Towards All-Optical Packet Networks

Raluca-Maria Indre
Orange Labs, France
ralucamaria.indre@orange.com

Thomas Bonald, Davide Cuda
Telecom ParisTech, France
{thomas.bonald, davide.cuda}@telecom-paristech.fr

Ludovic Noirie
Alcatel-Lucent Bell Labs, France
ludovic.noirie@alcatel-lucent.com

***Abstract*—We propose a technique for building optical packet networks that does not require any buffering, any signalling or header processing. Contentions are solved by means of an optical device that allows the first packet to go through while blocking others. Blocked packets are redirected back to their source nodes, thus notifying the latter about the packet status. We describe the design principles of the corresponding all-optical networks, assess their performance and power consumption and give examples of application in the context of access networks and data centers.**

I. Introduction

Today's optical technologies such as wavelength division multiplexing (WDM) are mainly used to provide high-capacity point-to-point links between electronic network nodes. While optical transmission is able to cope with increasing traffic demands, electronic switching is currently reaching fundamental limits in terms of processing speed, energy requirements and port count [1]. The gap between high-speed optical transmission and limited electronic processing can be bridged by moving some switching functionalities to the optical domain.

Ideally, to match IP traffic and achieve high utilization, optical switching should be performed at packet level. Proposed Optical Packet Switching (OPS) solutions mimic electronic IP networks. At each node, the information contained in the packet header is extracted and processed electronically to make the forwarding decision; meanwhile, the payload is stored optically using fiber delay lines (FDL). Unfortunately, contentions cannot be solved in the time domain, as in electronic networks, due to the lack of optical random access memory (RAM). Moreover, the OPS technology is limited by the power consumption required to perform O/E/O conversion of packet headers. Recent studies are questioning the ability of OPS to reduce power consumption when compared to present-day electronic routers [1].

In this paper, we propose an approach for building optical packet networks without resorting to any optical buffering, signalling or header processing. The proposed solution is particularly suitable for corporate networks, access networks and data centers. Buffering and routing operations are performed electronically at the network edge, while packets are transmitted end-to-end in the optical domain. Contentions are resolved through a simple first-come first-served policy: the first optical packet to arrive goes through while the others are blocked and sent back to their source node, as explained in the following.

The authors are members of LINCS (Laboratory of Information, Network and Communication Sciences), Paris, France. See `www.lincs.fr`.

II. Optical Devices and Networks

We propose to resolve contention by means of a new optical device, that we call a Dynamic Optical Combiner (DOC). A DOC has N input ports and a single output port. At any given time, at most one of the input optical signals is transmitted to the output. Specifically, the first input to become active is allowed to pass; during the transmission, any other input signal is blocked and redirected to the source using the same fiber on which the signal arrived. Importantly, the cost, complexity and power consumption of a DOC is practically independent of the data rates.

The feedback mechanism used by the DOC is similar in principle with the concept of protection against collision described in [2] for star-coupler based WDM systems. The feasibility of the feedback mechanisms has recently been demonstrated experimentally in the context of array waveguide grating router (AWGR)-based optical interconnects [3].

DOCs can be used to build a wavelength-routed optical packet switch, that we call *optical switch-combiner*. Specifically, the WDM signal received on each input fiber is first demultiplexed into W different optical signals, one for each wavelength channel. The optical signals are then separated into W wavelength planes, i.e., signals transmitted on the same wavelength but on different input fibers are grouped together and sent to a DOC. The latter will allow only one of the input signals to pass at any given moment. The output signal of each DOC is an aggregation of the data transmitted on all input fibers, on a given wavelength channel. This signal is then transmitted to some predefined output port(s).

Optical switch-combiners can be used to build wavelength-routed optical networks in which every wavelength is associated to a specific set of destinations. The network consists of a set of edge nodes that transmit and receive data in the form of optical packets and a set of optical switch-combiners that multiplex and forward these packets, possibly dropping some, as described previously.

Each edge node is equipped either with a *tunable* transmitter so as to select the appropriate wavelength, depending on the destination of the optical packet, or with a simpler, less expensive, *fixed* transmitter, configured on the appropriate wavelength, if the edge node transmits data to a single set of destinations. Each transmitter is characterized by some bit rate R, typically equal to 1 or 10 Gbit/s.

Edge nodes are also equipped with one or several fixed receivers, depending on the wavelengths on which they receive data. These wavelengths may be shared with other edge nodes in order to maximize the network utilization. In particular, an optical packet sent by some source on wavelength w to some specific destination is usually received by a set of nodes, all receiving data on wavelength w. These edge nodes convert the received signal to the electronic domain and process the packet header in order to identify the packets addressed to them. Note that the packet headers are only processed at the network edge to determine the destination; the switch-combiners multiplex and forward packets based on the activity of the input ports on the different wavelengths, without any header processing.

