

Dynamic Injection Control Module for Optical Packet Switching Fabrics

Assaf Shacham, Benjamin G. Lee, and Keren Bergman

Department of Electrical Engineering, Columbia University, New York, NY 10027

assaf@ee.columbia.edu; (212) 854-2768

Abstract A novel module that mediates packet injection into optical packet switching fabrics in the presence of contentions is presented. Functionality and design are discussed, and correct operation is verified for WDM packets with 12x10Gb/s data.

Introduction

Optical packet switched (OPS) networks offer the unique combination of encompassing the enormous capacity of the optical domain while providing versatile connectivity afforded by individual packet routing. The insertion of OPS networks to alleviate the electronic bottleneck has been proposed for a range of systems including core routers, local- and storage-area networks as well as high-performance computing systems [1]-[4]. Whereas electronic interconnection networks do not scale well when systems move toward high port counts, high data rates and large spatial volumes [5], fiberoptic technology can offer immense bandwidth utilizing wavelength division multiplexing (WDM), low latency on the order of the packet's time of flight and better scalability [6].

Systems based on optical packet switched interconnection networks will have to address new challenges presented specifically by photonic technology. The shortage of efficient optical memory requires new network elements to be designed that utilize other means to resolve traffic contention such as deflection routing and fibre delay lines. Retiming modules, such as packet synchronizers that align packets with system slots, and interface modules that control the packet injection into the system are critical for time-slotted OPS switching fabrics.

In this paper, we present a novel optical injection control module that receives an arbitrary stream of optical packets from a network terminal and retimes packets until an input port is available. The module, designed for compatibility with Data Vortex switching fabrics [1] maintains the wideband wavelength-parallel packets in the optical domain and effectively serves as a one-packet input buffer. Experiments are run to verify the functionality of the injection module and to measure the optical signal integrity of the routed packets.

Data Vortex Input Interface

The Data Vortex architecture is a synchronous distributed OPS architecture that employs deflection routing and a high degree of path diversity to virtually buffer packets within the network [1],[7]. A Data Vortex switching fabric is comprised of simple 2x2 switching nodes, interconnected by optical fibres, with each node processing a single packet per system time-slot. Some of the switching nodes also serve as

input ports to the system, while being connected to internal fibres as well (fig. 1). In the architecture definition, internal traffic has priority over incoming traffic, so when a switching node receives an internally routed packet, it is not available to receive an incoming packet, and the injection of the incoming packet must be avoided. The switching node that transmits the internal packet sends out an electronic signal that notifies the injecting element that injection is not allowed during that time slot.




Fig. 1 – The Data Vortex input nodes (DV) are interconnected by optical fibres (dashed). The injection control module (IC) mediates between the transmitters (TX) and receives a control signal from the Data Vortex nodes (dotted).

The injection control module mediates between the Data Vortex switching fabric and the packet-generating elements (e.g. a terminal in a LAN). It receives a stream of optical packets and an electronic control signal and decides whether to inject the packet right away or delay until it can be successfully injected. In cases of contention, where a packet is delayed and a new packet is received while injection is forbidden, the injection control gracefully discards one of the packets. The module should work under stringent timing constraints, where the availability of the input port is not known until the control signal is received, just a few nanoseconds before the beginning of the slot.

Module Description

The proposed module (fig. 2) is based on a 2x2 wideband photonic switching node implemented using semiconductor optical amplifiers (SOA) organized in a tree structure. The switching node controls a feedback fibre delay line (FDL) whose latency matches exactly to one system's slot time.

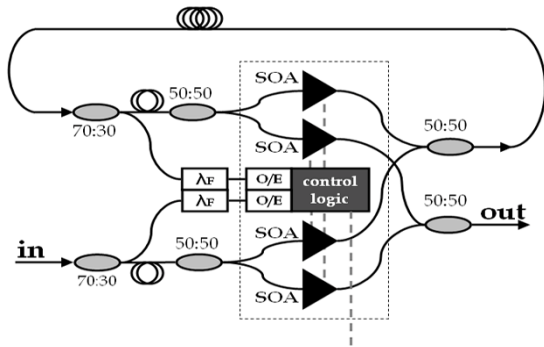


Fig. 2 – The injection module is an SOA based switching node with a feedback FDL. It is comprised of couplers (ellipses), filters (λ), p-i-n photodetectors (O/E), control electronics, and SOAs.

