

WDM Packet Routing for High-Capacity Data Networks

Qimin Yang, *Student Member, IEEE*, Keren Bergman, *Member, IEEE, Member, OSA*, Gary D. Hughes, and Frederick G. Johnson

Abstract—We present experimental and numerical studies of a novel packet-switch architecture, the data vortex, designed for large-scale photonic interconnections. The selfrouting multihop packet switch efficiently scales to large port counts (>10 k) while maintaining low latencies, a narrow latency distribution, and high throughput. To facilitate optical implementation, the data-vortex architecture employs a novel hierarchical topology, traffic control, and synchronous timing that act to reduce the necessary routing logic operations and buffering. As a result of this architecture, all routing decisions for the data packets are based on a single logic operation at each node. The routing is further simplified by the employment of wavelength division multiplexing (WDM)-encoded header bits, which enable packet-header processing by simple wavelength filtering. The packet payload remains in the optical domain as it propagates through the data-vortex switch fabric, exploiting the transparency and high bandwidths achievable in fiber optic transmission. In this paper, we discuss numerical simulations of the data-vortex performance and report results from an experimental investigation of multihop WDM packet routing in a recirculating test bed.

Index Terms—Optical networks, optical packet switch, wavelength division multiplexing (WDM) optical packet routing.

I. INTRODUCTION

WITH THE ever-growing demand for higher data rates and a wide variety of services, a fully transparent optical-switch network element presents an attractive way to overcome the electronic speed bottleneck. However, all-optical implementations encounter considerable challenges in the processing and buffering of optical data, which are essential to ensuring switch performance and to providing the necessary services and protections. Therefore, most optical-switch architectures actually employ hybrid optical-to-electronic technologies where the strengths of optics such as transparency and

high bandwidth can be fully exploited while the weaknesses can be bypassed in the electronic domain. Nevertheless, issues such as contention resolution, buffering, scalability, and latency remain the key challenges in photonic packet switching [1]–[4].

To achieve high throughputs, a packet switch generally processes multiple packets simultaneously. Therefore, whether the architecture is based on a single hop or multiple hops, contention is inevitable when multiple packets are competing for the same output port. To solve the contention problem, either buffering or deflection techniques can be implemented. Simple deflection methods without buffers (hot-potato routing) usually introduce severe performance penalties in throughput, latency, and latency distribution [5]. Current optical buffering techniques configured as optical fiber delay lines in a traveling or a recirculating geometry do not have random access capability [6], [7]. The straight delay-line buffers can introduce large timing errors and the recirculating rings are generally bulky and expensive. Therefore, it is still impractical to employ optical buffers within a switching network element, especially for very large-scale switches.

In recent research, the use of wavelength division multiplexing (WDM) in packet switches has been explored to achieve greater flexibility. Technologies such as wavelength routing (WR) and wavelength conversion (WC) are proposed to solve the difficulty of optical buffering [8]–[12]. These techniques are very attractive in maintaining the throughput performance as well as the optical data transparency without resorting to optical buffers. The wavelengths are exploited as logical buffers within the existing architectures. Generally, however, the additional WR and WC devices add performance penalties and make such systems complex and costly.

In this paper, we present experimental and numerical studies of a new architecture, the data vortex, as a candidate for large-scale packet switching [13]. The multihop packet switch tightly couples the deflection method and a virtual buffering mechanism to achieve hardware simplicity, scalability, and high throughput. The timing and control algorithm permits only one packet to be processed at each node in a given clock frame and, therefore, the need to process contention resolution is eliminated. The wavelength domain is additionally used to enhance the throughput and to simplify the routing strategy. Numerical studies of the traffic flow within the switch architecture have shown that the data vortex can efficiently scale to greater than 10 000 ports while maintaining a low packet switching latency and a narrow latency distribution.

The rest of the paper is organized as follows. Section II provides an overview of the data-vortex switch architecture, and in Section III, the numerical studies of the system performance are

Manuscript received January 23, 2001; revised May 30, 2001. This work is sponsored by the Defense Advanced Research Projects Agency and by the National Security Agency through an agreement with the National Aeronautics and Space Administration under Contract 960 199, and by the National Science Foundation under contract ECS98-00401.

Q. Yang is with the Electrical Engineering Department, Princeton University, Princeton, NJ 08544 USA (e-mail: qyang@ee.princeton.edu).

K. Bergman was with the Electrical Engineering Department, Princeton University, Princeton, NJ 08544 USA. She is now with Tellium, Inc., Oceanport, NJ 07757 USA.

G. D. Hughes is with the Laboratory of Physical Science, College Park, MD 20742 USA.

F. G. Johnson was with the Laboratory of Physical Science, College Park, MD 20742 USA. He is now with Little Optics, Inc., Annapolis Junction, MD 20701 USA.

Publisher Item Identifier S 0733-8724(01)08794-1.

presented. In Section IV, experimental results from the recirculating test bed are discussed. Finally, we present our conclusions in Section V.

