Heterogeneous Traffic Transport over Optical Networks

Kazem Sohraby[1], Patricia Morreale[2], Mohammad Fatehi[3], and Victor Lawrence[1]
sohraby@lucent.com

ABSTRACT

As all-optical networks emerge as the high-speed networking standard, the performance of optical networks for mixed traffic requires investigation. Currently, traffic on optical networks varies from monomedia or one type of session only, to multimedia, in which multiple monomedia sessions are presented as one multimedia session. Today's internet offers elementary examples of this, with modest video and audio clips accompanying data. The performance of multimedia sessions over optical networks has not met with user expectations, primarily due to network latency as a result of the underlying protocols in the network infrastructure. This protocol incompatibility has the most significant impact on multimedia IP traffic on optical networks. In this paper, a new structure for the transport of dissimilar packet and non-packet information over optical networks is presented. The use of edge devices permits the establishment of wavelengths between end points, which allows signaling to accompany multimedia transport. The use of "intelligence at the edge" of the network for processing results in less transmission overhead with this structure than is found in other methods. Several illustrations of the utility of this method are presented, and the impact on Quality of Service (QoS) is assessed.

Keywords: optical networks, edge devices, multimedia IP traffic

1. INTRODUCTION

Future networks will certainly carry a mix of several traffic types. The specific nature and mix of these traffic streams is not well known at this time, due to their evolutionary nature. However, the traffic streams are expected to range from *monomedia* where only one type of session is supported, to *multimedia,* where a session will be composed of multiple monomedia sessions. The former case is comparable to traditional circuit or packet networks where voice, data, or video sessions are established and traffic is carried. In the current internet, early versions of multimedia sessions are observed. For example, upon gaining access to a server, an end user is permitted to access video and audio clips (as well as establishing real-time connections such as voice and video) along with data exchange. However, the existing network providing this service is composed of legacy switches, cross-connects, and transmission equipment that do not meet the performance needs of the end users. Furthermore, in addition to the exponential growth in network traffic fueled by the introduction of new applications and services, service providers are faced with significant protocol incompatibility in the case of internetworking among heterogeneous systems. Such systems include access to cellular and wireless services and access to wireline applications and services. Many service providers face the expensive task of equipment re-deployment as a response to the future capacity growth and other means of combating the needs of their customers and end users. This approach is not feasible for the majority, and "re-deploy" is not truly a solution, but a work-around.

Given the breakthroughs in wavelength-division multiplexing (WDM) and dense WDM (DWDM) as well as the ability to reach terabit capacity and higher on fiber links, the high capacity requirements of future networks demand that they be built based on optical transmission, cross-connects, and possibly switching [1, 2, 3]. The addition of future new services should be accomplished in a much shorter time period and should not require revamping of the existing infrastructure. Future networks should be traffic transparent, and new traffic types should be accommodated with, at most, minor developments at the interface to the network infrastructure. Furthermore, it should be possible to incrementally add and/or remove capacity in places where traffic demand dictates it. In order to improve network performance, the components involved in processing of traffic should also be minimized.

[1] Kazem Sohraby is with the Performance Analysis Department of the Advanced Communications Technologies, Lucent Technologies, Holmdel, NJ. Victor Lawrence is the Director of Advanced Communications Technologies, Lucent Technologies, Holmdel, NJ.
[2] Patricia Morreale is with the Advanced Telecommunications Institute, Department of Computer Science, Stevens Institute of Technology, Hoboken, NJ.
[3] Mohammad T. Fatehi is with the Optical Networking Group, Lucent Technologies, Holmdel, NJ.

In this paper, a new approach to traffic transport is proposed. Traffic may be packet or circuit data and is transported over optical links. This method differs significantly from that of using SONET/SDH in the following two ways: first, it does not require as much transmission overhead as in SONET/SDH [4]. Secondly, it allows signaling to be used in the establishment of paths between end points, in addition to allowing the transport of heterogeneous traffic. This new method allows the transport of any mix of traffic from a variety of applications/services in transport entities called ***Digital Containers***, while the core network is transparent to the details of specific protocols carried within the Digital Containers.

