The Design of an All-Optical Packet Switching Network

Ruth Van Caenegem, Didier Colle, Mario Pickavet, and Piet Demeester, Ghent University

Konstantinos Christodoulopoulos, Kyriakos Vlachos, and Emmanouel Varvarigos, University of Patras

Leontios Stampoulidis, National Technical University of Athens

Diego Roccato, Telecom Italia Lab

Ruth Vilar, Fibernet, Spain

ABSTRACT

Recent advances in the all-optical signal processing domain report high-speed and nontrivial functionality directly implemented in the optical layer. These developments mean that the all-optical processing of packet headers has a future. In this article we address various important control plane issues that must be resolved when designing networks based on all-optical packet-switched nodes.

INTRODUCTION

Despite the recent economic slump, the demand for telecommunication services continues to grow steadily. Moreover, networks are quickly shifting from voice-centric to data-centric. In this developing scenario, the role of synchronous optical network/synchronous digital hierarchy (SONET/SDH) is likely to decrease, and the optical transport network will provide a transport infrastructure for legacy and new IP services directly.

In addition, the use of wavelength-division multiplexing (WDM) offers massive bandwidth through the parallel transmission of high-bit-rate channels on the same fiber, at very attractive costs per bit. Currently, we are witnessing the shift from purely point-to-point WDM systems to switching systems with real networking functionality at the optical level. Additionally, because of the migration of most services over IP and the evolution of optical circuit switching, core networks tend to have a two-level IP-over-optical architecture. Optical cross-connects (OXCs), however, offer granularity only at the wavelength level. For networks that must provide a large variety of service levels in a highly dynamic environment, this has a baleful impact on the dimensioning of the network and the size of the OXCs.

Packet switching (PS) offers the flexibility and bandwidth efficiency called for in this environment. Current PS networks transport IP packets in the optical domain but convert the packets to the electronic domain to route them at intermediate nodes. Despite the bandwidth flexibility provided by this approach, the high-speed optical/electronic/optical (OEO) conversions required are costly or even infeasible as we move toward higher speeds. A more viable long-term approach is optical packet switching (OPS), where the packet payload remains in the optical domain during the entire journey of the packet.

Even though the payload in the OPS paradigm remains in the optical domain, the header containing the routing information may still require electronic processing. Therefore, in the first OPS approaches, only the packet header is converted to the electronic domain. Often this header is modulated at a lower speed than the line speed, allowing for less costly OEO conversions. IP giga/tera routers are among the most expensive devices in an operator's network, and they face critical switching speed limitations. To switch faster and more efficiently, multiprotocol label switching (MPLS) uses local labels that do not refer to an absolute node address in the network. MPLS is a connection oriented protocol — before a packet is sent, a label switched path (LSP) is set up — in contrast to typical connectionless IP packet networks. For optical communications, generalized MPLS (GMPLS) [1] extends MPLS by defining labels that are no longer carried strictly in data form but can be identified by a time slot, wavelength, or port. Even though (G)MPLS forms a good and fast solution, it does not by itself solve the mismatch between the switching speed of the router and the data speed of the fiber (table lookup procedures are still time consuming). In an attempt to overcome this, research started focusing on all-optical packet switching (AOPS) and all-optical label swapping (AOLS), where the packet header (label) is processed (all) optically [2].

AOLS implements the packet-by-packet routing and forwarding functions of (G)MPLS directly in the optical domain. Ideally, this approach can route packets independent of bit rate, packet format, and packet length. Advantages of AOPS are particularly evident in core networks, where AOLS can be used to replace both OXCs
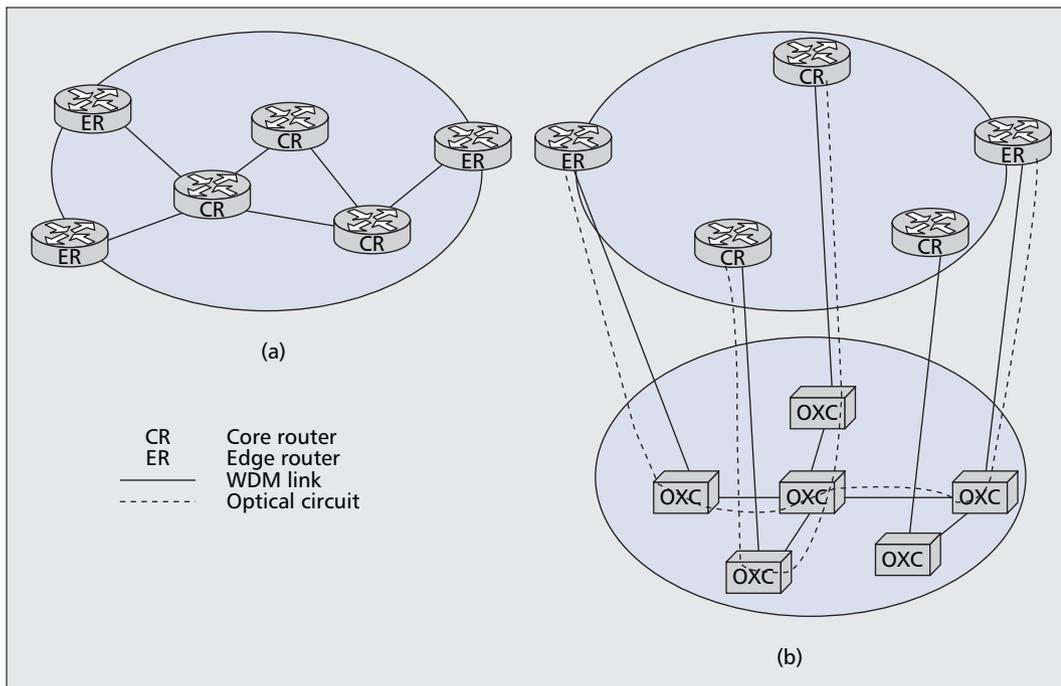
■ Figure 1. *Architectures of an optical transport network employing OPS nodes: a) pure OPS; b) OPS + optical circuit network.*