II. ARCHITECTURE OVERVIEW [13]

The data-vortex architecture was designed specifically to facilitate optical implementation by minimizing the number of the switching and logic operations, and by eliminating the use of internal buffering. The architecture employs a hierarchical structure, synchronous packet clocking, and distributed-control signaling to avoid packet contention and reduce the necessary number of logic decisions required to route the data traffic. As a result of this novel topology and data flow, the node design for routing the optical packets is greatly simplified. All packets within the switch fabric are assumed to be the same size, and are aligned in timing when they arrive at the input ports. The routing paths of the switching architecture are designed properly to maintain this synchronous slotted operation. A detailed description of the switch architecture is provided in [13].

The data-vortex topology consists of routing nodes that lie on a collection of concentric cylinders. The cylinders are characterized by a height parameter (H) corresponding to the number of nodes lying along the cylinder height, and an angle parameter (A), typically selected as a small odd number (<10), corresponding to the number of nodes along the circumference. The total number of nodes is $(H)(A)$ for each of the concentric cylinders. The number of cylinders (C) scales with the height parameter as $C = \log_2(H) + 1$. Because the maximum available number of input ports into the switch is given by $(H)(A)$, which equals the available number of output ports, the total number of routing nodes is given by $(H)(A)(\log_2(H) + 1)$ for a switch fabric with $(H)(A)$ input-output (I/O) ports.

To illustrate the packet-routing paths through the cylinders, we show the routing tours from a top view of the concentric cylinders and from a side view of each cylinder level. In Fig. 1, an example of $A = 5, H = 8, C = 4$ switch fabric is shown. Each cross point shown is the routing node, labeled uniquely by the coordinate (a, c, h) where $0 \leq a < A, 0 \leq c < C, 0 \leq h < H$. As shown, the packets are injected at the outermost cylinder ($c = 0$) from the input ports, and emerge at the innermost cylinder ($c = \log_2 H$) toward the output ports. Each packet is selfrouted in the fashion of binary-tree decoding as it propagates from the outer cylinder toward the inner cylinder. Every cylindrical progress fixes a specific bit within the binary header address. By allowing the packet to circulate, the innermost cylinder ($c = \log_2 H$) also alleviates the overflow of the output-port buffers.

Packets are processed synchronously in a highly parallel manner. Within each time slot, every packet within the switch progresses by one angle forward in the given direction either along the solid line toward the same cylinder or along the dashed line toward the inner cylinder. The solid routing pattern at the specific cylinder shown can be constructed as follows. First, we divide the total number of nodes along the height H into 2^c subgroups, where c is the index of the cylinders. The first subgroup is then mapped as follows. For each step, we map half of the remaining nodes at angle (a) from the top

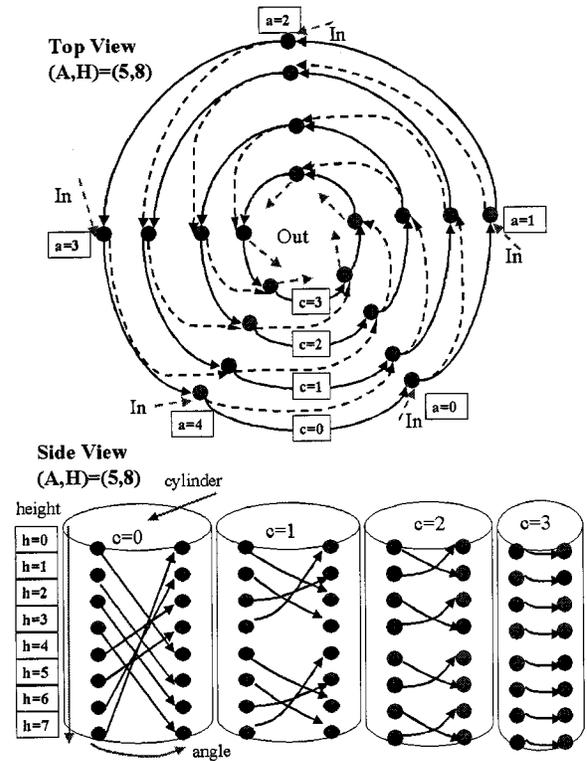


Fig. 1. Data-vortex topology with $A = 5, H = 8, C = 4$ with routing tours seen from the top and the side. Each routing node is labeled by the coordinate (a, c, h) , with $0 \leq a < A, 0 \leq c < C, 0 \leq h < H$.

to half of the remaining nodes at angle $(a + 1)$ from bottom in a parallel way. This step is repeated until all nodes of the first subgroup are mapped from angle (a) to angle $(a + 1)$. If multiple subgroups exist, the rest of them copy the mapping pattern of the first subgroup. This mapping rule can easily be verified from the example of $H = 8$ switch fabric shown. The solid routing paths are repeated from angle to angle, which provide permutations between “1” and “0” for the specific header bit. At the same time, due to the smart twisting feature of the pattern, the packet-deflection probability is minimized because of the reduced correlation between different cylinders. The dashed-line paths between neighboring cylinders maintain the same height index h because they are only used to forward the packets.