2. DIGITAL CONTAINER FUNCTION

Using the existing format for packet and circuit transport, this paper presents a new structure whereby heterogeneous packet and non-packet information can be transported in Digital Containers of preferably fixed size over optical links. Each Digital Container carries an identifier called *Optical Logical Channel Identification (OLCI)* field, which contains the destination address of the edge device to which the Digital Container is addressed. Each Digital Container carries multiple heterogeneous packets and/or non-packet information addressed only to the end users at the same destination edge device.

In this manner, except for occasional Digital Container routing (at the intermediate edge devices) and information extraction from the Digital Containers upon receipt (at the edge devices), individual packet processing at the edge devices, which is usually a processor intensive task, is avoided. Digital Container processing at the edge device is limited to the process of individual Digital Container delineation and possibly Digital Container routing along with the task of forwarding units of information within the Digital Container to the individual end devices.

The Digital Containers have the feature that they can be dedicated to a particular end device and/or shared among end devices connected to an edge device. When no data is available from an end user at the source edge device, or when there is no data to transmit on a particular wavelength, empty Digital Containers with a universal OLCI are transmitted so that transmission synchronization can be maintained between connected end devices. Digital Containers can also be used for the transport of maintenance, operations, and performance monitoring information. An OLCI for the same destination edge device can be used for the transport of information from the source edge device multiplexed from many users over an available wavelength. This feature allows bandwidth allocation from the pool of available wavelengths in an on-demand basis. Routing tables at intermediate edge devices allow looking up for the OLCI of the incoming Digital Containers and determining whether the Digital Container should be forwarded to another edge device on a wavelength or be forwarded to an end user at the destination edge device. In the latter case the packets are identified as in the current routing/switching methods. By allowing Digital Container processing rather than individual packet processing the edge device capacity in terms of number of packets will increase - which is required at optical transmission speeds.

In this paper, the emphasis is on the concept that the optical link is responsible for reliable transport of information and that packet processing is the responsibility of the end user devices. This separation of functionality has been shown to be crucial for combining scalablity with flexibility [5]. When packets within a Digital Container are handed to end user, processing at the end user devices commences and remains the same as in legacy systems. No changes are needed to the hardware or software at the end user devices. Figure 1 shows the interface between an edge device, which multiplexes traffic from end user ports, and the optical cross-connect (OXC). OXC is used for wavelength connection and traffic transport.

