

# Performances of the Data Vortex Switch Architecture Under Nonuniform and Bursty Traffic

Qimin Yang, *Member, IEEE*, and Keren Bergman, *Member, IEEE*

**Abstract**—The Data Vortex switch architecture has been proposed as a scalable low-latency interconnection fabric for optical packet switches. This self-routed hierarchical architecture employs synchronous timing and distributed traffic-control signaling to eliminate optical buffering and to reduce the required routing logic, greatly facilitating a photonic implementation. In previous work, we have shown the efficient scalability of the architecture under uniform and random traffic conditions while maintaining high throughput and low-latency performance. This paper reports on the performance of the Data Vortex architecture under nonuniform and bursty traffic conditions. The results show that the switch architecture performs well under modest nonuniform traffic, but an excessive degree of nonuniformity will severely limit the scalability. As long as a modest degree of asymmetry between the number of input and output ports is provided, the Data Vortex switch is shown to handle very bursty traffic with little performance degradation.

**Index Terms**—Bursty traffic, Data Vortex switch architecture, photonic packet switch, switching fabric, uniform traffic.

## I. INTRODUCTION

TRANSMISSION of data in optical fibers has enabled the delivery of enormous bandwidth in today's communication networks, especially with new technology developments in dense wavelength division multiplexing (DWDM) and Raman fiber amplifiers [1]–[4]. The challenges in optical networking have recently migrated from transmitting high-capacity optical signals over long distances to effectively switching and managing that data [5]. These functions, currently performed in the electronic domain, have built an emerging large bottleneck to the scalability and growth of optical networks.

Photonic or optical switches present an attractive solution to the electronic bottleneck with promises of transparency, high capacity, and little electromagnetic interference (EMI). However, to deliver the necessary system performances, most existing switching architectures require intense processing and buffering, which can be realized easily in the electronic domain, but is still considerably challenging within the optical domain [6], [7]. Therefore, new switch architectures must be developed to accommodate the optical or photonic implementation, where the advantages of the optics such as bandwidth and transparency can be fully exploited while the disadvantages of the optics can be avoided.

Manuscript received May 31, 2001; revised March 29, 2002. This work was supported by the Defense Advanced Research Projects Agency and by the National Security Agency through an agreement with NASA under Contract 960199 and by the National Science Foundation under Contract ECS98-00401.

The authors are with the Lightwave Technology Laboratory, Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: qimin\_yang@hmc.edu).

Digital Object Identifier 10.1109/JLT.2002.800330

In [8], a multiple-level minimum-logic architecture, the Data Vortex, was proposed for large-scale low-latency packet-switching fabric. This novel architecture employs synchronous timing and distributed traffic-control signaling to avoid packet contention; therefore, the system achieves great simplicity, scalability, and high throughput. The hierarchical routing topology is carefully designed at each level so that packet deflection probability and deflection-induced latency are both minimized. The hierarchical routing procedure also facilitates the use of a wavelength-header-encoding technique to further simplify the routing function and the switching latency. Using DWDM within the data payloads further enhances high data throughput. In previous work, we have investigated the basic routing functionality, control signaling mechanism, and wavelength routing technique within an experimental test bed [9], [10]. In addition, our simulation results have shown that the Data Vortex packet switch achieves high scalability, low latency, and narrow latency distribution under uniform and random traffic conditions [11]. However, in practice, packet-switching systems are generally subject to nonuniformly distributed and/or bursty traffic. These factors may contribute additional congestion or deflections within the architecture; therefore, they may affect the switching latency performance and the throughput performance of the system. In this paper, we report on the robustness of the Data Vortex switch performances under these nonideal factors.

The rest of the paper is organized as follows. In Section II, an overview of the Data Vortex architecture is described. In Section III, we characterize the system performances in terms of successful injection probability, mean latency, and latency distribution. The performance results under different cases of nonuniform traffic and bursty traffic will be discussed and compared. Finally, we present our conclusions in Section IV.

## II. ARCHITECTURE OVERVIEW

The Data Vortex switching topology can be arranged as a collection of richly connected routing nodes on multiple fiber cylinders, as seen in Fig. 1. The switch fabric size is characterized by two parameters  $A$  and  $H$ , representing the number of nodes along the *angle* and *height* dimensions, respectively.  $A$  is typically set to be a small odd number ( $< 10$ ), and is independent of the choice of  $H$ . The available number of input/output (I/O) ports is given by  $H \times A$ . The number of cylinder levels  $C$  scales as  $C = (\log_2 H + 1)$ . In Fig. 1, a switch fabric of  $(A, H) = (5, 8)$  is shown with a top view of the routing tours and a side view of the interconnection patterns at each of the  $C = 4$  cylinders. Each cross point shown is a routing node, which can be labeled uniquely by the coordinates  $(a, c, h)$ , where  $0 \leq a < A$ ,  $0 \leq c < C$ , and  $0 \leq h < H$ .

