A Novel Optical Buffer Architecture for Optical Packet Switching Routers

Assaf Shacham, Benjamin A. Small, and Keren Bergman
Department of Electrical Engineering, Columbia University, New York, NY 10027; assaf@ee.columbia.edu

Abstract *We present an optical packet buffer architecture designed to implement queues in optical packet switching routers. Correct functionality and error-free routing are experimentally demonstrated on a prototype buffer.*

Introduction

One of the key challenges in the implementation of all-optical routers lies in packet buffering. Contemporary internet routers use very large packet buffers, which store millions of packets, to efficiently utilize expensive long haul links. These large buffers are clearly impractical for implementation using photonic technology. Recent studies have indicated that by sacrificing some of the link utilization, buffer sizes can be reduced dramatically to the capacity of approximately 20 packets [1]. Photonic packet buffers of this capacity may have a realistic potential and several modules have been suggested [2]-[5].

The aforementioned modules can either store a single packet [2],[3] or require a complex and non-scalable control scheme [4],[5]. An architecture that will facilitate the construction of first-in-first-out (FIFO) buffers of different capacities while maintaining simplicity and scalability is therefore sought.

In this paper we introduce a novel modular optical packet buffer architecture comprised of primitive building-block modules, each of which is capable of storing a single packet. The modular structure facilitates simplicity and allows for the simultaneous storage of multiple packets, while maintaining FIFO ordering. We experimentally demonstrate a prototype building-block module with correct functionality and error-free storage of optical packets.

Buffer architecture

The presented optical packet buffer is comprised of identical buffer building-block modules organised as a chain (fig. 1). Each building-block module has two input ports and two output ports and is capable of storing a single packet on a fibre delay line (FDL). A pair of ports (*Up-in* and *Up-out*) connects the module to the next module in the chain, and the *Down-in* and *Down-out* ports are connected to the previous module. In the root module the *Down-in* and *Down-out* ports are used as the system input and output ports, respectively. Each module is also connected to the next module in the chain by an electronic cable, for the transmission of *Read* signals.

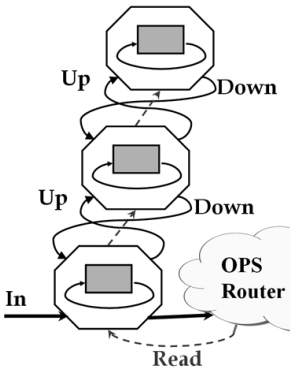

Fig. 1 - system architecture

Writing packets to the buffer is performed implicitly: optical packets arrive into the *Down-in* port of the root node aligned with system timeslots. The packets are either stored locally if the internal FDL buffer is empty, or routed to the *Up-out* port to be stored in the next module in the chain. Each packet is forwarded in this manner up the chain, and is stored in the first empty module it encounters, which is necessarily the last position in the queue.

The read process is completely unrelated to the write process: when a *Read* signal is received, the locally stored packet is transmitted from the *Down-out* port and a *Read* signal is sent to the next module in the chain to retrieve the next packet. With each *Read* signal all the packets in the chain move a step closer to the output port, while maintaining the packet sequence.

This distributed modular structure has several advantages: (1) no central management is required – all modules follow an identical simple set of rules; (2) the buffer capacity can be increased simply by connecting additional modules at the end of the chain; (3) packet dropping, in the case of overflow, is cleanly executed by routing packets to the *Up-out* port of the last module, which is not connected.

Building-block module structure

Each building-block module in the buffer is implemented as an SOA-based 3×3 non-blocking optical switch (fig. 2). At each of the three inputs (*Up-in*, *Down-in*, and *Buffer-in*), the packet's power is split by a coupler. A portion of the packet's power is directed to a low-speed *p-i-n* photodetector acting as an envelope detector. The envelopes of the packets

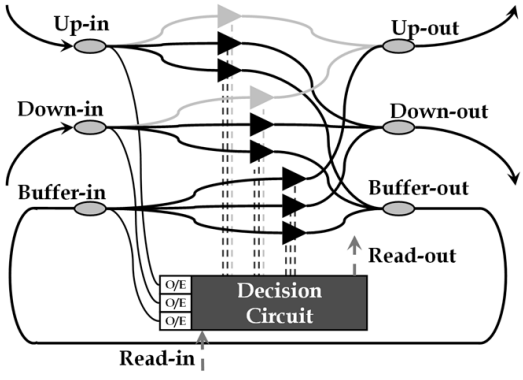

Fig. 2 - module block diagram

from all input ports along with the input electronic *Read* signal are used in an electronic decision circuit to set the state of the SOA-based switch. The decision rule used by the electronic circuitry is represented by the truth table in Table 1. The table should be read as follows: the four left columns {U,B,D,R} are inputs representing the existence of packet on the input ports (*Up-in*, *Buffer-in, Down-in*) and the *Read* signal, respectively. The next nine columns represent the switching states of the SOAs (e.g. U2B means "*Up-in* to *Buffer-out*", etc.) Finally, the last column (ER) represents when a *Read* signal is transmitted to the next module in the chain.