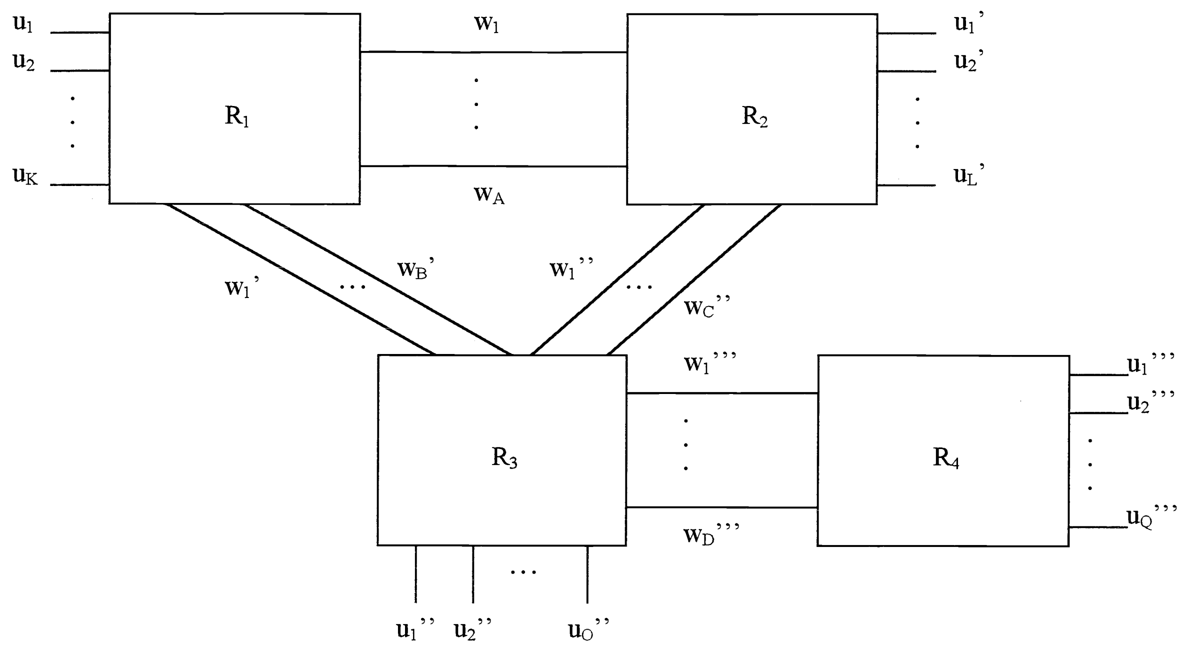

Figure 1. Network Architecture

For the purpose of connection establishment among end devices, signaling among intermediate optical cross-connects and end devices is devised. This allows establishment of concatenated end-end wavelengths between source and destination end devices. In this case Digital Container routing at the intermediate cross-connects is also avoided. With the exception of possible buffering of Digital Containers at the source end device, and processing of Digital Containers at the destination end device, no other processing at intermediate optical cross-connects is performed. Therefore, a greater measure of QoS can be provided to the end users

.

In the following, we first present the network architecture view from this optical transport method using WDM/DWDM on fiber links. The process of incorporating access traffic onto the core network facilities is then described.

3. NETWORK ARCHITECTURE

Figure 1 shows the principal elements of the core network architecture. It consists of Optical Cross-Connects (OXC) R_1, R_2, R_3 and R_4 interfacing with the electronic edge devices u_1, u_2, ... and interconnected with fiber links. Each fiber link multiplexes several wavelengths in DWDM mode. For example, the link interconnecting R_1 and R_2 carry wavelengths W_1, W_2, ..., W_A. Each wavelength may have a capacity of several Gbps or higher. The total capacity of a fiber link is in the order of several terabits per second. The electronic end devices consist of circuit switches, packet switches, and or dedicated end user links such as DS-1, DS-3, E1, which terminate at the customer premises equipment. The switching nodes as end user systems may consists of mobile switching systems in the case of wireless and cellular systems, base stations, or multiplexers which combine traffic from multiple end user devices and transport them onto the optical cross-connects.

When an optical cross-connect cannot be reached directly from another OXC, then the concatenated wavelengths interconnect the edge devices that are connected to them. For example, users connected to the OXC R_4 can only be reached via R_3. Thus traffic from R_1 or R_2 for the end users connected to R_4 can only be routed via R_3. For the most part traffic originated at the OXC R_1 and destined for R_2 or R_3 may use the wavelengths on the direct link between the cross-connects. However, in cases where such direct links does not exist, provision is made for the routing table at the edge devices to receive the Digital Containers from the incoming wavelength and transmit them onto the outgoing wavelengths toward the destination edge device. Thus, for example, traffic from R_1 destined for R_4 are routed to R_3 (either directly or via R_2) and are subsequently routed to R_4. The process of *routing at the edge* device is by using the Optical Logical Channel Identifier (OLCI). A table look-up at R_3 allows the incoming Digital Containers to determine the next edge device to which the Digital Containers are to be forwarded to. Optionally, the outgoing wavelength is also indicated in the table.

In an alternate method, the wavelengths are allocated end-end between source-destination end devices and thus routing of individual Digital Containers at intermediate edge device is avoided. In the current implementation of the Digital Container transport presented here, the OLCI field is not updated as the Digital Container travels through the network. However, the option of updating the contents of incoming OLCI and re-writing its contents at intermediate edge device is not excluded. A field in the Digital Container indicates whether the entire Digital Container is dedicated to a particular end user (addressed by the OLCI) or contains a number of packets (ATM/IP/Frame Relay, etc.) for end users at the same destination edge device.