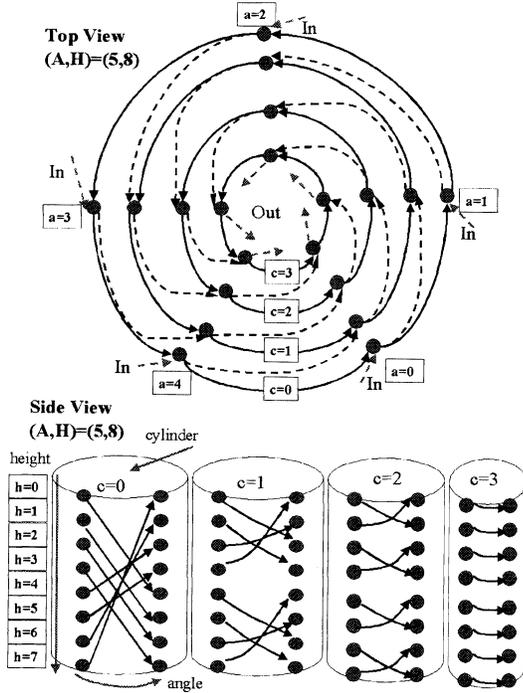


Fig. 1. Data Vortex topology  $(A, H) = (5, 8)$  with routing tours seen from the top and the side. Each node is labeled uniquely by the coordinate  $(a, c, h)$ , where  $0 \leq a < A$ ,  $0 \leq c < C$ , and  $0 \leq h < H$ .

Packets are processed synchronously in a highly parallel manner. Each packet is of a fixed length, and is routed in a slotted manner. Within each clock cycle, every packet in 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 paths along the same cylinder are shown in Fig. 1 from the side view of each cylinder. These connection patterns are carefully designed and repeat from angle to angle to minimize the packet deflection probability. At a specific cylinder  $C$ , a node  $(a, c, h)$  is connected to the node  $((a + 1) \bmod A, c, h')$  where the conversion from  $h = (h_{C-2}, \dots, h_{C-2-i}, \dots, h_1, h_0)$  to  $h' = (h'_{C-2}, \dots, h'_{C-2-i}, \dots, h'_1, h'_0)$  can be obtained as follows (where  $0 \leq i \leq C - 2$ ):

**For**  $(0 \leq c \leq C - 2)$

{

**If**  $h_{C-2-c} = 0$

**Then**

$$h' = (h_{C-2}, \dots, \overline{h_{C-2-c}}, h_{C-3-c}, \dots, h_1, h_0)$$

**Else if**

$$h_{C-2-c} \cap h_{C-3-c} \cap \dots \cap h_{j+1} = 1 \quad \text{and} \quad h_j = 0$$

where  $-1 \leq j \leq C - 3 - c$

**Then**

$$h' = (h_{C-2}, \dots, \overline{h_{C-2-c}}, \overline{h_{C-3-c}} \dots \overline{h_{j+1}}, \overline{h_j}, h_{j-1}, \dots, h_1, h_0)$$

}

**For**  $(c = C - 1)$

$$\{h' = h = (h_{C-2}, \dots, h_1, h_0)\}$$

We set a dummy parameter  $h_{-1} = 0$  for convenience. The formula can be easily verified in the example for  $H = 8$  shown in Fig. 1. The dashed-line paths between neighboring cylinders maintain the same height index  $h$  as they are used to forward the packets.

As shown, 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 self-routed in the fashion of binary-tree decoding as it propagates from the outer cylinder toward the inner cylinder. Every cylinder progress fixes a specific bit within the binary header address. The innermost cylinder ( $c = \log_2 H$ ) also allows the packet to circulate around when the output buffers are busy. To avoid packet contention, the switching architecture employs a synchronous and distributed control mechanism to properly schedule the neighboring packet flow [9]. As a result, each node encounters at most one packet at a time and no optical buffering will be necessary within the Data Vortex switch fabric. This also greatly simplifies the routing procedure at each hop and facilitates the photonic implementation of the architecture. Although packet deflection occurs due to the need of traffic control, the probability of that event and its incurred latency penalty are minimized. This is achieved because packets are provided multiple paths to the destination and the topology provides as low as two hops of latency penalty in the case of deflection. Since packets are always allowed to stay on the same cylinder if they are deflected, the ‘‘angle’’ dimension virtually provides a buffering mechanism for the packets while eliminating the potential packet conflict. As shown in later sections, this design is essential in sustaining promising system performances under various traffic conditions. A more detailed description of the switch architecture and this control mechanism is provided in [8]–[11].

