenner, J. Nohava, E. Kalweit, S. Bounnak, T. Marta, and B. Walterson, "FAST-Net optical interconnection prototype demonstration program," in *Optoelectronic Interconnects V*, R. T. Chen and J. P. Bristow, eds., Proc. SPIE (to be published).