4. DIGITAL CONTAINER FORMAT

Digital Container is a unit of information identifiable by the end devices. Contents of the Digital Containers are transparent to the OXC and the end devices (with very few exceptions.) This approach is a major departure from the traditional internet packet switching and routing, but is similar (but not identical) to the Synchronous Optical NETwork (SONET) transport mechanism [6, 7].

In SONET a variety of techniques are provided for provisioning and maintenance as well as payload structure within its 125-microsecond frame. SONET does not provide an explicit routing process at the cross-connects and/or switch/routers. Nor does SONET provide signaling for path setup between source/destination end devices for which communications should take place. Furthermore, in order to extract the packets/information from SONET, the payload is processed through embedded pointers in its path and section headers. With Digital Container transport, however, a signaling mechanism for path set up between end devices is devised, and for simplicity in packet processing, payload of a Digital Container is addressed to the same destination edge device, which supports end users that use the network. Thus, no intermediate processing of the payload is required.

SONET is regarded by some as over-kill for the reliable and almost error-free transmission media provided the optical transmission. In this paper, we consider packet re-transmission and error recovery to be a rare event on an optical link. Most recovery techniques can be delegated to the end devices and end user applications, particularly in the case of multimedia.

Three types of Digital Containers are proposed: *dedicated Digital Container* to an end user, *shared Digital Container* which is used to carry payload for several end users, and *signaling Digital Container*. The dedicated Digital Container is also referred to as a *free format Digital Container*, which carries information for a single end user in an entire Digital Container. This information is not shared with other end users. Shared Digital Container carries information from end users of the same destination edge router. The Digital Container is shared in a statistical multiplexing manner. In the signaling Digital Container case, the contents of a Digital Container is dedicated to all the end users at the destination edge router. Thus the Digital Container contents are not read/modified at the intermediate edge devices. A Digital Container is preferably of fixed size (such as 125 microseconds). Depending on the packet sizes at the source edge devices, a number of possibly heterogeneous packets can be combined into a Digital Container. Since packets are associated with end user applications in legacy systems, with their own error recovery and re-transmission mechanisms, the transport method on optical links based on Digital Containers does not require such techniques.

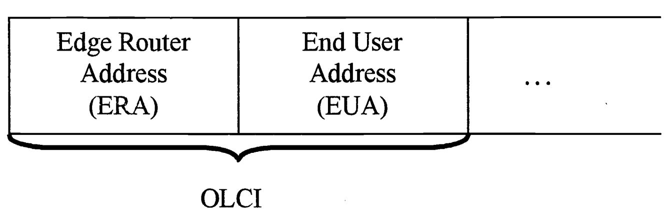

Figure 2. Digital Container OLCI

As shown in Figure 3b, when the Digital Container is dedicated to a particular end user, such as when it carries video bit stream, its identification field (OLCI) indicates address of the end user. This field is shown in Figure 2. The OLCI consists of 2 sub-fields: Edge Device Address (EDA), and End User Address (EUA). For a dedicated Digital Container a Payload

Control Field (PCF) is set to 1 indicating that the Digital Container carries only 1 "unit" of information dedicated to a specific end user. The "payload" format of the Digital Container in this case is free. When PCF is set to a value greater than 1, then this indicates that the Digital Container is allocated to multiple end users at the Edge Device (see Figure 3a). Depending on the Digital Container size, there can be unused portions of the Digital Container that may not contain packet or end user information.

Boundary delineation on a packet and assessment of the packet contents is the responsibility of the end user and is only minimally performed at the destination edge router during packet sorting and forwarding to the end user devices. The destination edge device extracts the individual packets from the Digital Container and forwards them to the end users connected to it only based on the individual packet addresses. The Digital Container boundary can be determined using the OLCI fields and the fixed Digital Container size. The OLCI and Digital Container size, in addition to other available synchronization information, can be used for network synchronization and clock recovery when necessary.

OLCI	PCF=N	OVF	Packet #1	Packet #2	...	Packet #N

Figure 3a. Digital Container Format (PCF > 1)

OLCI	PCF=1	OVF=0	Information Bits (Such As Video Stream)

Figure 3b. Digital Container Format (PCF=1)

OLCI	PCF=0	OVF=0	Signaling Type (ST)	

ST Values = 0, 1, 2, ...