The avoidance of contention, and thereby reduction of processing necessary at the nodes, is accomplished with separate control bits. Control messages pass between nodes before the packet arrives at a given node to establish the right of way. Specifically, a node A on cylinder, c has two input ports: one from a node B on the same cylinder c and one from a node C on the outer cylinder $c - 1$. A packet passing from B to A causes a control signal to be sent from B to C that blocks data at C from progressing to A . The blocked packet is deflected and remains on its current cylinder level. As mentioned above, the routing paths along the angle dimension provide permutations between “1” and “0” for the specific header bit. Therefore, after two node hops, the packet will be in a position to drop to an inner cylinder and maintain its original target path. The control messages thus permit only one packet to enter a node in any given time period. Because the situation of two or more packets contending

for the same output port never occurs in the data vortex, it significantly simplifies the logic operations at the node and, therefore, the switching time, contributing to the overall low latency of the switching fabric. Similar to convergence routing [14], the control mechanism and the routing topology of the data-vortex switch allow the packets to converge toward the destination after each routing stage. The fixed priority given to the packets at the inner cylinders by the control mechanism allows the routing fairness to be realized in a statistical sense. The data vortex has no internal buffers; however, the switch itself essentially acts as a delay-line buffer. Buffers are located at the input and output ports to control the data flow into and out of the switch. If there is congestion at an output buffer, the data waiting to leave to that buffer circulates around the lower cylinder and, thus, is optimally positioned to exit immediately as soon as the output ports are free.

III. PERFORMANCE

The performance of the data-vortex architecture is investigated with numerical simulations. We assume a uniform, random traffic model in which all the launched packets have a balanced, uniform distribution of their destination addresses. In addition, each packet injection is statistically independent of injections occurring at other input ports or other clock cycles. The input-port injections are determined by a load parameter, which is defined as the fraction of time a new packet injection is attempted. For example, for a load parameter of 0.5, on average, a new packet injection is attempted at every other time slot. For simplicity, only the height information is encoded as the packet destination address, and all of the available angles at the output side accept the switched data packets. System performance characteristics such as injection rate, latency, and latency distribution are evaluated under various loads and various switch sizes.

A. Injection Rate

We use an injection rate parameter to represent the sustained throughput of the switch. The injection rate is defined as the ratio of successful packet injections into the switch over the number of injection attempts. Generally, the maximum throughput supported by the switch will be less than the product of the input-ports number and the bandwidth supported by each port because packets may share the same routing paths. The actual throughput can be studied by numerical simulation of the traffic flow within the switch under various load conditions for various switch fabric sizes.

Considering the number of I/O ports, the switch can be operated in two different modes: symmetric or asymmetric. As mentioned, we always assume that all of the available output angles are operated as the receiving ports. At the input side, however, we either choose to inject packets to all the angles in the symmetric mode or choose to inject packets into a fraction of the angles in the asymmetric mode.

In Fig. 2, the injection rate is plotted as a function of the input loads for various switch heights with $A = 5$. The switch is operated in the symmetric mode, i.e., all five angles are used as input ports. The switch fabric in this symmetric mode has $(5 \times H)$

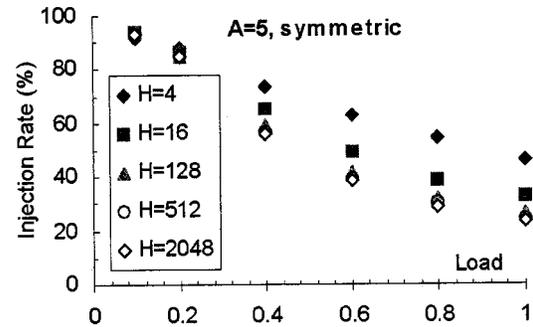


Fig. 2. Injection rate of the symmetric data vortex under various loads for different switch heights. The packet-switch size is given by $(5 \times H)(5 \times H)$.

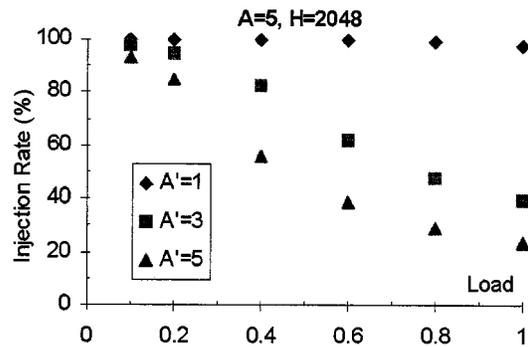


Fig. 3. Injection rate of the asymmetric data vortex under various loads for different input angles. The packet-switch size is given by $(A' \times 2048)(A \times 2048)$.

input and $(5 \times H)$ output ports. As shown, the difference in injection rate for various switch fabric sizes is more pronounced under heavier loads. However, as the switch size increases to $H > 16$, the degradation in injection rate slows down rapidly even for heavy loads. For example, in a fully loaded (load = 1.0) switch with 640 (5×128) I/O ports, about 26.1% of the new injections are successful. For the same load condition, a symmetric switch with >10 k (5×2048) I/O ports can still maintain about 23.7% successful injections. Thus, the data vortex is shown to scale to large port counts (>10 k) while maintaining a reasonable rate of successful packet injections even under fully loaded conditions.

