

Load-Balanced Optical Switch for High-Speed Router Design

Bing Hu and Kwan L. Yeung
Dept. of Electrical and Electronic Engineering
The University of Hong Kong
Hong Kong, PRC
E-mail: { binghu, kyeung }@eee.hku.hk

Abstract — A load-balanced (electronic) switch architecture is attractive because it does not require a central scheduler, and provides close to 100% throughput. Although the switch fabrics in a load-balanced switch are amenable to optics (e.g. using AWGR), packets must be electrically handled/buffered between the two switch fabrics (i.e. at middle-stage ports). This makes the all-optical packet transmission from an input linecard to an output linecard impossible. In this paper, a very simple-yet-effective all-optical implementation of a load-balanced switch is proposed. It comprises N linecards connected by an N -wavelength WDM fiber ring. Each linecard i is configured to receive on channel λ_i . To send a packet, it can select and transmit on an idle channel based on where the packet goes (e.g. using λ_j if the packet goes for linecard j). The ring network is engineered such that the amount of time a packet should be buffered at a middle-stage port exactly matches the propagation delay that this packet experiences en route. This intelligently solves the problem of lacking optical buffers. Following a pre-determined sending schedule, we show that our proposed load-balanced optical switch can give an unbeatable delay and throughput performance under different traffic patterns.

Keywords—Load-balanced optical switch, optical buffer, fiber ring network

I. INTRODUCTION

With the continuous growth of bandwidth in fiber links and the ever increasing Internet traffic, the need for building high speed routers is urgent in order to keep pace with the increased transmission rate. A major bottleneck in high-speed router design is its switch fabric, as the fabric needs to handle the aggregated traffic loads coming from all the input ports. Input-queued switch is preferred over output-queued switch for its scalability. But an input-queued switch generally requires a sophisticated central scheduler [1-2]. Numerous efforts have then been made to reduce the scheduler complexity. Various hardware designs have also been proposed for simplifying the scheduler implementation, and some of them have been incorporated into real products. The general technique adopted is to trade extra communication overhead for less computational requirements [2].

To this end, the load-balanced switch [3] is very attractive because it does not require a central scheduler, and its two stages of switch fabrics (see Fig. 1) can be very simple. Indeed, each switch fabric in the load-balanced switch only needs to realize N switch configurations, whereas a non-blocking input-queued switch [1-2] needs to realize the entire $N!$ possible configurations. But the load-balanced switch also faces its own

challenge: in-order packet delivery cannot be guaranteed due to the variable delay a packet can experience at the intermediate buffer between the two switch fabrics, i.e. middle-stage ports. Among various efforts in solving the packet mis-sequencing problem, the feedback-based two-stage switch architecture [4] provides an elegant solution while not sacrificing the switch's throughput-delay performance.

On the other hand, to support the increasing number of ports and higher data rates on each port, the clock frequency for electronic switch fabrics must be substantially increased. However, by increasing the operational frequency, the inherent electromagnetic compatibility, power density problems, and the layout complexity, as well as larger power consumption and dissipation issues, make the electronic switch fabrics unattractive for high-speed switch design [5-9]. As a result, the use of optical technologies for switching is gaining more interests, both in the research and industrial communities. Indeed, employing optical technologies for switching presents interesting aspects [5-6]: huge available bandwidth, reduced power consumption and dissipation, intrinsic flexibility in supporting different interconnection topologies and especially, a wavelength switching cost quite independent from the data bit-rate (a big contrast to that in the electronic domain).

Therefore, in high-speed router design we should not only address the performance bottleneck due to central scheduler, but also make use of the state-of-art optical technologies for switch fabric design. Notably, in the 100Tb/s router project [5], optical implementation of a load-balanced switch [3] was actively considered. A hybrid electro-optic switch fabric based on a three-stage Clos network architecture with optical MEMS [8] is used. But all-optical packet transmission from an input port/linecard to an output port/linecard is not possible, as packets must be temporarily stored and processed in electrical domain inside the hybrid switch. Besides, to tackle the packet mis-sequencing problem a large re-sequencing buffer of N^2+1 packets at each output port is required. This significantly prolongs the packet delay. Another good attempt [9] is to implement the feedback-based two-stage switch [4] using AWGR (Arrayed Waveguide Grating Router [7]). Despite of its much simpler scheduler [4], packets must still be electrically handled/buffered at middle-stage ports.