### III. PERFORMANCE

In previous work, the Data Vortex performance has been studied under uniform and random traffic conditions. In those cases, each input port has the same probability (i.e., offered traffic load) of receiving a new packet. This probability is independent of whether or not there have been packets at previous clock cycles or at other input ports. At the same time, we assume the destinations of the incoming packets have a uniform distribution to the available output ports. System performances are evaluated in terms of successful injection probability, mean latency, and latency distribution, which can be derived statistically from a packet-flow model of the architecture [10], [11]. The results have shown that under uniform and random traffic conditions, the Data Vortex packet switch is able to scale to very large I/O port counts ( $>10$  K) while maintaining a reasonable fraction of successful injections, low mean latency, and narrow latency distribution. By choosing part of the available angles as active input ports, an asymmetric I/O mode is usually preferred during the operation because it

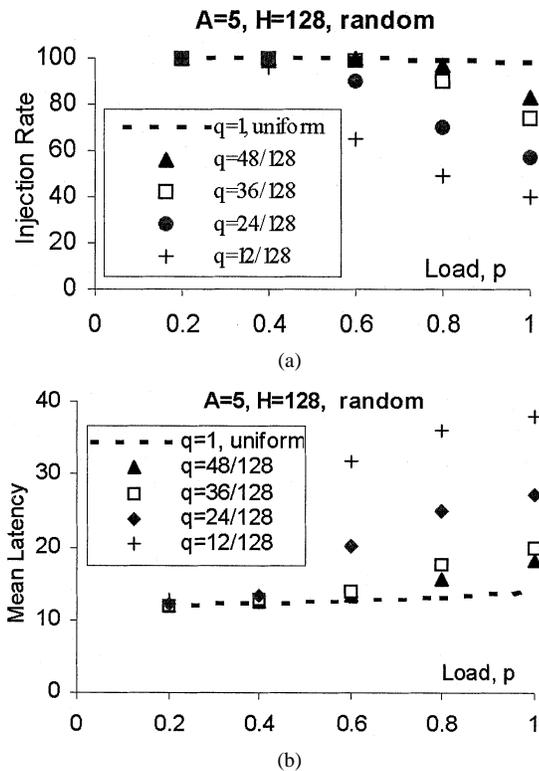


Fig. 2. (a) Successful injection probability versus load for various nonuniform distributions. The switch fabric size is  $128 \times (5 \times 128)$ . (b) Mean latency versus load for various nonuniform distributions. The switch fabric size is  $128 \times (5 \times 128)$ .

significantly improves the system performances with a modest increase in the system complexity.

In practical systems, however, the incoming traffic tends to be nonuniformly distributed and/or bursty. These factors can contribute to additional congestion or deflections within the architecture. It is, therefore, important to study the architecture performance under these realistic conditions and to evaluate its robustness under different traffic parameters.

#### A. Nonuniform and Random Traffic

To model the traffic load at each input port, a load parameter  $p$  is used to characterize the probability of packet injection into a port. Packet injections are still random and assumed to be independent of the injections at other clock cycles or at other input ports. The output target height is encoded in binary as the packet destination address. To represent a nonuniform distribution, a specific output address or a group of addresses are chosen as the packet targets throughout the traffic flow simulation. To characterize the degree of nonuniformity, we define a parameter  $q$  that is given by the ratio of the chosen address group size over the switching fabric height  $H$ . Within the chosen group of output targets, the probability of reaching each of them is assumed to be the same. Given this nonuniform target addressing, packet congestion may form in some areas of the switching fabric and may result in a deteriorated latency and throughput performance.

In Fig. 2(a) and (b), the successful injection probability and mean latency performances are shown, respectively, for a Data Vortex packet switch with  $A = 5$  and  $H = 128$  under random

injections. The successful injection probability is the ratio of successful injection attempts over all injection attempts at the input ports. The mean latency is the average number of hops that a packet travels through the multilevel switching fabric. Another important factor, the asymmetric I/O mode, is defined by a parameter  $\theta = A_{in}/A_{out} = A_{in}/A$ , as the ratio of the active input angles over the active output angles. Here, we assumed that all the available angles at the output side are active receiving ports. The I/O mode in this case is  $\theta = 1/5$ , therefore, the switch fabric size is given by  $128 \times (5 \times 128)$  with 128 input ports and  $(5 \times 128)$  output ports. As shown in Fig. 2, the performances are shown as functions of the per-port load  $p$  under various degrees  $q$  of nonuniform target distribution. The uniform case ( $q = 128/128 = 1$ ) is shown as a reference for comparison.

As the results show, the Data Vortex switch fabric performs well under a modest degree of nonuniform distribution ( $q > 30\%$ ), and only very small degradations are observed while the system is lightly loaded ( $p < 0.4$ ). Both the successful injection probability and the mean latency degrade further as the nonuniform degree increases (i.e.,  $q$  decreases) and as the traffic load  $p$  increases. As mentioned above, the topology of the Data Vortex switch fabric provides a virtual buffering mechanism within the angle dimension. The routing redundancy provided is able to smooth out the traffic flow before the nonuniform traffic-induced congestion accumulates. However, this only works effectively if the system is lightly loaded or if the degree of the nonuniform distribution is not excessive. As the traffic load increases, especially as the degree of nonuniformity increases beyond the limit ( $q < 20\%$ ), the redundancy provided by the topology cannot sustain the same degree of smoothing effect toward the varying traffic densities. As a result, the probability of packet deflection starts to increase considerably and packets tend to accumulate latency hopping within the switching fabric. As more packets are accumulated inside the switch, it is harder to inject new packets successfully at the input side, therefore, the successful injection probability or the throughput of the system degrades as well.