By injecting packets into only a fraction of the available input ports, or operating asymmetrically, better injection rates and latency performances can be achieved. The downside, however, is that the switch cost rises because the same resources are now shared among a smaller number of input ports.

To study the performance improvement under the asymmetric operation, a switch of $H = 2048$ and $A = 5$ is used. In this case, only A' out of A input angles are active. The asymmetric switch now has $(A' \times H)$ input ports and $(A \times H)$ output ports.

In Fig. 3, the successful injection rate is shown as a function of the traffic load for varying numbers of A' . As indicated, the injection rate performance is improved from 23.7% in the symmetric mode ($A' = 5$) to 39.2% for $A' = 3$ and to as much as 97.5% for $A' = 1$ under a fully loaded condition. Therefore, the throughput performance can be improved dramatically even though the hardware cost increases by a factor

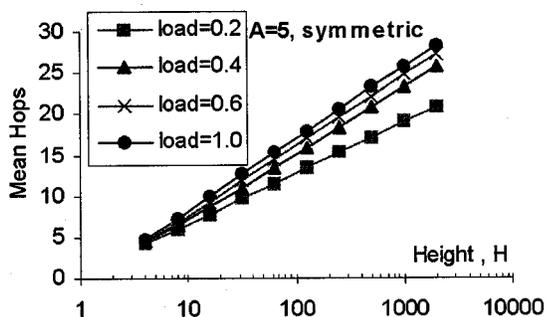


Fig. 4. Latency performance versus switch height H for different traffic loads. Symmetric mode with $A = 5$.

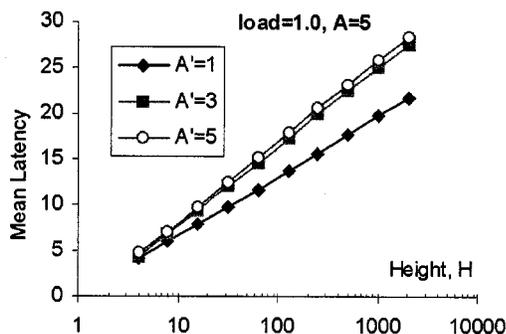


Fig. 5. Latency performances versus switch height for different asymmetric modes. Traffic load = 1.0.

of five given the same number of input ports. However, because each node only requires simple routing function and has potential in large-scale integration, the hardware implementation of the switching nodes can be potentially economic. As shown later, the asymmetric I/O operation also improves the mean latency performance.

B. Latency

In addition to the efficient scalability of the data vortex, the switch fabric also exhibits low latencies. The efficient routing schemes and small deflection penalty act to reduce the time of flight latency. We study the switch latency, defined as the mean number of hops propagated by the packets, as a function of the switch size and under various loads. Figs. 4 and 5 show the latency as a function of the switch height, H with $A = 5$ for various loads and various asymmetric modes. The x axis is plotted in a log scale; therefore, the mean latency is linearly proportional to $\log(H)$ as shown, and the slope of the lines depends on the scale of the input loads as well as the degree of asymmetry.

For the symmetric case (with $A = 5$) shown in Fig. 4, the overall latency performance is improved (i.e., smaller slope) with reduced load simply because the switch is less crowded and the average deflection probability is decreased.

In Fig. 5, A' out of A angles at the input side are active injecting ports. The load at each injecting port is set to 1.0. The more asymmetric cases (smaller A' for a given A) have lower overall traffic within the switch and, therefore, smaller latencies due to the reduced deflections. The absolute latency in the

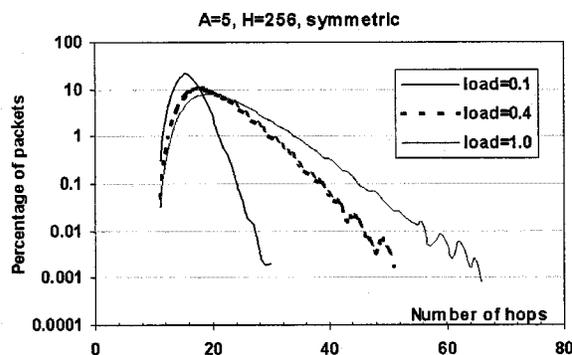


Fig. 6. Latency distribution under different loads with $A = 5$, $H = 256$ in a symmetric mode.

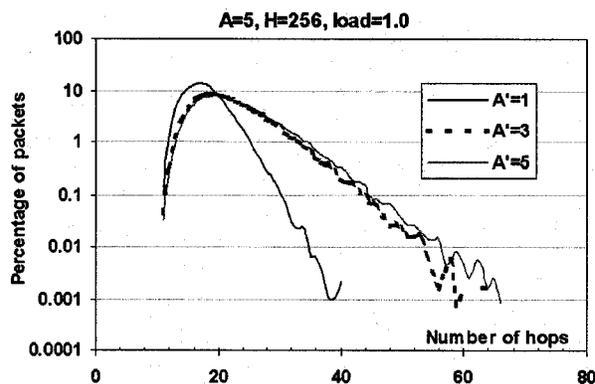


Fig. 7. Latency distribution under various asymmetric modes with load = 1.0.

data-vortex switch is low. For example, as shown in Fig. 5, in a switch fabric with 2-k input 10-k output ports under a fully loaded (load = 1.0) condition, the mean latency is only 22 hops. Because each node hop requires very little processing, the physical delay time at each hop can be very small, e.g., less than 10 ns (as will be shown in the experimental section). Thus, with the data-vortex architecture, one could construct a switch fabric with more than 2 k–10 k optical I/O ports and an overall packet latency in the 200-ns to 300-ns range.