Recently, Fasnet [6], an optical switch fabric comprising N switch linecards connected by two counter-rotating WDM fiber rings, was proposed. In Fasnet, one ring is used for transmission,

while another is for reception. The N wavelengths in transmission ring are switched to the reception ring at a folding point between the two rings. Only a special input port (called master input) can generate frame headers (called locomotives). Other input ports put their packets at the end of the frames as frame headers pass by. This switch architecture does realize all-optical packet transmission from one linecard to another. But the throughput-delay performance of Fasnet is rather limited [6]. (Please also refer to our simulation results in Section IV.)

In this paper, a very simple-yet-effective optical implementation of a load-balanced switch is proposed. In our design, N linecards are connected by a single N -wavelength WDM fiber ring. Each linecard i is configured to receive on channel λ_i . To send a packet, it can select and transmit on an idle channel based on where the packet goes (i.e., using λ_j if the packet goes for linecard j). The packet will move along the ring and arrive at the destination linecard without being electrically buffered/processed en route. The ring network is engineered such that the amount of time a packet should be buffered at a middle-stage port of a load-balanced electronic switch exactly matches the propagation delay that this packet experiences in the fiber ring. Following a pre-determined sending schedule, we show that the WDM ring network functions as the optical counterpart of the load-balanced (electronic) switch [4].

The rest of the paper is organized as follows. In the next section, we recap the operation of the load-balanced electronic switch [4]. In Section III, the load-balanced optical switch is detailed. Simulation results are presented in Section IV and we conclude the paper in Section V.

II. LOAD-BALANCED ELETRONIC SWITCH

In a basic load-balanced two-stage switch [3] (Fig. 1), the first switch fabric converts the incoming non-uniform traffic into uniform, and the second switch fabric delivers packets to their correct outputs. Each of the two switch fabrics is configured by a pre-determined and periodic sequence of N switch configurations. The only requirement is that each input visits each output (of the same switch fabric) exactly once in the sequence. Accordingly, the switch fabric only needs to realize N switch configurations instead of $N!$ for an input-queued switch. Besides, the two switch fabrics are free to choose their own sequences of configurations. Since middle-stage ports introduce different delays to different packets, it is difficult to ensure that packets of the same flow arrive at an output in order.

A recently proposed load-balanced switch, known as feedback-based two-stage switch [4], elegantly solves the packet mis-sequencing problem. Fig. 1 depicts an $N \times N$ feedback-based two-stage switch, where $VOQ_1(i,k)$ represents the VOQ (Virtual Output Queue) at input port i with packets destined for output k , and $VOQ_2(j,k)$ denotes the VOQ at middle-stage port j with packets destined for output k . Note that each $VOQ_2(j,k)$ is equipped with only a single packet buffer¹.

In the feedback-based two-stage switch, the two sequences of configurations used by the two switch fabrics, i.e. a joint sequence, should have two additional properties, *staggered symmetry* and *in-order packet delivery*. An example is shown in Fig. 2, where the dashed lines show the configurations used by the first switch and the solid lines for the second switch. The first sequence is constructed such that input port i is connected to middle-stage port j at time slot t , where $j = (i + t) \bmod N$. The second sequence is constructed directly from the staggered symmetry, which says that for any middle-stage port j , if it is connected to output k at time slot t , then at next slot $(t+1)$ input k is connected to the same middle-stage port j .

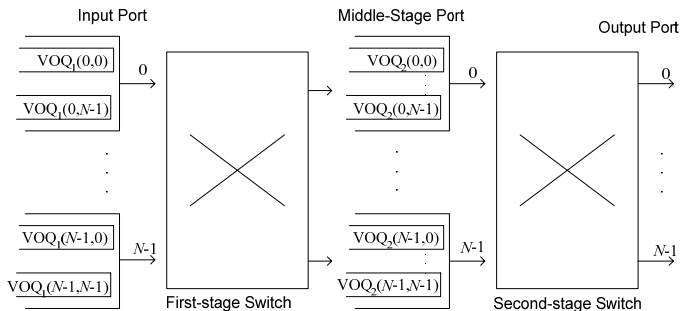


Fig. 1: Feedback-based two-stage switch.