The edge nodes generate traffic in the form of *flows* (e.g., data transfers, voice calls, video streams) that compete for access to the network. A Medium Access Control (MAC) is thus required to enforce some form of fairness in the sharing of network resources. We consider a simple access scheme in which every flow delays the transmission of each of its packets by some random backoff time, as in IEEE 802.11 networks for instance. Specifically, after every packet transmission a flow goes into an idle state of random duration. When the backoff period ends, the flow becomes eligible for transmission if the transmitter is free; otherwise, the flow waits for another backoff time.

The proposed MAC regulates access to each transmitter and to each wavelength. Note that this fully distributed MAC does not prevent from packet collisions inside the network; it is the role of switch-combiners to resolve the contentions by letting one packet pass and by redirecting the colliding packets back to the corresponding sources. These packets are then retransmitted after some random backoff time, according to the above access scheme.

III. Performance evaluation method

We assess network performance through some variants of the Engset model at the packet level and processor-sharing queues at the flow level. Specifically, n flows sharing a common transmitter according to the above access scheme can be modeled using the Engset model [4] with a single circuit. Indeed, the flows *sense* the transmitter and can immediately determine whether the channel is occupied. The flow only transmits if the transmitter is sensed idle; otherwise the flow waits for another random backoff time before re-attempting a transmission. On the other hand, n flows sharing a common wavelength channel at the output of the DOC *cannot sense* the state of the channel, i.e., they are not aware of the wavelength occupancy. This corresponds to the generalized Engset model [5] which is no longer tractable and some approximations are needed to get explicit expressions. To this end, we propose to use a simple fixed-point approximation [6]. This approximation allows to derive expressions for both homogeneous and heterogeneous sources. We show the accuracy of the proposed approximation by comparing the results obtained using it with results obtained through simulations.

We then apply the previous packet-level contention models to estimate the throughput and delay performance of the proposed solution. The flow-level model corresponds to a processor-sharing queue with state-dependent service rate. We use this flow-level model to assess the throughput achieved by elastic flows and the delay incurred by real-time traffic such as video streaming or voice applications. Our results show that performance is comparable to that of electronic networks of same capacity, both in terms of throughput and packet delay. The key advantages of our solution are the scalability in the input data rates and a huge gain in power consumption. Indeed, we show that the power consumption of the proposed solutions dynamically varies with the network load, bringing reductions of one order of magnitude or more.

IV. Case studies

We illustrate the practical interest of our solution in the case of access networks and data centers. We show that an access network of 4,000 users and an aggregate capacity of 80 Gbit/s can be built using 4 switch-combiners, each equipped with two DOCs for resolving contention on the upstream wavelength channels. We also show that the proposed optical networks can be used to build a data center of 12,000 servers using only 80 wavelength channels, each operating at 1 Gbit/s. By reutilizing wavelengths over non-overlapping segments of the data center, the total traffic which can be sustained by the considered data center is 6 Tbit/s. Using both simulations and the methodology described in the previous section, we show that performance in terms of throughput and delay is excellent for both the data center and the access network.

V. Conclusion

We have propose a novel optical device based only on off-the-shelf components, able to resolve packet contention optically without requiring any electronic signalling or header processing. We show that this simple component can be used as a building block for all-optical packet networks. We analyse the performance of the proposed optical networks using variations of the Engset model at the packet level and processor sharing queues at flow level. The practical interest of these networks is illustrated in the context of access networks and data centers. Future work will be focused on the control plane and on recovery issues in case of failures.

References

[1] R. Tucker, "Scalability and energy consumption of optical and electronic packet switching," *Lightwave Technology, Journal of*, vol. 29, aug., 2011.

[2] B. Glance, "Protection-against-collision optical packet network," *Lightwave Technology, Journal of*, vol. 10, no. 9, pp. 1323 –1328, sep 1992.

[3] R. Proietti *et al.*, "All-optical physical layer nack in awgr-based optical interconnects," *Photonics Technology Letters, IEEE*, vol. 24, no. 5, pp. 410 –412, march1, 2012.

[4] T. Engset, "On the calculation of switches in an automatic telephone system," in *Tore Olaus Engset: The man behind the formula*, A. Myskja and O. Espvik, Eds., 1998.

[5] J. Cohen, "The generalized engset formulae," *Philips Telecommun. Rev.*, vol. 18, 1957.

[6] H. L. Vu, A. Zalesky, E. Wong, Z. Rosberg, S. Bilgrami, M. Zukerman, and R. Tucker, "Scalable performance evaluation of a hybrid optical switch," *Lightwave Technology, Journal of*, vol. 23, no. 10, pp. 2961 – 2973, oct. 2005.