C. Latency Distribution

The latency distribution of the switched packets is an important aspect of the system performance. It is easier for the receivers to restore the order of packets if their switching latencies are narrowly distributed around the mean values. At the physical layer, packets are also similarly degraded because they went through small variations of routing-path lengths. The latency distribution is shown in a histogram format with the percentage of the total number of packets plotted as a function of corresponding number of hops. The latency distribution is studied under various traffic loads and for different asymmetric modes. The results are shown in Figs. 6 and 7, respectively.

Fig. 6 shows the latency distributions in a symmetrically operated switch of $A = 5$, $H = 256$ under various traffic loads. Lighter traffic loads result in narrower distribution curves because packets are less likely to interfere with each other or accumulate large differences of deflection penalties. The x coordinate of the distribution curve peak, which also reflects the mean

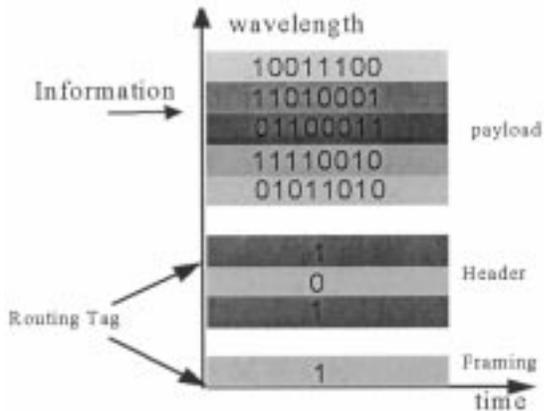


Fig. 8. WDM-TDM payload and WDM header.

latency of the system, is shifted to the left as the degree of load is reduced.

Fig. 7 compares the latency distributions for various asymmetric modes. For all cases in Fig. 7, the load for injecting ports is set to 1.0, and the same switch size is studied as in Fig. 6. For smaller A' , or the more asymmetric I/O case, a narrower distribution performance is achieved. This is simply due to the fact that fewer packets are injected into the switch, leading to a lower number of deflections and less variation in the number of deflections. We note that the improvement from $A' = 5$ to $A' = 3$ in latency distribution is minor compared with the $A' = 1$ case. In addition, we find that these results are similar for different switch sizes. Therefore, the overall latency of the data-vortex switch is considered narrowly distributed. For example, in a $A = 5, H = 2048$ switch, the 1% latency distribution tail is extended at a range of 40–50 hops, even for the worst traffic conditions.

IV. EXPERIMENTAL TEST BED

The throughput of the data vortex can be enhanced by using WDM to encode the packet payload. In addition, encoding the header bits by WDM can greatly simplify the routing strategy at the node. Because each of the cylinders decodes a specific header bit in a binary tree fashion, passive wavelength filtering can be implemented in the nodes to perform the routing operations. Using WDM in the header field also reduces the switch latency. Each incoming packet into the switch is then constructed as in Fig. 8. It consists of a payload field, i.e., the WDM data and a header field, which includes a framing bit to indicate the presence of a packet and the WDM header bits that are the destination address of the packet.

To experimentally demonstrate the routing functions of the data-vortex switch and measure the physical layer performance, we constructed a routing-node test bed [15], [16]. The test bed is configured in a recirculating loop fashion to emulate packets propagating through multiple node hops.

The experimental setup is shown in Fig. 9. Each routing node has two in (as the North and West) and two out (as the East and South) optical ports. At the packet generator, two continuous wave (CW) lasers LD1 (1547 nm) and LD2 (1550 nm) are used to carry framing and header information, and three other CW

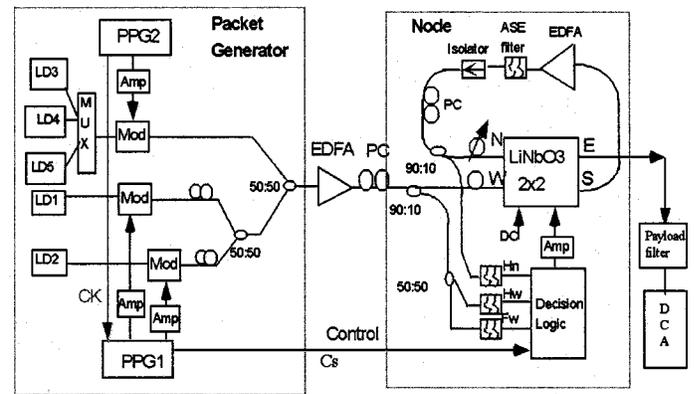


Fig. 9. Experimental setup. PPG: Pulse pattern generator. EDFA: erbium-doped fiber amplifier. PC: polarization controller. DCA: digital communication analyzer.