This staggered symmetry property provides an efficient feedback path for reporting middle-stage port occupancy to input ports. Since each $VOQ_2(j,k)$ only has a single packet buffer, an N -bit vector is sufficient to denote the occupancy of the N $VOQ_2(j,k)$'s at each middle-stage port. This vector is piggybacked onto the data packet sent to output k , and is then made available to input k at negligible cost, because both input k and output k reside on the same switch linecard. Due to the staggered symmetry property, input k will connect to middle port j in the next time slot. Therefore, the received occupancy vector provides just-in-time feedback to the local packet scheduler at input k .

On the other hand, the in-order packet delivery property [10] ensures that packets of the same flow always experience the same amount of middle-stage port delay (no matter via which middle-stage port), and is bounded between $[1,N]$ time slots. Refer to Fig. 2. Take packets from input 0 to output 1, i.e. flow(0,1), as an example. At $t=0$, a packet is sent to join middle-port 0 and it will remain in its buffer $VOQ_2(0,1)$ until it is delivered to output 1 at $t=2$. (Again notice that each $VOQ_2(j,k)$ is of size one.) The amount of delay it experienced is 2 slots. In fact, this fixed delay of 2 slots applies to every packet of flow(0,1) no matter which middle-stage port the packet traverses.

Fig. 3 shows another joint sequence of configurations, where input i is always connected to output i in every time slot. This joint sequence does not have the staggered symmetry property, but it can still ensure in-order packet delivery. For example, for packets belonging to flow(0,1), they always experience 1-slot

¹ This feature greatly facilitates the optical implementation of the

feedback-based two-stage switch.

middle-stage port delay. (That is, if a packet from input 0 joins a middle-stage port at slot t . It is guaranteed to arrive at output 1 in $t+1$.) Due to the lack of staggered symmetry property, the middle-stage port occupancy cannot be reported back to input ports by piggybacking onto data packets (sent in the second switch). Accordingly, dedicated feedback packets (incurring additional cost) are required [12]. Therefore, a joint sequence with in-order packet delivery property only is inferior to a joint sequence with both staggered symmetry and in-order packet delivery properties [4]. But in our load-balanced optical switch (detailed in the next section), the joint sequence in Fig. 3 is a better choice because the need for explicit feedback packets/mechanism is removed. (For more details on joint sequence design, please refer to [10].)




Fig. 2: A joint sequence with both staggered symmetry and in-order packet delivery properties.




Fig. 3: A joint sequence with in-order packet delivery property only.

When the occupancy vector arrives at an input port (either by piggybacking using the joint sequence in Fig. 2, or dedicated feedback packets using Fig. 3), the local packet scheduler at each input port can select the best packet to send. It was shown that [4] a longest queue first (LQF) local scheduler is efficient, where a packet from the longest $VOQ_1(i,k)$ is selected for sending if the corresponding middle-stage $VOQ_2(j,k)$ is empty.

In each time slot, depending on the pre-determined configuration in the second stage switch fabric, up to N packets can be delivered to their respective outputs. Since in-order packet arrival is ensured and packets arrived at the line rate of the output port, no output buffer is required (in Fig. 1).

III. LOAD BALANCED OPTICAL SWITCH

In this section, we present our load-balanced optical switch (LBOS) design, which allows a packet to be sent all-optically from an input port to its destination output port, without being electrically buffered/processed at a middle-stage port. It is important to know that switch linecards (where input/output ports locate) on a router must operate in electrical domain, because optical processing, e.g. for table lookup, is not possible. So, our LOBS only targets at the optical implementation of the highlighted parts in Fig. 4.

