Transportation & Switching of Native Ethernet Frames Across MPLS/GMPLS Managed and Controlled Optical Data Networks

M. A. Ali, A. Hadjiantonis, A. Khalil, G. Ellinas
Dept. of Electrical Engineering,
The Graduate Center and the City College of the City University of New York,
{ali, antonis, akhalil, gellinas}@ee.ccny.cuny.edu

Keren Bergman
Dept. of Electrical Engineering,
Columbia University,
bergman@ee.columbia.edu

ABSTRACT- In this talk, we outline and propose simple optical networking solutions that constitute a credible starting foundation en-route to the rapid adoption and deployment of fully automated optical data networking paradigms

The phenomenal growth in Internet traffic along with the abundance of transmission capacity propelled by the explosion of wavelength-division multiplexing (WDM) technology and the rise of optical communication systems has signaled the beginning of a new networking paradigm where the vision of integrating data and optical networking seems to be indispensable. While IP has emerged as the primary network layer technology, recent advances in optical technology are now transforming the optical layer from a static transmission facility into a dynamically reconfigurable network. There is an emerging consensus that a simplified, two-tiered architecture combining the strengths of both IP and optical transport in which IP can be implemented directly over WDM, bypassing all of the intermediate layer technologies, is the key to realizing such a vision [1-4].

Several industrial organizations including the Internet Engineering Task Force (IETF) and the Optical Internetworking Forum (OIF) have already proposed several architectural options on how IP routers must interact with the optical layer to achieve end-to-end connectivity, including overlay, augmented, and peer-to-peer models (interconnection models) [2-4]. The simplest is to treat the optical layer as completely separate from the IP layer. In this "overlay" model, the optical layer provides point-to-point connections (lightpaths) to the IP domain. The client routers request high-bandwidth connections from the optical network, via a well-defined signaling and routing interface, referred to as User to Network Interface (UNI), and are provided with no knowledge of the optical network topology or resources. A more sophisticated model that offers a tighter integration between IP and optical layers (peer model) collapses the two layers into a single integrated layer managed and traffic engineered in a unified manner. IP-based Multi-protocol label switching (MPLS) and its extension Generalized MPLS (GMPLS) have been proposed as the integrating structure between IP and optical layers [2-4].

Despite the phenomenal advances and development in IP-optical networking integration over the last several years, a mixture of both the fiscal and technological hurdles described above have dramatically slowed down the adoption of automated optical data networking paradigms. In fact, the lack of true economic drivers for a network that can only support provisioning of coarse/fixed granularity bandwidth (full wavelength) has been the main impediment to the adoption and deployment of automated optical data networking paradigms. Compounding the problem is that the optical control plane of the overlay model as well as most of the GMPLS-based routing and signaling algorithms, which have been reported by the standards bodies and corresponding research activities, were developed under the assumption that the optical layer can only provision connection requests at the full wavelength capacity.

Indeed, in spite of the euphoria about GMPLS-based routing and signaling standards and the long waited carriers promise of delivering rapid provisioning and restoration (competitive in speed to SONET rings) of optical channels to their customers, little has changed about the way that carriers have traditionally offered services to their customers today. For instance, carriers still offer individual static Point-to-Point connections, one at a time, and have to go through a lengthy provisioning process for each connection addition or deletion. This process is manual and generally takes several months to accomplish. In addition, the vision of an integrated data and optical networking paradigm and the corresponding promise of delivering rapid provisioning/ restoration has never materialized.

0-7803-8557-8/04/$20.00©2004 IEEE

This talk addresses the techno-economic limitations that have prevented the industry from achieving a large part of the promise of IP-optical networks integration and erected a market barrier to the wide scale deployment that has been expected by this point. We propose simple and efficient optical networking solutions that constitute a credible starting foundation en-route to the rapid adoption and deployment of optical data networking paradigms. Specifically, we outline and propose the key functionalities of a novel, simple, and scalable optical networking paradigm that can efficiently supports fully automated optical networking service (layer 1) at any bandwidth granularity. The proposed networking paradigm offers the potential of transporting and switching native layer 2 jumbo Ethernet frames across an MPLS/GMPLS managed and controlled "optical cloud". Transporting native Ethernet frames directly over an optical cloud will set the stage for seamless transport of Ethernet frames from the access network through the Metro and public wide area networks to another access network.