channels (1552 nm, 1553 nm, 1555 nm) are used to build the WDM payload. Electrical data of framing and header bits are generated from the pulse-pattern generator 1 (PPG1) HP81200 with a packet period of 32 ns. The modulated optical header and framing bits are aligned with the corresponding optical payload channels in time. The payload information is modulated at 10 Gb/s from the PPG2 (MP1763B, Anritsu Corp., Tokyo 106–8570, Japan), which includes a 4.8-ns guard-time band at the edge of packet boundary to allow for the switching transients. An erbium-doped fiber amplifier (EDFA) is used to compensate the loss through the modulating and multiplexing stages. A polarization controller (PC) is inserted at the input of the node to compensate for the polarization sensitivity of the switch.

The packet routing is realized by a LiNbO₃ 2 × 2 crossbar switch [17]. By properly applying the switching signal to the crossbar switch, the input optical packet can be directed to one of the two output ports. Based on this configuration, the electrical switching signal is derived from the routing information of the incoming packets [West's framing (Fw), West's header (Hw), and North's header (Hn)] and the control signal (Cs), which originates from the competing node at the neighbor cylinder. In the experimental setup, however, the control signal is programmed by PPG1 directly, and aligned with the corresponding routing information at the decision circuit. Fw, Hw, and Hn are tapped and filtered from the optical packets and converted into electronics feeding the decision circuit.

The recirculating configuration is achieved by connecting the output of the node (the South port) back as the North input. Within the loop, an EDFA with amplified spontaneous emission (ASE)-peak filtering is used to compensate the loss through the switches and through the taps and splices. The loop delay is designed to be five packets long to keep the routing synchronous. An optical delay line is inserted after the tap off, allowing the appropriate temporal alignment of the optical packets with the switching decision. Depending on the associated destinations and control conditions programmed, each packet will recirculate the node a varying number of times. A packet sequence entering the test bed will, therefore, emerge out of the original order because different packets travel a different number of hops through

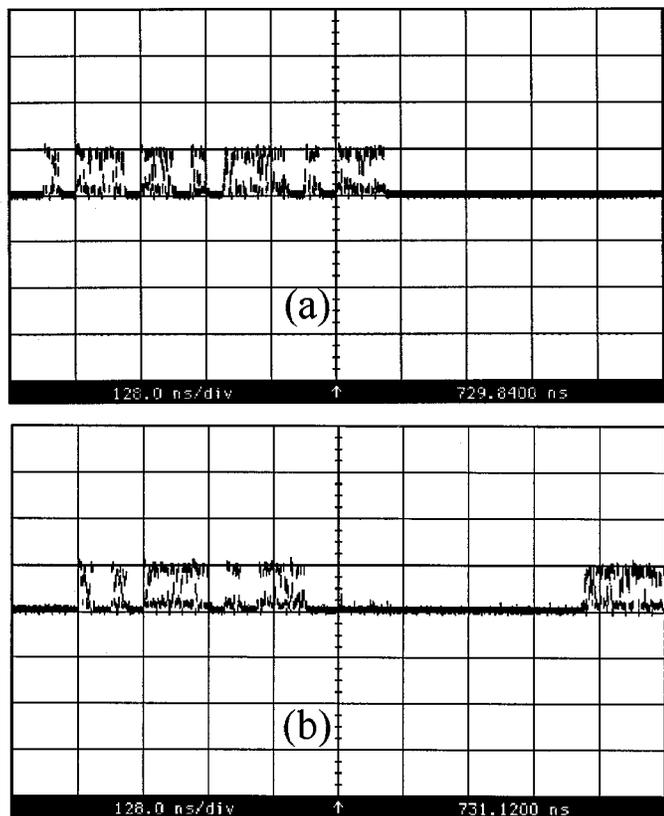


Fig. 10. (a) Input-packet sequence, four packets per division. “0010 1110 1101 0111 1010 1110 0000 0000 0000 0000”. (b) Reordered-output sequence, four packets per division. “0000 1010 1111 0101 1100 0000 0000 0000 0001 1111”.

the node. As an example, we tested a 40-packet sequence given by “0010 1110 1101 0111 1010 1110 0000 0000 0000 0000”, where “0” presents an empty packet slot, and “1” presents an existing packet slot. According to the specific header address and control for each packet, the output sequence is expected to be “0000 1010 1111 0101 1100 0000 0000 0000 0001 1111”. As shown in Fig. 10, the experimental results are consistent, demonstrating the correct routing function of the node.

To verify the successful multihop routings, we also examine the detailed payload of each packet from input to output. For example, the fifteenth packet in the input sequence is programmed to travel through six node hops and arrive as the fortieth packet at the output sequence. The end portions of this fifteenth-input packet and fortieth-output packet are shown to match in Fig. 11.

The eye diagram results of the filtered payload channel at 1555 nm are given in Fig. 12. The output shows the collective eye diagrams of all of the packets that have propagated through varying number of node hops. For a maximum of six node hops, we find clean and open eye diagrams. Payloads at the other two channels have similar results.