In Fig. 3(a) and (b), the successful injection probability and the mean latency performance measures are shown, respectively, for different switch heights under a load = 1.0 condition. The same angle parameter  $A = 5$  and same I/O mode  $\theta = 1/5$  are used, and the switch fabric size is given by  $H \times (A \times H)$ . Performances are compared for the same degree of nonuniform distribution, i.e., as a function of  $q$ . The mean latency performance is normalized to the number of the cylinder levels  $C$  for comparison among different switch sizes. As the results show, smaller switch fabrics perform slightly better than the larger switch fabrics because of the stronger smoothing effect on the traffic distribution provided in the smaller switches. However, as the traffic becomes closer to the uniform distribution ( $q \rightarrow 1$ ), the differences become smaller because the performances are less dependent on the smoothing effect of the architecture when it is subject to uniformly distributed traffic.

#### B. Bursty and Uniform Traffic

In practical systems, traffic is also likely to be bursty since multiple packets injected in consecutive slots will tend to be targeted for the same destination. These consecutive slots are

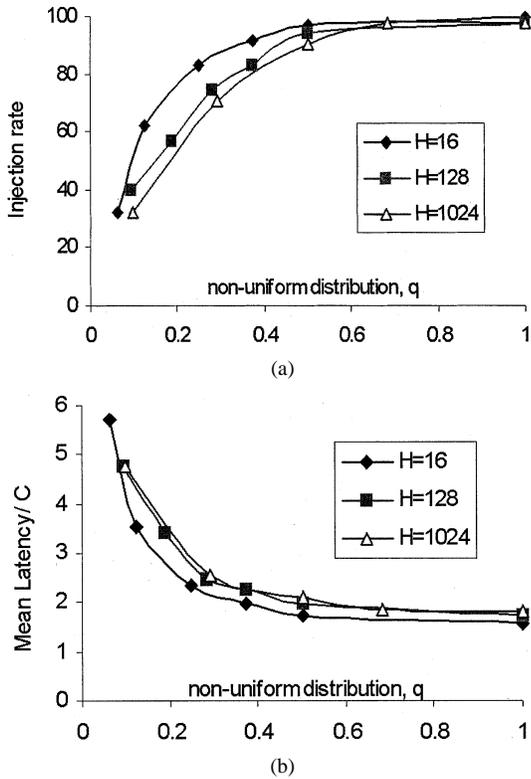


Fig. 3. (a) Successful injection probability versus nonuniform distribution for various switch sizes. The switch fabric size is  $H \times (5 \times H)$ . (b) Mean latency versus nonuniform distribution for various switch sizes. The switch fabric size is  $H \times (5 \times H)$ .

considered as an active ON period at the input port, alternating with an OFF period in which multiple time slots are idle. The length of each period has a certain distribution characterizing the burstiness of the traffic pattern. To simplify the modeling, we focus here on a uniform distribution for the output targets, in which each new burst has the same probability of targeting any of the available  $H$  output heights.

To simulate the bursty traffic, each ON/OFF period is modeled by [12]

$$\begin{aligned} T_{\text{on}} &= \left\lfloor \frac{1}{U^{1/\alpha_{\text{on}}}} \right\rfloor \\ T_{\text{off}} &= \left\lfloor \frac{1}{U^{1/\alpha_{\text{off}}}} \right\rfloor \end{aligned} \quad (1)$$

where  $U$  is a random variable uniformly distributed over  $[0, 1]$ , and  $\lfloor x \rfloor$  indicates the floor function. The parameters  $\alpha_{\text{on}}$  and  $\alpha_{\text{off}}$  specify the variability of the ON and OFF time periods. For infinite long period,  $T_{\text{on}}$  and  $T_{\text{off}}$  follow the rounded Pareto distributions. The resulted traffic ranges from rather smooth cases ( $\alpha \gg 2$ ) to highly bursty cases ( $1 < \alpha < 2$ ), which in general contain much longer burst periods. Since the long OFF burst period case ( $1 < \alpha_{\text{off}} < 2$ ) leads to a light traffic load where the degradation in the performance is much less an issue, we only vary the parameter  $\alpha_{\text{on}}$  between different ranges of burstiness while setting the parameter  $\alpha_{\text{off}}$  to generate short OFF periods. Each input port is modeled independently. Different bursts generated are independent of each other, and have a uniform distribution to any of the available target addresses. Traffic loads are averaged over different input ports during the total simulation time.