A. Switch architecture

Our proposed load-balanced optical switch (LBOS) is

depicted in Fig. 5, where N linecards are connected by an N -wavelength WDM fiber ring. Each linecard i is configured to receive (only) on its dedicated wavelength channel λ_i . This creates a strict one-to-one correspondence between N linecards and N receiving channels $\lambda_1, \lambda_2 \dots \lambda_N$. To send a packet to output linecard j , an input linecard i needs to transmit the packet onto channel λ_j when λ_j is idle. Due to the ring topology, the transmitted packet is guaranteed to arrive at linecard j on channel λ_j some time later (depending on how far away the two linecards are).




Fig. 4: Feedback-based two-stage switch with decoupled electrical/optical view.




Fig. 5: The implementation of a 4x4 load balanced optical switch.

Fig. 6 shows the internal structure of a linecard i , where the electrical buffers for implementing $VOQ_1(i,k)$'s (in Fig. 4) are not shown. A linecard consists of three major modules: a fixed receiver (on channel λ_i), a tunable transmitter (implemented using a laser array) and a wavelength monitor. From Fig. 6, incoming optical signals first pass through the EDFA (Erbium Doped optical Fiber Amplifier) for compensating the loss in propagation. Then the optical fiber/energy will be split to enter a fixed optical filter and a demultiplexer, respectively. The filter will remove/block channel λ_i from the fiber and pass all other channels to the fiber delay line (FDL). FDL introduces a fixed delay of t_d , which gives the linecard sufficient time to identify an idle channel for carrying its packets to another linecard (as described below).

The optical signals entered the demultiplexer will be demultiplexed into N individual channels. Wavelength λ_i will be directed to the high bit-rate burst mode receiver², which

² If no further data processing is required at the output port/linecard, we can directly forward the optical signals carried in λ_i (i.e. without passing through the burst mode receiver) to the next-hop router.

demodulates the data from λ_i . The remaining $N-1$ wavelengths are directed to the dc-coupled photodiode array. A threshold comparator is used to detect if a wavelength channel is busy or idle³. Based on the monitoring results and the backlogged traffic at $VOQ_1(i,k)$'s, the LQF (longest queue first) scheduler selects a packet to send using the tunable transmitter. The tunable transmitter is implemented by an array of fixed lasers [11], which is generally more cost-effective than a single fast tunable laser.

Finally the optical signals from the tunable transmitter merge back with the optical signals coming out from the fiber delay line, and the combined signals are forwarded to the next hop/linecard along the fiber ring. The merging process is contention-free because the linecard only sends the data in an idle channel as detected by the monitor.

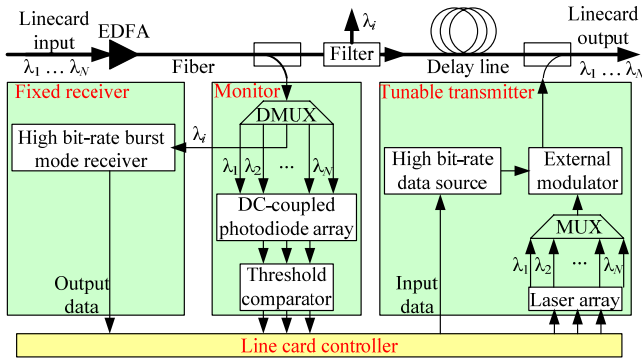


Fig. 6: The linecard i structure

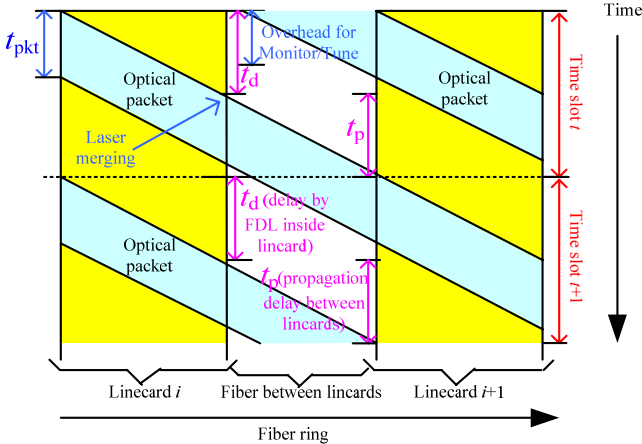


Fig. 7: Time diagram for load balanced optical switch

B. Switch operations

In Fig. 5, let the length of the fiber connecting every two adjacent linecards be equal, and have a propagation delay of t_p . By also taking the FDL inside a linecard into account, we define the duration of a time slot as t_d+t_p , as shown in Fig. 7. Then, one round trip time along the ring is N time slots. We assume the whole system is synchronized, and in each time slot, at most one packet can be transmitted and/or received by each

linecard. Therefore, a fiber ring on a specific wavelength can store up to N optical packets. With N wavelengths, the fiber ring can store up to N^2 packets – same capacity as that rendered by all the middle-stage ports in Fig. 4.