For practical large-scale switching fabrics, LiNbO₃ switch-based routing node will not be an economical implementation due to the difficulty of integration and a relatively large insertion loss. Therefore, semiconductor optical amplifier gate-based node with a similar switching speed is more likely to be the solution. In addition to integration advantage, it provides a higher ON-OFF ratio and is able to compensate the loss within the same

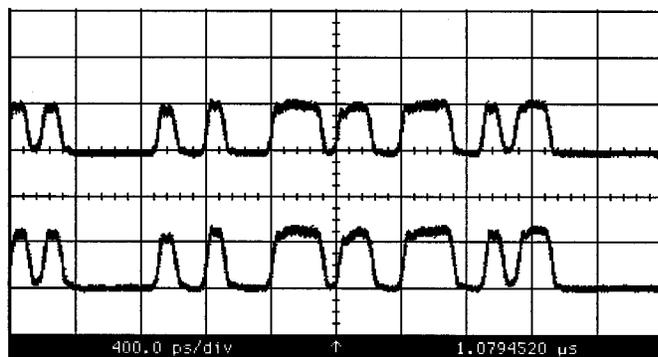


Fig. 11. Fifteenth input packet and fortieth output packet, four bits per division.

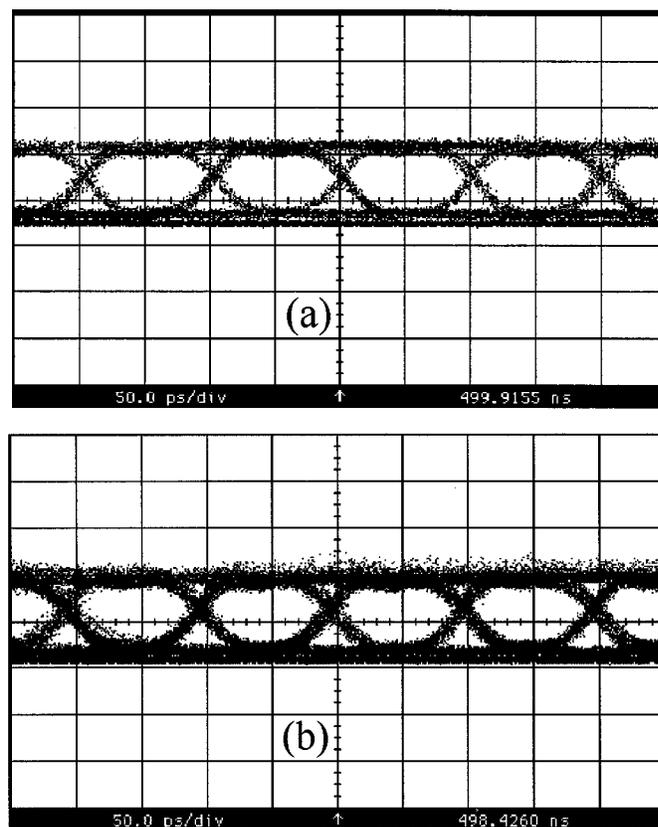


Fig. 12. (a) Input-packet eye diagram. (b) Eye diagram of collective-output packets.

device. A similar routing-node test bed is currently being built and under research. To better study the switching scalability, physical models are also being built for detailed studies of the noise and crosstalk accumulation in such multihop systems.

V. SUMMARY

We have investigated, in detail, the performance and experimental implementation of the data vortex, a new photonic packet-switch architecture. The simplified bufferless packet-switch design enables the exploitation of the high bandwidth and transparency afforded by optical transmission. The numerical studies show that the data vortex is highly scalable to large

port counts (>10 k) while maintaining low latencies (22 hops) and a narrow latency distribution. A wavelength-encoding method is additionally employed for the header bits, allowing the use of passive wavelength filtering for the routing operation. Under modest asymmetric I/O operation of the switch, sustained high switching throughput and a narrow latency distribution can be achieved. Therefore, for large-scale packet-switched application, it is feasible to build the data-vortex architecture with a slight penalty of the switching complexity. An experimental test bed was constructed to demonstrate the basic function of the routing node as well as successful routing performance for six-cascaded nodes.