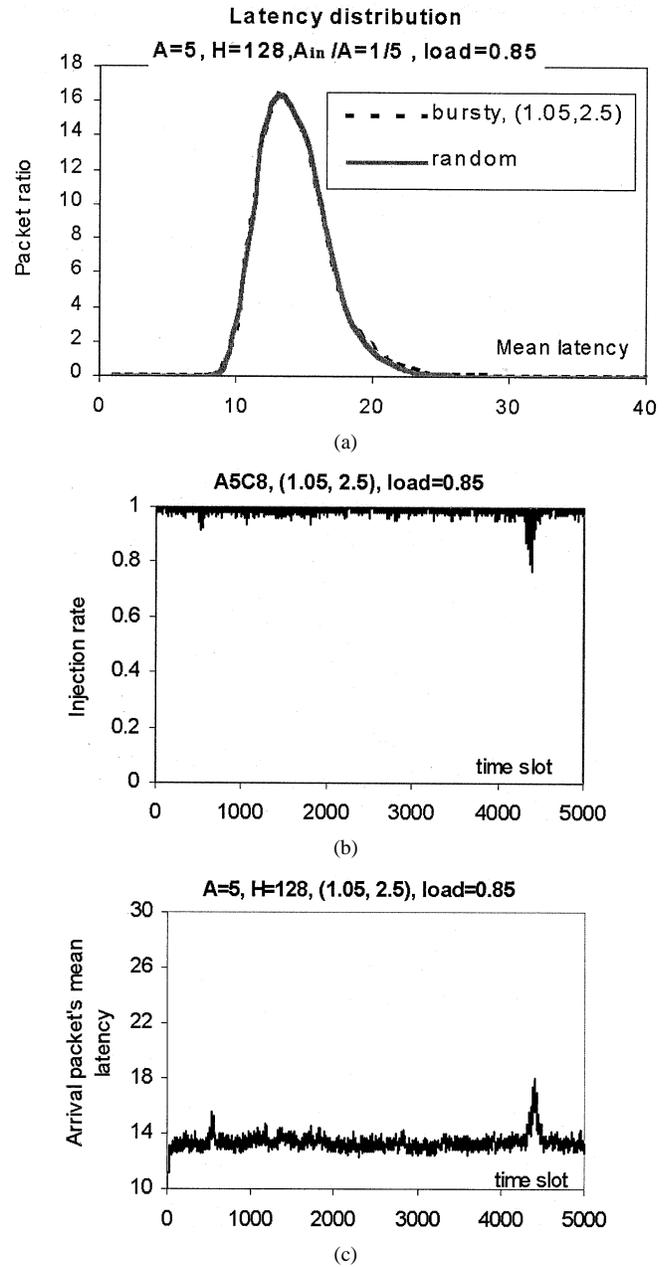


Fig. 4. (a) Latency distribution for ( $\alpha_{\text{on}} = 1.05$  and  $\alpha_{\text{off}} = 2.5$ ). The switch size is  $128 \times (5 \times 128)$ . (b) Successful injection probability during the time slots for ( $\alpha_{\text{on}} = 1.05$  and  $\alpha_{\text{off}} = 2.5$ ). The switch size is  $128 \times (5 \times 128)$ . (c) Mean latency of arrival packets during the time slots for  $\alpha_{\text{on}} = 1.05$  and  $\alpha_{\text{off}} = 2.5$ . The switch size is  $128 \times (5 \times 128)$ .

The bursty system performance measures are studied under various operating conditions, which include different bursty parameters ( $\alpha_{\text{on}}$ ,  $\alpha_{\text{off}}$ ) and different asymmetric I/O modes ( $\theta$ ). Since the traffic varies over different time periods, we need to monitor the system performance during each clock cycle to investigate the bursty traffic effect. In order to observe the heavy tails of the traffic distribution, we need to measure the network performance over a very long period of time.

For the first case, a very bursty pattern with  $\alpha_{\text{on}} = 1.05$  and  $\alpha_{\text{off}} = 2.5$  is studied. Given the simulating run of 5000 time slots, the averaged ON period in this case is about six packet slots and the averaged OFF period is about 1.3 packet slots. The resulting traffic load is approximately 0.85 at the input ports. In