Let the packet duration, or its transmission time, be t_{pkt} . From Fig. 7, linecard i starts to receive a packet at the beginning of slot t and it takes t_{pkt} seconds to completely receive the packet. Meanwhile, the wavelength monitoring circuitry identifies idle wavelength channels on the incoming fiber, and the local LQF scheduler selects a packet (e.g. destined to linecard j) to send (if λ_j is idle and the corresponding $VOQ_1(i,j)$ is the longest). Laser λ_j in the laser array (in Fig. 6) is then activated to carry the packet. The optical signals from the transmitter merge back with the signals coming out from the FDL at t_d seconds after the beginning of the slot. It takes another t_p seconds for the first bit of the packet to arrive at linecard $i+1$. This marks the end of time slot t and also the beginning of slot $t+1$. Note that $t_p \geq t_{pkt}$ must be true, and the optical packet sent by linecard i will arrive at linecard j in optical domain after $(j-i) \bmod N$ time slots.

C. Equivalence with electronic switch

In the following we show that the proposed LBOS (in Fig. 5) is equivalent to the load-balanced electronic switch (in Fig. 1) configured using the joint configuration sequence in Fig. 3. In particular, we focus on how optical buffers (i.e. by storing packets along the ring on different channels as they are propagating) mimic the services provided by the electrical buffers $VOQ_2(j,k)$'s at middle-stage ports in Fig. 1.

In a load-balanced electronic switch, there are altogether N^2 packet buffers at middle-stage ports, whereas the same amount of optical buffer space is provided by the fiber ring, with a round trip time of N slots and carrying N parallel wavelengths. In an electronic switch, the N electrical $VOQ_2(j,k)$'s ($j=0, 1, \dots, N-1$) destined to the same output port k , are equivalent to the N optical buffers available on λ_k (in a round trip time) for each input port and at the beginning of a time slot, the N electrical $VOQ_2(j,k)$'s ($k=0, 1, \dots, N-1$) from a connected middle-stage port j , are equivalent to the N optical buffers simultaneously presented to a linecard on the N wavelength channels ($\lambda_1, \lambda_2 \dots \lambda_N$).

Recall that in a load-balanced electronic switch, all middle-stage $VOQ_2(j,k)$'s are resided away from the linecards. So dedicated occupancy packets [12] or piggybacking vectors [4] must be used to report the queue status. In our LBOS, the “optical” $VOQ_2(j,k)$'s circulate around the ring to visit each linecard in turn. Therefore, their occupancy can be examined locally by the wavelength monitor on each linecard as “optical” $VOQ_2(j,k)$'s pass by – so no explicit feedback mechanism is required. This also helps to explain why the staggered symmetry property of a joint sequence (as that in Fig. 2) is not needed in our load-balanced optical switch.

Next we examine if the ring network in Fig. 5 can exactly mimic the configuration patterns rendered by the joint sequence in Fig. 3 for an electronic switch. Note that input i and output i are collocated on the same linecard i where the monitoring

³ Though λ_i is not received by monitor, it is always considered idle for linecard i .

process takes place. This is equivalent to the examination of the electrical $VOQ_2(j,k)$'s at each middle-stage port in Fig. 3, where a middle-stage port is always connected to the same linecard (i.e. input i and output i).