REFERENCES

- [1] R. S. Tucker *et al.*, "Photonic packet switching: An overview," *IEICE Trans. Electron.*, vol. E82-C, pp. 202–212, Feb. 1999.
- [2] A. Pattavina *et al.*, "Techniques and technologies toward all-optical switching," *Opt. Networks*, vol. 1, Apr. 2000.
- [3] D. K. Hunter, M. C. Chia, and I. Andonovic, "Buffering in optical packet switches," *J. Lightwave Technol.*, vol. 16, pp. 2081–2094, Dec. 1998.
- [4] J. R. Sauer, "A multi-Gb/s optical interconnect," *SPIE, Digital Optical Computing II*, vol. 1215, p. 198, 1990.
- [5] F. Forghieri, A. Bononi, and P. R. Prucnal, "Analysis and comparison of hot-potato and single-buffer deflection routing in very high bit rate optical mesh networks," *IEEE Trans. Commun.*, vol. 43, pp. 88–98, Jan. 1995.
- [6] I. Chlamtac, A. Fumagalli, L. G. Kazovsky, P. Melman, W. H. Nelson, P. Poggiolini, M. Cerisola, A. N. M. M. Choudhury, T. K. Fong, R. T. Hofmeister, C.-L. Lu, A. Mekittikul, D. J. M. Sabido, C.-J. Suh, and E. W. M. Wong, "CORD: Contention resolution by delay lines," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 1014–1029, June 1996.
- [7] Y. Yamada, K. Sasayama, and K. Habara, "BER measurements in frequency-division-multiplexed optical loop buffer at data rate of 2.5 Gbit/s," *Electron. Lett.*, vol. 30, pp. 1869–1871, Oct. 27, 1994.
- [8] S. L. Danielsen, P. B. Hansen, and K. E. Stubkjaer, "Wavelength conversion in optical packet switching," *J. Lightwave Technol.*, vol. 16, pp. 2095–2108, Dec. 1998.
- [9] S. L. Danielsen, C. Joergensen, B. Mikkelsen, and K. E. Stubkjaer, "Optical packet switched network layer without optical buffers," *IEEE Photon. Technol. Lett.*, vol. 10, pp. 896–898, June 1998.
- [10] T. Durhuus, B. Mikkelsen, C. Joergensen, S. Lykke Danielsen, and K. E. Stubkjaer, "All-optical wavelength conversion by semiconductor optical amplifiers," *J. Lightwave Technol.*, vol. 14, pp. 942–954, June 1996.
- [11] P. Ohlen, "Noise and crosstalk limitation in optical crossconnects with reshaping wavelength converters," *J. Lightwave Technol.*, vol. 17, pp. 1294–1301, Aug. 1999.
- [12] A. Bononi, G. A. Castanon, and O. K. Tonguz, "Analysis of hot-potato optical networks with wavelength conversion," *J. Lightwave Technol.*, vol. 17, pp. 525–534, Apr. 1999.
- [13] C. Reed, "Multiple level minimum logic network," U.S. Patent 5 996 020, Nov. 30, 1999.
- [14] B. Yener, Y. Ofek, and M. Yung, "Topology design of loss-free switched-based LANs," in *Proc. IEEE INFOCOM'95*, 1995, pp. 88–96.
- [15] Q. Yang, M. F. Arend, G. D. Hughes, F. G. Johnson, and K. Bergman, "WDM/TDM optical packet switched network for supercomputing," in *Optics Computing 2000*, vol. 4089, 2000, pp. 555–561.
- [16] Q. Yang and K. Bergman, "WDM routing in photonic packet switch," in *LEOS 2000*, San Juan, Puerto Rico, 2000, MC2.
- [17] W. A. Payne and H. S. Hinton, "Design of lithium niobate based photonic switching systems," *IEEE Commun. Mag.*, vol. 25, no. 5, pp. 37–41, 1987.



Qimin Yang (S'99) was born in Ningbo, China, in 1972. She received the B.S. degree in electrical engineering from Zhejiang University, Hangzhou, China, in 1974 and the M.S. degree in electrical engineering from Beijing University of Posts and Telecommunications, Beijing, China, in 1997. She is currently pursuing the Ph.D. degree in optics and optoelectronic engineering at Princeton University, Princeton, NJ.

Her research interests include packet-switching architecture and high-speed WDM optical networks.



Keren Bergman (SM'87–M'93) received the B.S. degree in electrical engineering from Bucknell University, Lewisburg, PA, in 1998 and the M.S. and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, MA, in 1991 and 1994, respectively.

She joined Princeton University, Princeton, NJ, as an Assistant Professor in 1994, teaching both undergraduate and graduate courses in optics. Her research has included fundamental studies of quantum noise and nonlinear soliton propagation in

fibers, high bit-rate lasers, and ultrafast optical networks. She recently left the academic track to join the optical networking team at Tellium, developing the Aurora Full-Spectrum.

Dr. Bergman is a member of the Optical Society of America (OSA), for which she was recently elected Director at Large, and the IEEE Lasers and Electro-Optics Society (LEOS), for which she served on the board of governors from 1996 to 1999 and as the chair of the Princeton/Central NJ Chapter from 1998 to 2001; she is currently the Vice President for Membership (Americas). She received an AT&T Bell Laboratories Fellowship while a graduate student at MIT in 1991. She received the National Science Foundation CAREER Award in 1995, the ONR Young Investigator and Princeton's Wentz Award for excellence in teaching and scholarship in 1996, and the Caltech President's Award in 1997.



Gary D. Hughes works for the National Security Agency (NSA), Fort Meade, MD. He is the challenge leader responsible for computing research at the NSA. His work focuses on computer communications, including interconnects, switching, and networking.

Frederick G. Johnson received the B.A. degree from Oberlin College, Oberlin, OH, in 1987 and the Ph.D. degree from the Institute of Optics at the University of Rochester, Rochester, NY, in 1994. He was a researcher at the Laboratory for Physical Sciences, College Park, MD, from 1994 to 2001. He now works at Little Optics, Inc., Annapolis Junction, MD.