Figure 3c. Digital Container Format with PCF = 0

In cases where a packet is larger than the remaining size of a Digital Container, the packet excess free overflows to the next Digital Container. Digital Container overflow is indicated in Digital Containers that are not used for signaling by setting the Overflow Field (OVF). Digital Containers from the same source edge device can be forwarded to the same destination edge device using any available wavelengths between them as long as a path exists.

The final Digital Container type referred to as signaling is identified by the PCF field, which is set to zero. The OVF field for this type is ordinarily set to zero also. In this case, an additional field called Signaling Type (ST) contains values that refer to the message signaling type encapsulated in the Digital Container. In addition to the optical network signaling encompassed by the Digital Containers, end user applications and devices may utilize their own signaling technique, after a path between source and destination is set up. For example, when using the existing optical network infrastructure, call set-up messages can use the Digital Containers for SS7 signaling message exchange among the end devices (such as circuit switches at the source and destination OXCs) for call set-up, teardown, etc.

In the following sections, the procedure for end-end optical wavelength set up between source-destination edge devices, and the process of logical channel establishment when Digital Containers are processed at intermediate edge devices is described. The latter is used when end-end set of optical wavelengths that can be concatenated in order to create a path between source-destination edge devices is not available.

5. ADDRESSING AND EDGE DEVICE ROUTING

Since packets carried in a Digital Container have their own addresses, when a Digital Container is addressed to a destination edge device, at the destination individual packets are forwarded to the end user devices. End user devices are identified through the port to which the user device is connected, or a virtual address contained in the packet (as is the case with ATM cells). When the entire Digital Container is dedicated to an end user, the address of the end user is indicated in the second part of the OLCI field, as shown in Figure 2. As shown in this figure, the first part of the OLCI field contains the Edge Device Address (EDA) and the latter part the End User Address (EUA). When the Digital Container is initially generated at the source edge router, the EDA field is set to the final destination edge device address. Alternatively, EDA field can be set to the address of the edge router next to be visited in the network. In the former case as the Digital Container hops through the network, the EDA field remains the same and is not modified. In the latter case as the Digital Container arrives at the next edge device, its EDA field is updated based on a routing table. The former method requires less processing at the intermediate edge devices and may be preferred for certain applications/services. In the case where the Digital Container is shared among multiple end users, the EUA field is unused. In this case, a field next to the OLCI, which is called Packet Counter Field (PCF), indicates the number of packets contained in the Digital Container. When the entire Digital Container is allocated to the same end user PCF field is set to 1 and the Digital Container EUA is set to the address of the end user device. The details of edge device processing of a Digital Container are not presented here. It is assumed that the decision making process such as for routing is performed off-line or in real-time. An overflow (OVF) field is included in certain Digital Containers. The values of this field is determined as follows:

OVF = 0 — For signaling (PCF=0) and free format (PCF=1) Digital Containers.

OVF=1 — When PCF>1 and the Digital Container contains the first part of an incomplete packet along with other packets.

OVF=2 — When PCF $\geq$1 and the Digital Container contains the latter portion of the incomplete packet. If a packet size exceeds the size of two Digital Containers, then the Digital Container containing the first portion of packet sets PCF>1 and OVF=1, the next Digital Container contains the remainder of packet (and no other packet) with PCF=1 and OVF=2. The third Digital Container containing the remainder of the Digital Container (and possibly other packets) will have PCF>1 and OVF=2.

In order to prevent the last part of an incomplete packet carried with the first part of another incomplete packet in the same Digital Container, the Digital Container containing the last part of an incomplete packet is released for transmission without other packets (partially filled Digital Container). When a Digital Container is used in free forward mode, PCF=1, always OVF=0. In the case where end-end optical wavelengths between source and destination edge devices have been set up, such intermediate step processing at edge devices are avoided [8]. Packets are loaded into the Digital Containers at the source edge device and are extracted at the destination edge device.