Three salient features characterize the proposed networking paradigm: 1) the optical layer must own and manage both the physical connectivity (optical resources) and logical connectivity (IP resources); 2) The optical layer can independently and selectively provision/restore diverse traffic granularity (on a per-call basis including both full-lambda and/or sub-lambda flows). This approach enables the optical layer to provide selective restoration by providing different levels of restoration (differentiated resilience) for different classes of service. 3) Customer traffic (packets/frames) must remain invisible to the network service provider. This is in contrast to Layer 2 (Ethernet, ATM, FR) and Layer 3 (IP) services that involve carrier breaking into customer's packets/frames;

We introduce two networking innovations to support the transport of native Ethernet frames across an optical cloud:

- The proposed networking paradigm migrates most of the networking functionalities and intelligence down to the optical layer including traffic engineering, switching, performance monitoring, and selective provisioning/restoration of diverse traffic granularity, all supported entirely on the optical layer's terms. While both technology trends and ongoing debate and activities within the standard bodies point compellingly to network intelligence moving up to IP layer (favoring the intelligence of routers over optical switches) [5-8], we argue that moving the networking functionality and intelligence down to the optical layer (favoring the intelligence of optical switches over routers), is more compelling in terms of simplicity, scalability, overall cost savings, and the feasibility for near-term deployment.
- Customers' packets/frames (traffic flows) presented to the ingress nodes of the carrier's core network are encapsulated into Ethernet jumbo frames with frame sizes exceeding 9000 bytes [9-10]. These jumbo frames are transported natively over the optical layer through the WDM optical channels. The core ingress nodes insert a stack of labels/fields into the jumbo frame and are stripped off at the egress nodes before the customer's packets/frames are handed over to the appropriate customer premises equipment (CPE). Each label corresponds to a color/lightpath, and the numbers of labels are equal to the number of lightpaths transporting the frame. These fields have local significance within the carrier's backbone domain. As such, the customer traffic remains invisible to the network service provider. Based on the outermost label in the stack, these frames can be switched electronically using commercially available GigE switches [9-10] capable of switching jumbo frames (layer 2 switching).

References

[1] B. Rajagopalan et al, "IP over Optical Networks: Architecture Aspects," IEEE Communications Magazine, pp. 94-102, September 2000.

[2] D. Awduche et al, "Multiprotocol Lambda Switching," Internet Draft, work in progress, November 1999.

[3] E. Mannie *et al.*, "Generalized Multi-Protocol Label Switching (GMPLS) architecture," IETF Internet draft, Mar. 2002.

[4] "IP over Optical Networks: A Framework," draft-ietf-ipo-framework-03.txt, IETF Draft, January 2003.

[5] Y. Ye, C. Assi, S. Dixit, and M. A. Ali, "A Simple Dynamic Integrated Provisioning/Protection Scheme in IP over WDM Networks," IEEE Communication Magazine, vol. 39, no. 11, pp. 174-182, Nov. 2001.

[6] A Greenberg, G. Hjálmtýsson and J. Yates, "Smart Routers - Simple Optics: A network architecture for IP over WDM," Optical Fiber Comm. Conf., ThU3, March 2000.

[7] Jennifer Yates *et al.*, "IP Control of Optical Networks: Design and Experimentation," Optical Fiber Comm. Conf., (OFC) 2001.

[8] G. Bernstein, J. Yates, and D. Saha, "IP-Centric Control and Management of Optical Transport Networks," IEEE Communications magazine, October 2000.

[9] Phil Dykstra, "Gigabit Ethernet Jumbo Frames," http://sd.wareonearth.com/phil/ jumbo.html

[10] Vish Ramamurti *et al.*, "Initial implementations of point-to-point Ethernet over SONET/SDH transport," IEEE Communications magazine, Vol. 42, pp. 64-70, March 2004.