Consider a 4x4 load-balanced electronic switch using the joint configuration sequence in Fig. 3. Assume a packet, destined for output port 2, is sent from input 0 at time slot $t=0$. As input 0 is connected to middle-stage port 0 at time slot 0 in Fig. 3, this packet will arrive at middle-stage port 0 at the end of the slot and joins $VOQ_2(0,2)$. At $t=1$, middle port 0 is connected to output 1 (not the desired output 2). The packet in $VOQ_2(0,2)$ remains in the buffer. At $t=2$, middle port 0 is connected to output 2. The packet in $VOQ_2(0,2)$ is sent to arrive at output 2. Ignoring the propagation delay, this packet has waited for two slots in $VOQ_2(0,2)$.

We consider the same example using the optical implementation in Fig. 5. That is at $t=0$, a packet destined for output 2 is sent from linecard/input 0. Then at the end of time slot 0, this packet is put on channel λ_2 by linecard/input 0. From Fig. 5, we can see that the total propagation delay (including FDL inside linecards) from linecard 0 to 2 is two time slots. That means the packet will arrive at linecard 2 two slots later.

We can easily generalize the above argument to find a one-to-one mapping between every instance of joint sequence in Fig. 3 and the corresponding operation on the optical ring network in Fig. 5. This shows that our LBOS can exactly mimic the configuration patterns for a load-balanced electronic switch.

IV. PERFORMANCE EVALUATIONS

In this section, we study the performance of our proposed Load-balanced optical switch (LBOS), using simulations. We implement iSLIP algorithm [2] (with a single iteration), which serves as a benchmark for single-stage input-queued switches and output-queued switch, which serves as a lower bound. For more comparison, Fasnet [6] is also implemented because its scheduling algorithm is also distributed and it has similar hardware complexity as ours. In simulating Fasnet, we use the best parameter settings as reported in [6], i.e. fairness quota $Q=100$ and maximum accumulated quota $M \times Q=500$. To be fair, the same simulating environments as Fasnet [6] are adopted, so the propagation delay between the adjacent linecards is 100 ns ($t_p=100$ ns) and each linecard introduces a delay (by FDL) 100 ns ($t_p=100$ ns). Consequently in our simulation the duration for one time slot is 200 ns. Although we only present simulation results for switch with size $N=32$ linecards below, the same conclusions and observations apply for other switch sizes.

Under the above simulation environment, due to the propagation delay between linecards, both LBOS and Fasnet can only utilize 50% of one time slot (see Fig. 7). On the other hand, for (single-rack) switches such as the iSLIP and output-queued switches, they can fully utilize the whole time slot because their propagation delays can be ignored. To have a fair comparison, we allow iSLIP and output-queued switches to send 2 packets per time slot.

A. Uniform traffic




Fig. 8: Delay vs input load, under uniform traffic

Uniform traffic is generated as follows. At every time slot for each input, a packet arrives with probability p (input load p) and destined to each output with equal probability. From Fig. 8, we can see that our load-balanced optical switch (LBOS) gives less-than-20-slot delay performance for input load up to $p=0.8$. When p is reasonably large (>0.7), our LBOS also beats iSLIP. Compared with Fasnet, our LBOS gives significantly smaller delay. When $p=0.6$, Fasnet requires 149.5 time slots, and ours only 17.6 slots, improving the delay performance 88.2%.

B. Uniform bursty traffic




Fig. 9: Delay vs input load, under uniform bursty traffic

Bursty arrivals are modeled by the ON/OFF traffic model. In the ON state, a packet arrival is generated in every time slot. In

the OFF state, no packet arrivals are generated. Packets of the same burst have the same output and the output for each burst is uniformly distributed. Given the average input load of p and average burst size s , the state transition probabilities from OFF to ON is $p/[s(1-p)]$ and from ON to OFF is $1/s$. Without loss of generality, we set burst size $s = 30$ packets.

From Fig. 9, we can see delay builds up quickly with input load, which is due to the bursty traffic nature. Nevertheless, our LBOS performs better than Fasnet. At $p=0.8$, Fasnet requires 546.9 units, whereas 157.4 for LBOS only and delay is cut down 71.2%. When $p>0.6$, our LBOS outperforms iSLIP again.

C. Hot-spot traffic




Fig. 10: Delay vs input load, under hot-spot traffic

Packets arriving at each input port in each time slot follow the same independent Bernoulli process with probability p . Packet destinations are generated as follows. For input port i , packet goes to output i with probability $1/2$, and goes to other outputs with same probability $1/[2(N-1)]$.