Table 1

U	B	D	R	U2U	U2B	U2D	B2U	B2B	B2D	D2U	D2B	D2D	ER
0	0	0	0										
0	0	1	0								1		
0	1	0	0					1					
0	1	1	0					1		1			
1	0	0	0		1								
1	0	1	0		1					1			
1	1	0	0										
1	1	1	0										
0	0	0	1										
0	0	1	1									1	
0	1	0	1						1				1
0	1	1	1						1	1			1
1	0	0	1			1							1
1	0	1	1			1				1			1
1	1	0	1										
1	1	1	1										

To assist in understanding the decision rule, the following examples are given: line {0110} represents a case where a new packet received from the *Down-in* port (D=1), while another packet is locally stored in the buffer (B=1). In that case the locally stored packet is routed back to *Buffer-out* (B2B=1) while the new packet is routed to the next module through the *Up-out* port (D2U=1). Line {0101} represents the case where a *Read* signal is received (R=1) and a packet is locally stored (B=1). In this case, the local packet is sent to the *Down-out* port (B2D=1) and a *Read* signal is sent to the next module (ER=1) to retrieve the next packet in the queue.

Careful inspection of Table 1 reveals that some states can never occur if all the modules follow the rules. More importantly, two paths (*Up-in* to *Up-out* and *Buffer-in* to *Up-out*) are never used so the module can be implemented using only seven SOAs. The forbidden states and the unused SOAs appear shaded in fig. 2 and in Table 1.

Experimental demonstration

To verify the feasibility of implementing the optical packet buffer, a prototype of the buffer building-block module is constructed using SOAs, passive optical elements, and 155 Mb/s *p-i-n* photodetectors. The decision circuit is implemented using a Xilinx complex programmable logic device (CPLD).

The SOA switching elements have several advantages. They compensate for coupling and other losses, thus maintaining the packet power, which is crucial for ensuring that packets have a long lifetime within the buffer. Because SOAs amplify over a wide wavelength range, this buffer implementation is transparent bit-rate and modulation format and can also store multiple-wavelength packets.

A stream of 45.0-ns long optical packets, spaced by 6.4 ns gaps are generated, and modulated at 10 Gb/s. Each packet is also modulated with a 7-bit label on a separate wavelength to allow for its experimental identification at the input and output ports. *Read* signals are also generated in order to read packets from the module. In this experiment, since only a single module is implemented, the *Up-out* port is connected to the *Up-in* port through an optical fibre, emulating the presence of an additional building-block module.

The waveforms in fig. 3 confirm the correct functionality of the module: packet *A* (label 1101111) is written and immediately read, with a 40-ns latency; packet *B* (1101011) is buffered for one timeslot; packet *C* (1100111) is read after 3 slots, or 192 ns; packet *D* (label 1011011) is buffered for 2 slots.

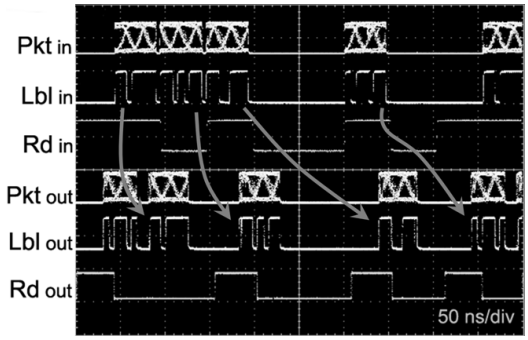

Fig. 3 - waveforms of the input and output signals for functional demonstration

All packets were measured to be error-free (bit error rate < 10^{-12}) at the buffer output.

Conclusions

A novel architecture for optical packet buffers is presented. The architecture is modular, and the simple construction of the building-block modules can straightforwardly scale to the implementation of buffers of different capacities simply by connecting additional modules. A prototype building-block module is constructed, its functionality is experimentally verified, and error-free storage of 10 Gb/s packets is demonstrated for buffering times between 40 ns and 192 ns.

References

1 Enachescu *et al.* "Routers with Very Small Buffers" IEEE INFOCOM'06,

2 Shacham *et al.*, PTL, 17 (2005), pp. 2778-2780.

3 Yeo *et al.*, PTL, 16 (2004), pp 2559-2561.

4 Langenhorst *et al.*, JLT, 14 (1996), pp. 324-335.

5 Chlamtac *et al.*, JSAC, 14 (1996), pp. 1014-1029.