6. HANDLING HETEROGENEOUS TRAFFIC

DESTINATION EDGE DEVICE At the destination edge device, Digital Containers destined for the end users connected to it are forwarded to the end user ports as represented in the destination edge device and the packet addresses. Upon receiving a Digital Container with PCF>1, the edge device takes the responsibility of removing the individual packets from the Digital Container content and forwarding them to the end user ports on the edge device.

In order to properly identify individual packets, at the source edge device where packets are loaded onto the Digital Containers, a unique byte (such as a flag in the form '01111110') is added between each two consecutive packets. At the destination edge device, as the unique pattern is encountered, it is assumed to be the boundary between two packets and each one is sent to the appropriate processor for handling. For this purpose, N copies of each packet are reproduced and handed to

N separate processes. Each process in the electronic domain is responsible for handling a particular packet type. For example, the process of identification for two types of packets, ATM and IP, might be as follows:

One copy is submitted to the ATM cell recognition processor. In this case the process of delineating the cell is performed as it is done in an ATM switch or destination device. The number N represents the total number of different type packets handled by the edge device. If a packet is not recognized by any of the N processes, then two possibilities are examined:

1. A packet of a type unrecognizable by the destination edge device is contained in the Digital Container.
2. A packet or combination of packets were corrupted during Digital Container transmission or processing and are not recognized by the destination edge device.

In either case, a decision may be made as to either request retransmission of the entire Digital Container or portions of it or to discard the Digital Container entirely. Such issues will be handled by the exception routine at the edge device.

Among other packet candidates for which processing at the destination edge device is suggested are: SONET frames contained in a Digital Container, packets from wireless end points, and TDM samples. The process (hardware/software) which handles the different Digital Container contents is assumed to be the responsibility of the edge device. Furthermore, all such processes occur in parallel in order to minimize the task of processing packets at the destination. Packet ordering, and handling of network transport level tasks is the responsibility of the end user.

SOURCE EDGE DEVICE At source edge device, packets are loaded into Digital Containers destined for a given destination edge device. As packets appear at edge device input ports from end user devices, the edge device determines whether an end-end wavelength path is available. If so, the packet is included in the outgoing Digital Containers and handed to the output port at the edge device for the wavelength. Incoming packets may be included in a Digital Container on a first-come, first-served basis (in which no priority is defined), or precedence can be given to the higher priority packets. A unique bit pattern (such as '01111110' used as the HDLC flag) is inserted between two consecutive packets for delineation. As packets are assigned to the different Digital Containers, the Digital Container is released at every frame interval whether it is full or not. If a packet size is such that portion of it is included in a Digital Container and the remainder is in the next Digital Container, it is so included and assumed to be the responsibility of the destination edge device and user application to reassemble them. In this case, the Digital Container in which the first part of the incomplete packet is loaded into has its OVF set to 1. The OVF is set to 2 for the subsequent Digital Container(s) carrying the remainder of the packet. Ordinarily when free format Digital Container is used, or there is no packet overflow out of, or into a Digital Container, OVF is set to zero.

Voice traffic from incoming TDM trunks are loaded onto the Digital Containers either as packet (after conversion at the source edge device, such as in the form of Voice over IP, or Voice over ATM, etc.), or directly as TDM voice samples. The former case is handled as a regular packet in a Digital Container. The latter is handled in a free format Digital Container (PCF=1). In the latter mode, handling of DS1, DS3, and various SONET rates is expected.

7. QUALITY OF SERVICE

During connection establishment, wavelengths at different cross-connects are concatenated so that an end-to-end connectivity between source and destination edge routers can be maintained. Thus, ordinarily there is no buffering of Digital Containers at the intermediate edge devices. On occasions where sufficient bandwidth may not be available (all Digital Containers are full) and/or no end-to-end set of wavelengths are available between source and destination edge devices, Digital Containers are stored and forwarded at the intermediate edge routers. This requires conversion from optics to electronics in order to determine the OLCI and the destination edge device. This is the case for services where no QoS is guaranteed, or desired by the application. It is expected that with WDM and availability of several wavelengths on a fiber strand and particularly with proper network engineering it is possible to find dedicated end-end wavelengths (or a concatenation of wavelengths) between source and destination edge devices [9, 10]. The process of effectively determining the available dedicated end-to-end wavelengths or appropriate concatenation of wavelengths is part of ongoing research.