From Fig. 10, again we can see LBOS gives comparable and consistently good performance. It delivers close-to-100% throughput, and significantly outperforms Fasnet, which is also distributed and has similar hardware complexity as ours.

V. CONCLUSIONS

A load-balanced (electronic) switch architecture is attractive because it does not require a central scheduler, and provides close to 100% throughput. Although the switch fabrics in a load-balanced switch are amenable to optics (e.g. using AWGR), packets must be electrically handled/buffered between the two switch fabrics (i.e. at middle-stage ports). This makes the all-optical packet transmission from an input linecard to an output linecard impossible. In this paper, a very simple-yet-effective all-optical implementation of a load-balanced switch is proposed. It comprises N linecards connected by an N -wavelength WDM fiber ring. Each linecard i is configured to

receive on channel λ_i . To send a packet, it can select and transmit on an idle channel based on where the packet goes (e.g. using λ_j if the packet goes for linecard j). The ring network is engineered such that the amount of time a packet should be buffered at a middle-stage port exactly matches the propagation delay that this packet experiences en route. This intelligently solves the problem of lacking optical buffers. Following a pre-determined sending schedule, we show that our proposed load-balanced optical switch can give an unbeatable delay and throughput performance under different traffic patterns.

REFERENCES

- [1] N. McKeown, V. Anantharam and J. Walrand, "Achieving 100% throughput in an input queued switch," *INFOCOM 1996*, April 1996, San Francisco, USA
- [2] N. McKeown, "Scheduling algorithms for input-queued cell switches," *PhD. Thesis*, University of California at Berkeley, 1995.
- [3] C. S. Chang, D. S. Lee and Y. S. Jou, "Load balanced Birkhoff-von Neumann switches, part I: one-stage buffering," *Computer Communications*, Vol. 25, pp. 611 – 622, 2002.
- [4] K. L. Yeung, B. Hu and N.H. Liu, "A novel feedback mechanism for load balanced two-stage switches," *ICC 2007*, June 2007, Glasgow, Scotland.
- [5] I. Keslassy, S.T. Chuang, K. Yu, D. Miller, M. Horowitz, O. Solgaard and N. McKeown, "Scaling the Internet Routers using Optics," *ACM SIGCOMM'03*, Aug. 2003, Karlsruhe, Germany.
- [6] A. Bianco, E. Carta, D. Cuda, J. M. Finochietto and F. Neri, "A distributed scheduling algorithm for an optical switching fabric," *ICC 2008*, May 2008, Beijing, China.
- [7] P. Bernasconi, C. Doerr, C. Dragone, M. Capuzzo, E. Laskowski and A. Paunescu, "Large $N \times N$ waveguide grating routers," *Journal of Lightwave Technology*, Vol. 18, No. 7, pp. 985-991, July 2000.
- [8] R. Ryf et al., "1296-port MEMS transparent optical crossconnect with 2.07 petabit/s switch capacity," *Optical Fiber Comm. Conf. and Exhibit (OFC) '01*, Vol. 4, pp. PD28 -P1-3, 2001.
- [9] X. Wang and K. L. Yeung, "Load balanced two-stage switches using arrayed waveguide grating routers," *IEEE Workshop on High Performance Switching and Routing*, May 2007, New York, USA.
- [10] B. Hu and K. L. Yeung, "On joint sequence design for load-balanced two-stage switch architecture," *IEEE Workshop on High Performance Switching and Routing*, May 2008, Shanghai, China.
- [11] A. Carena, V. D. Feo, J. Finochietto, R. Gaudino, F. Neri, C. Pigioli and P. Poggiolini, "RINGO: an experimental WDM optical packet network for metro applications," *IEEE Journal on Selected Areas in Communications (JSAC)*, Vol. 22, No. 8, pp. 1561-1571, Oct. 2004
- [12] H. I. Lee, B. C. Lee and S. W. Seo, "A load balancing scheme for two-stage switches maintaining packet sequence," *ICC 2006*, June 2006, Istanbul, Turkey.