The Agile All-Photonic Network: An Architectural Outline

G. v. Bochmann T. Hall O. Yang
School of Information Technology & Engineering
Ottawa University
Email: {bochmann,thall,oyang}@site.uottawa.ca

M.J. Coates L. Mason R. Vickers
Department of Electrical and Computer Engineering
McGill University
Email: {coates,mason,vickers}@ece.mcgill.ca

I. Introduction

In contrast to current generation core networks, all-photonic networks have the property that both transmission and switching is performed in the optical domain. Data entering the network is converted into an optical signal, and this signal is not converted back to the electronic domain until the signal has reached the exit-point of the network. The absence of the need for any conversion leads to two important advantages: first, optical switches have (at least potentially) far greater switching capacity than electronic domain switches, removing one of the bottlenecks in a high speed network; second, the switches are transparent to data format and bit rate. The primary disadvantage is that the tasks necessary for packet switching (buffering, addressing or labelling) are very difficult to implement in the optical domain.

In this paper, we outline the architecture of an agile all-photonic network (AAPN), focusing primarily on topology aspects, switching mechanisms, and resource allocation. The design assumes that optical switching technologies will mature to the point where rapidly-reconfigurable, medium/high port-count, all-optical space-switches are deployable in the network core. The aim of the AAPN design is to extend the photonic portion of the data path as close to the end-user as possible. We believe that any practical, architectural paradigm that strives to achieve this goal must feature: (i) *agility* - the ability to perform time domain multiplexing to dynamically allocate bandwidth to traffic flows as the demand varies; and (ii) the concentration of control and routing functionality at the electronic edge switches that surround the photonic core. The latter approach obviates the need to resort to technologies that we believe will not be realized in the near future, including optical memory, optical header recognition, and practical, fast wavelength conversion. The agility of the network is critical if the photonic core is to be extended to small-capacity switches near the edge of the network.

II. Topology

The majority of current core networks have a (logical) mesh topology, which is both robust and distributes traffic load over many switches. In the AAPN, the photonic core switches have enormous capacity, so it is appropriate to rethink topology design. One of the main challenges in the AAPN is control of the network; because there is no buffering, we cannot afford to have frequent contention for resources at switches. At a given point in time, wavelength on an outgoing link on any photonic switch must be allocated to a single connection through the network. If the core of the network is a mesh of photonic switches, then the coordination of the switches becomes a very challenging problem. As demonstrated in the PetaWeb project [1], [2], which explored the design of an all-photonic network using slowly reconfigurable photonic switches, the overlaid star topology can compare favourably with mesh architectures. The overlaid star topology, depicted in Figure 1(a), connects edge nodes together via several central photonic core nodes. The overlay of several stars provides robustness in the case of link or core node failure, and also helps mitigate the potential for backhauling.

The overlaid star forms a logical mesh by connecting each edge node to every other edge node. However, data traversing the network only passes through one photonic switch, resulting in a major simplification of the control problem of ensuring that contention is rare. From the control point of view, each star can be managed independently of the others because there is no data interaction. For each star, resource allocation is concentrated at a single point in the network.

The edge nodes are not necessarily of equal capacities in terms of traffic; the number of wavelengths they support for traffic can vary. However, for protection purposes, each edge node is a member of every star in the overlaid structure and the division of its capacity between the stars is fixed on a wavelength basis, with dedicated fibres for each of the stars. The division of edge node traffic between the different star layers of AAPN need not be into equal parts. However, each edge node must be capable of supporting some level of traffic to every other edge node in the network (on all stars). This last condition places a

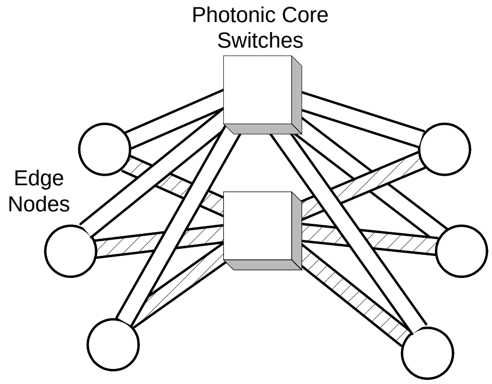

(a) Overlaid Star Architecture

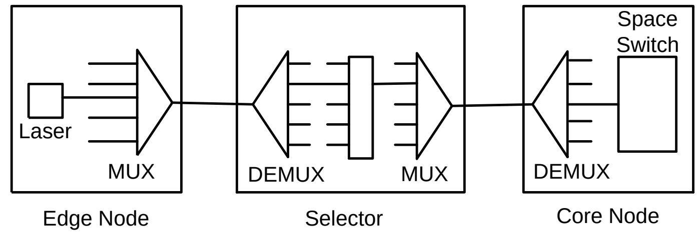

(b) *Half path* from edge node to core node.

Fig. 1. Topological aspects of the AAPN. (a) The overlaid star logical topology. (b) A depiction of the half path from an edge node to the core node in the AAPN. (DE-)MUX indicates wavelength (de-)multiplexer.

lower bound on the photonic port capacity required for an edge node. The actual traffic supported by the edge node can be less than the port capacity, but overall network efficiency then drops.

The data path within the core node is purely photonic. Signals on the incoming ports are demultiplexed, and the extracted components passed to a layered switch, with each layer switching the time slots for a given wavelength. There are therefore as many layers in the photonic switch as there are wavelengths on the network. Current technology trends indicate that a 64 port x 64 port switch is as large as one individual layer can reach in a crossbar structure, if reasonable switching speed is to be achieved. A method of port sharing is thus required to allow a Core Node to support more than 64 edge nodes. We achieve port sharing using a Selector (1 x N switch) that combines the traffic of multiple edge nodes onto a single fibre. Port sharing may be physically collocated with the core switch, or remotely located.

The *half path* of the network, that is the portion from a single edge node to the core node, is depicted in Figure 1(b). Each edge node has multiple lasers, and is thus capable of simultaneously transmitting (and receiving) data on multiple wavelengths. These signals are multiplexed onto a single fibre. The selector demultiplexes the incoming signal into individual wavelengths, and passes each wavelength to an selector. In total, N edge nodes are connected to the selector. On each wavelength, the selector extracts the signal arriving from one edge node, and forwards it to a multiplexer, which recombines all the wavelengths on a single fibre. This signal is passed to one of the ports of the central node.

III. Switching and Bandwidth Allocation

In order to retain flexibility, we design the AAPN to support three different switching modalities. In the first type, an entire wavelength is dedicated to the connection between two edge nodes. This is appropriate for scenarios where there is persistent traffic of substantial volume between relatively large edge nodes. The allocation of wavelengths is reassessed periodically (on the order of hundreds of seconds) to determine if the wavelength is being efficiently used. The second modality involves time-slot reservation; wavelengths are divided into *frames* comprising a fixed number of time slots, and each connection between two edge nodes is allocated a certain number of time slots in the frame. The number of allocated time slots is adjusted as the traffic demand between the two nodes changes. This modality, which enables bandwidth adjustment on the order of seconds, is designed to address connections with slowly-varying demand. Finally, some wavelengths in the network are dedicated for the third switching modality, optical burst switching, which is used to address unanticipated bursts in traffic demand.

References

[1] R. Vickers and M. Beshai, "PetaWeb architecture," in *Networks 2000 Symposium*, Toronto, Canada, 2000.

[2] F.J. Blouin, A.W. Lee, A.J.M. Lee, and M. Beshai, "Comparison of two optical-core networks," *J. Optical Networking*, vol. 1, no. 1, pp. 56–65, Jan. 2002.