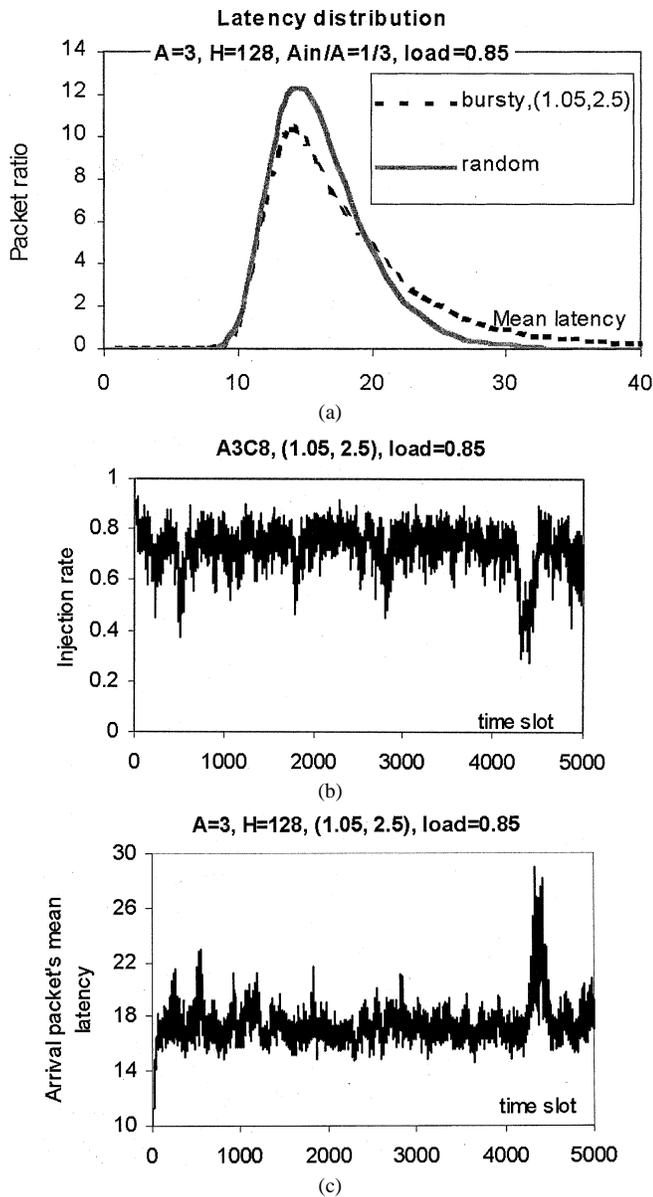


Fig. 5. (a) Latency distribution for  $\alpha_{on} = 1.05$  and  $\alpha_{off} = 2.5$ . The switch size is  $128 \times (3 \times 128)$ . (b) Successful injection probability during the time slot for  $\alpha_{on} = 1.05$  and  $\alpha_{off} = 2.5$ . The switch size is  $128 \times (3 \times 128)$ . (c) Mean latency of arrival packets during the time slot for  $\alpha_{on} = 1.05$  and  $\alpha_{off} = 2.5$ . The switch size is  $128 \times (3 \times 128)$ .

Fig. 4(a), the latency distribution is shown for the I/O mode of  $\theta = 1/5$  with  $A = 5$  and  $H = 128$  (dashed line). It is compared to the random case (solid line) in which each time slot is statistically independent and the same traffic load of 0.85 is used. As the results show, there is almost no degradation caused by the traffic burstiness. The monitored performance measures over the time slots illustrate more clearly how the switching fabric behaves under the bursty traffic. In Fig. 4(b) and (c), the successful injection probability and packet mean latency is shown, respectively, at each time slot. There is slight performance degradation around the 500th time slot and 4500th time slot because more ON bursts are overlapped during the period in this specific case. In general, however, the performance variation is very small over the time due to the smoothing effect provided by the switching fabric topology.