The module routes the wavelength-parallel Data Vortex packets described in [7] where the header and the payload are encoded on separate dedicated wavelengths and are transmitted as one unit in a single time-slot. When a packet is received, the *frame* signal, denoting the existence of a packet, is extracted using a wavelength filter, detected, and driven into the control electronics, implemented as a Xilinx complex programmable logic device (CPLD). The CPLD uses the frame signal from the received packet and from the delayed packet, along with an externally received control signal to dynamically decide, on a packet-by-packet basis, where to route each packet. Four SOA gates, driven by the CPLD, control the four possible input-output paths and allow the module to handle two packets (one from the input and one from the FDL) simultaneously.

Experimental results

Two experiments are performed to verify the correct functionality of the injection module. In the first experiment, a stream of optical packets, each 38 ns long, is driven into the input port of the module while the control signal is intermittently switched. The waveforms in fig. 3 indicate that packets are delayed and are only injected when the control signal indicates that injection is allowed (i.e. when it is low).

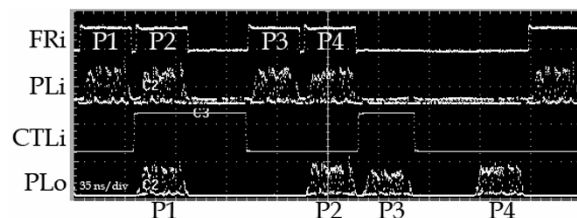


Fig. 3 – The waveforms shown are (top down): input frame optical signal, input payload, input control signal, and output payload (after a latency of 38 ns). Four packets, annotated P1 through P4, are injected into the module and only proceed to the output when the control signal is low.

In the second experiment, the signal integrity of the packets routed through the module is verified. Bit error rate (BER) measurements are taken on one wavelength (1552.3nm) to measure the power penalty caused by a hop through the module. The measured power penalty is smaller than 0.2 dB (fig. 4).

Then, a 120 Gb/s wavelength-parallel payload (12×10 Gb/s) is routed through the module. The payload wavelengths span over 17nm (1543.1nm to 1560.4nm) with a minimum spacing of 0.8nm. After 3 loop passes a BER better than 10^{-9} is achieved on all wavelengths. Under typical Data Vortex operating conditions, the probability that a packet will need to traverse the loop more than 3 times before entering the OPS system is less than 10^{-6} [1],[8]. Eye diagrams after 3 loop passes are shown in fig. 4.

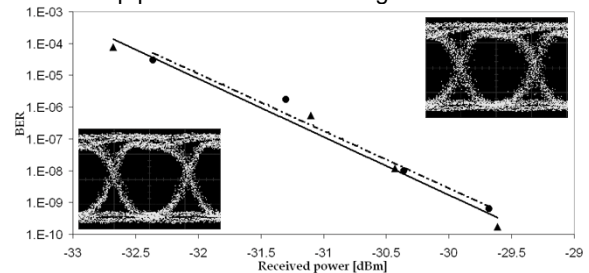


Fig. 4 – BER curves for the module after 1 hop (●,dashed), and back-to-back (▲,solid) indicate a power penalty < 0.2dB. Eye diagrams at 10 Gb/s ($\lambda=1554.4$ nm) for back-to-back (bottom left) and after 3 loop passes (top right).

Conclusions

We have demonstrated a dynamic injection control module which is essential to the operation of optical packet switching fabrics. The module design is detailed and experiments verifying the functionality and the signal integrity are reported. Future work will include inter-operability testing with a Data Vortex switching fabric and priority mechanisms that control packet dropping.

References

- 1 Q. Yang et al, *J. Lightwave Technol.*, **19** (2001), 1420-26.
- 2 Hemenway et al., *J. Opt. Netw.*, **3** (2004), 900-913.
- 3 J. Gripp et al., *J. Lightwave Technol.*, **21** (2003), 2839-2850.
- 4 M. Y. Jeon et al., *J. Lightwave Technol.*, **21** (2003), 2723-2733.
- 5 W. J. Dally, *Principles and Practices of Interconnection Networks*, Morgan Kaufmann, 2004.
- 6 G. I. Papadimitriou et al., *J. Lightwave Technol.*, **21** (2003), 384-405.
- 7 B.A. Small et al., OFC 2005, OWK1.
- 8 A. Shacham, "Data Vortex Single Angle Injection Study", internal manuscript, April 2005.