8. CONCLUSIONS

In summary, using the existing format for packet and circuit transport, this paper presents a new structure in which heterogeneous packet and non-packet information can be transported in *Digital Containers* of preferably fixed size over optical links. Each Digital Container carries an identifier called the *Optical Logical Channel Identification (OLCI)* field, which contains the destination address of the edge device to which the Digital Container is addressed. Each Digital Container carries multiple heterogeneous packets and/or non-packet information addressed only to the end users at the same destination edge device.

Using this approach, with the exception of occasional Digital Container routing at the intermediate edge devices, and for extracting information from the Digital Containers upon receipt at the edge devices, the processor intensive task of individual packet processing at the edge devices is avoided. Digital Container processing at the edge device is limited to the process of individual Digital Container delineation and possibly Digital Container routing along with the task of forwarding units of information within the Digital Container to the individual end devices.

Theoretical consideration of this effort has concluded that this would be an ideal manner in which to provide heterogeneous traffic transport over optical networks, while providing dynamic support for new service provisioning. Scalability, without service degradation, for both services and users served seems assured. While all-optical networks are expected to be the norm in the future, we must be able to provide new services in today's hybrid network environment of electronic and optical network infrastructures. Current research efforts include gathering network performance data for the multimedia service delivery using this methodology.

9. REFERENCES

1. Bonenfant, Paul and Antonio Rodrigues-Moral, "Optical Data Networking", *IEEE Communications*, March 2000, Vol. 38, No. 3, pp. 63-70.
2. Elmirghani, Jaadar M.H., and Hussein T. Mouftah, "Technologies and Architectures for Scalable Dynamic Dense WDM Networks", *IEEE Communications*, February 2000, Vol. 38, No. 2, pp. 58-66.
3. Yao, Shun, Biswanath Mukherjee, and Sudhir Dixit, "Advances in Photonic Packet Switching: An Overview", *IEEE Communications*, February 2000, Vol. 38, No. 2, pp. 84-94.
4. Anderson, Jon, James S. Manchester, Antonio Rodriguez-Moral, and Malathi Veeraraghavan, "Protocols and Architectures for IP Optical Networking", *Bell Labs Technical Journal*, January- March 1999, pp. 105-124.
5. Li, Bo, and Yang Qin, "Traffic Scheduling in a Photonic Packet Switching System with QoS Guarantee", *Journal of Lightwave Technology*, Vol. 16, No. 12, December 1998, pp. 2281-2295.
6. Al-Salameh, Daniel Y., Mohammad T. Fatehi, William J. Gartner, Stan Lumish, Bruce L. Nelson, and Kamal K. Raychaudhuri, "Optical Networking", *Bell Labs Technical Journal*, January-March 1998, pp. 39-61.
7. Alferness, Rod, Paul A. Bonenfant, Curtis J. Newton, Kevin A. Sparks, and Eve L. Varma, "A Practical Vision for Optical Transport Networking", *Bell Labs Technical Journal*, January-March 1999, pp. 3-18.
8. Gambini, Piero, Monique Renaud, et. al., "Transparent Optical Packet Switching: Network Architecture and Demonstrators in the KEOPS Project", *IEEE Journal on Selected Areas in Communications*, Vol. 16, No. 7, September 1998, pp. 1245-1259.
9. Ramaswami, Rajiv, and Galen Sasaki, "Multiwavelength Optical Networks with Limited Wavelength Conversion", *IEEE/ACM Transactions on Networking*, Vol. 6, No. 6, December 1998, pp. 744-754.
10. Banerjee, Dhritiman, and Biswanath Mukherjee, "A Practical Approach for Routing and Wavelength Assignment in Large Wavelength-Routed Optical Networks", *IEEE Journal on Selected Areas in Communications*, Vol. 14, No. 5, June 1996, pp. 903-908.