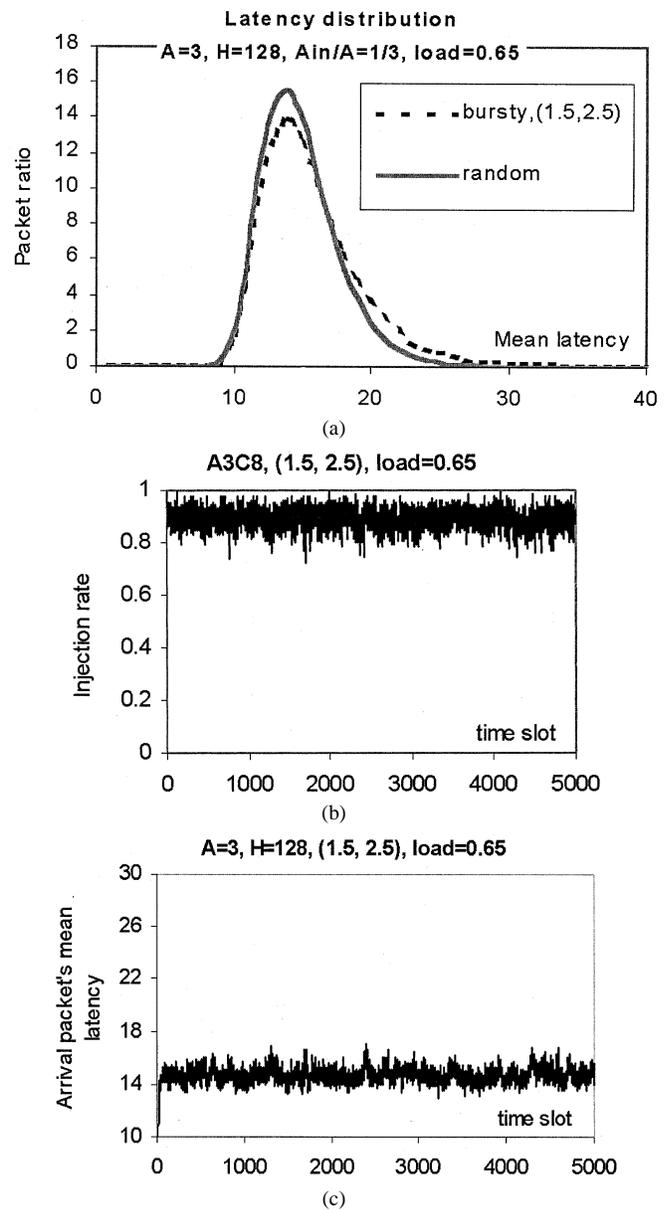


Fig. 6. (a) Latency distribution for  $\alpha_{on} = 1.5$  and  $\alpha_{off} = 2.5$ . The switch size is  $128 \times (3 \times 128)$ . (b) Successful injection probability during the time slot for  $\alpha_{on} = 1.5$  and  $\alpha_{off} = 2.5$ . The switch size is  $128 \times (3 \times 128)$ . (c) Mean latency of arrival packets during the time slot for  $\alpha_{on} = 1.5$  and  $\alpha_{off} = 2.5$ . The switch size is  $128 \times (3 \times 128)$ .

Thus, the asymmetric I/O mode is critical in providing enough smoothing effect for the bursty traffic. For comparison, an example of  $\theta = 1/3$  with  $A = 3$  and  $H = 128$  is studied under the same degree of traffic burstiness ( $\alpha_{on} = 1.05$ ,  $\alpha_{off} = 2.5$ ). As shown in Fig. 5(a), with  $\theta = 1/3$ , a much wider latency distribution and longer tail results in comparison with the random case. Therefore, if an insufficient asymmetry is provided between the number of input and output ports, the switch performance can be degraded severely under bursty traffic.

Similarly, the performance measures over the time slots illustrate the degradations more clearly. As shown in Fig. 5(b) and (c), both the successful injection probability and the mean latency fluctuate widely over time, especially during the period around the 500th and the 4500th slots. The degradations are

directly proportional to the number of input ports, which overlap their long ON period, and to the length of the overlapped period. Thus, compared with the  $\theta = 1/5$  case, the I/O asymmetry of  $\theta = 1/3$  provides less routing redundancy and, therefore, less smoothing effect on the traffic congestion.

Given the same degree of asymmetric I/O mode, different bursty modes will result in different degrees of performance degradation. For comparison, a second bursty mode with  $\alpha_{\text{on}} = 1.5$  and  $\alpha_{\text{off}} = 2.5$  is studied under the I/O asymmetry of  $\theta = 1/3$ . The resulting latency distribution curve, successful injection probability, and mean latency over the time slots are shown in Fig. 6(a)–(c), respectively. In the second bursty case, the averaged ON period is about 2.6 packet slots and the OFF period is kept at near 1.3 packet slots. The offered traffic load is approximately 0.65. Due to a smoother input traffic pattern, there are less long ON periods that overlap and the successful injection probability and mean latency fluctuate less over time compared with the case shown in Fig. 5. Since the overall traffic load is also smaller in this case, the smoothing effect is more effective and the overall performance degradation is diminished. As shown in Fig. 6(a), the latency distribution curve shows a relatively small tail compared to the random case of the same traffic load. Therefore, under a modest bursty traffic, the switch performance degrades only slightly even if a small I/O asymmetry is provided.

In addition, we find that the above performance comparisons for bursty traffic are quite independent of the switch size. This indicates that as long as a modest I/O asymmetry ( $\theta = 1/5$ ) is provided, the switch scalability will not be affected by the bursty nature of input traffic.

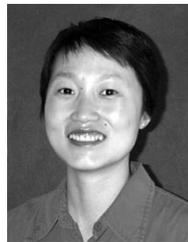
#### IV. CONCLUSION

In summary, we have studied the performances of the Data Vortex switch under nonuniform and bursty traffic conditions. The results show that the architecture generally maintains robust throughput and latency performance under these factors due to the inherent traffic smoothing effect. Excessive nonuniform target distribution ( $q < 20\%$ ), however, will result in significant degradation in the successful injection probability and latency distribution, especially under heavily loaded conditions. If subjected to the bursty input traffic, the switch fabric scalability is maintained well with little performance degradation as long as a modest asymmetric I/O mode ( $\theta = 1/5$ ) is provided.

#### REFERENCES

- [1] G. Wellbrock, D. Pitchforth, M. Hinds, and S. Chi, "Terabit field trial (100×10 Gb/s over 4-span NDSF route)," in *Proc. Optical Fiber Communication Conf. (OFC'2000)*, vol. 4, Mar. 2000, pp. 296–298.
- [2] S. Bigo, A. Bertaina, M. W. Chbat, S. Gurib, J. DaLoura, J.-C. Jacquinet, J. Hervo, P. Bousselet, S. Borne, D. Bayart, L. Gasca, and J.-L. Beylat, "320-Gb/s (32 × 10 Gb/s WDM) transmission over 500 km of conventional single-mode fiber with 125-km amplifier spacing," *IEEE Photon. Technol. Lett.*, vol. 10, pp. 1045–1047, July 1998.
- [3] H. Ono, M. Yamada, T. Kanamori, S. Sudo, and Y. Ohishi, "1.58- $\mu\text{m}$  band gain-flattened erbium-doped fiber amplifiers for WDM transmission systems," *IEEE J. Lightwave Technol.*, vol. 17, pp. 490–496, Mar. 1999.

- [4] K. S. Jepsen *et al.*, "Network demonstration of 32  $\lambda \times$  10Gb/s across 6 nodes of 640×640 WSXCs with 750-km Raman-amplified fiber," in *Proc. Optical Fiber Communication Conf. (OFC'2000)*, vol. 4, Mar. 2000, pp. 272–274.
- [5] "Special Issue on Optical Networks," *IEEE J. Lightwave Technol.*, vol. 18, Dec. 2000.
- [6] D. K. Hunter, M. C. Chia, and I. Andonovic, "Buffering in optical packet switches," *J. Lightwave Technol.*, vol. 16, pp. 2081–2094, Dec. 1998.
- [7] C. 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. Mekkittikul, D. J. MN. Sabido, IX, 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.
- [8] C. Reed, "Multiple level minimum logic network," U.S. Patent 5996020, Nov. 30, 1999.
- [9] Q. Yang and K. Bergman, "Traffic control and WDM routing in the data vortex packet switch," *IEEE Photon. Technol. Lett.*, vol. 14, pp. 236–238, Feb. 2002.
- [10] Q. Yang, M. F. Arend, G. D. Hughes, F. G. Johnson, and K. Bergman, "WDM/TDM optical packet switched network for supercomputing," in *Proc. SPIE Optics in Computing 2000*, vol. 4089, pp. 555–561.
- [11] Q. Yang, K. Bergman, G. D. Hughes, and F. G. Johnson, "WDM packet routing for high-capacity data networks," *J. Lightwave Technol.*, vol. 19, pp. 1420–1426, Oct. 2001.
- [12] A. Bononi, L. Tancevski, and L. A. Rusch, "Large power swings in doped-fiber amplifiers with highly variable data," *IEEE Photon. Technol. Lett.*, vol. 11, pp. 131–133, Jan. 1999.



**Qimin Yang** (M'02) was born in Ningbo, China, in 1972. She received the B.S. degree from Zhejiang University, Hangzhou, China, in 1994 and the M.S. degree from Beijing University of Posts and Telecommunications, Beijing, China, in 1997, both in electrical engineering, and the Ph.D. degree in optics and optoelectronic engineering from Princeton University, Princeton, NJ, in 2002.

She was a Postdoctoral Research Associate with the Lightwave Research Laboratory, Columbia University, New York, NY. Currently, she is with Harvey

Mudd College, Claremont, CA. Her research interests are in packet-switching architecture and high-speed WDM optical networks.



**Keren Bergman** (S'87–M'93) received the B.S. degree from Bucknell University, Lewisburg, PA, and the M.S. and Ph.D. degrees from the Massachusetts Institute of Technology (MIT), Cambridge, all in electrical engineering, in 1988, 1991, and 1994, respectively.

She joined Princeton University, Princeton, NJ, as an Assistant Professor in 1994. In 2000, she joined the optical networking team at Tellium, developing a MEMS-based optical cross connect. Since January 2002, she has been an Associate Professor in the Department of Electrical Engineering at Columbia University, New York, NY. Her research interests range from fundamental studies of quantum noise and non-linear soliton propagation in fibers to high-bit-rate lasers and ultrafast optical networks.

Dr. Bergman is a Member of the Lasers and Electro-Optic Society (LEOS), has served as Chair of the Princeton/Central New Jersey Chapter (1996–1999) and the Board of Governors (1998–2001), and is the current Vice President for Membership (Americas). She is also active with the Optical Society of America (OSA) and has served on several program committees. She was elected Director at Large for the OSA in 2000. She was awarded an AT&T Bell Laboratories Fellowship in 1991 while at MIT. She received the National Science Foundation CAREER Award in 1995, and the ONR Young Investigator Award in 1996. She received the President's Award from the California Institute of Technology, Pasadena, in 1997, for her work on optical networks with the Jet Propulsion Laboratory.