and IP core routers. With regard to OXCs, AOLS is a multiclient transport platform used by IP, SDH, gigabit Ethernet (GbE), fiber channel (FC), and asynchronous transport mode (ATM) clients to manage the bandwidth more efficiently. With regard to IP routers, AOLS offers an aggregation layer. The IP network is simplified (avoiding core devices) through the transport infrastructure realized by OPS nodes. From a networking perspective, an all-optical node is defined as a high-throughput packet switching node. Due to the wide use of optics, the architecture scales well with the payload bit rate and number of ports. However, processing capabilities are rather limited, and the node essentially limits itself to a forwarding function based on the label of the incoming packets. In metro-regional networks, transport functionality is currently realized by means of different solutions (SDH, transmission over dark fiber, WDM rings, etc). Switching devices (GbE switches, IP routers, ATM switches) perform complex operations related to service functionality (e.g., accounting and bandwidth limitation). At present, this functionality is beyond the capabilities of all-optical devices.

This article focuses on the design of an all-optical network and the construction of AOPS nodes, and presents the appropriate control plane protocols. It also presents a dimensioning study of the data plane and elaborates on the opportunities offered by an all-optical network from an operator's view. This work is performed in the framework of the Information Society Technologies (IST) Label Swapping Employing Optical Logic Gates in Network Nodes (LASAGNE) project, whose aim is to bring to the field of AOLS and AOPS the use of optical gates and memory elements to perform intelligent functionality requiring speed and processing power.

NETWORK LAYERS AND CHARACTERISTICS

This section describes and compares different options for the all-optical network data plane and control plane design.

ALTERNATIVES FOR THE OPTICAL CORE NETWORK

Figure 1 illustrates two options for the design of a core network made up of all-optical packet nodes:
• **Pure packet-switched approach:** The OPS nodes are connected directly using WDM links.
• **Mixed packet-circuit approach:** An optical packet layer and an underlying optical circuit layer exist. OPS nodes are connected using optical WDM circuits provided by the circuit layer.

The first approach has the advantage of simplifying the network architecture. It is constructed of fewer network layers and is considered an ultimate long-term scenario. Nevertheless, in a number of situations the use of circuits has advantages. Some operators deliver lambda services, which are layer 1 (L1) point-to-point connections. Interesting is an L1 virtual private network (VPN), representing L1 multipoint connectivity, which is realized by means of a number of point-to-point, static or dynamic circuits. Because of their high profitability, these services are not expected to disappear in the short-medium term.

Another scenario for the mixed packet/circuit architecture is the migration OCS towards OPS. Even if the OPS network is supposed to carry every type of client traffic, the connection of client equipment to OPS edge routers would

require a certain amount of time to evolve, as well as the definition of a migration policy. In this transient phase, some clients may use the OPS network, while others may still connect through a circuit-based backbone.

OPS REFERENCE PROTOCOL ARCHITECTURE: CORE AND EDGE FUNCTIONALITY

In our discussion of the OPS solution, we distinguish between core and edge functionality. The all-optical processing of optical packets (i.e., label processing, label swapping, packet routing, and related issues), obviously must be implemented in every core node. The header of an optical packet contains the information required to help the packet reach its destination. The basic forwarding entity is a virtual connection, called the OPS LSP, according to GMPLS, an optical analogy of an MPLS LSP or an ATM virtual circuit. The basic operation at the ingress side of the all-optical network is the insertion of
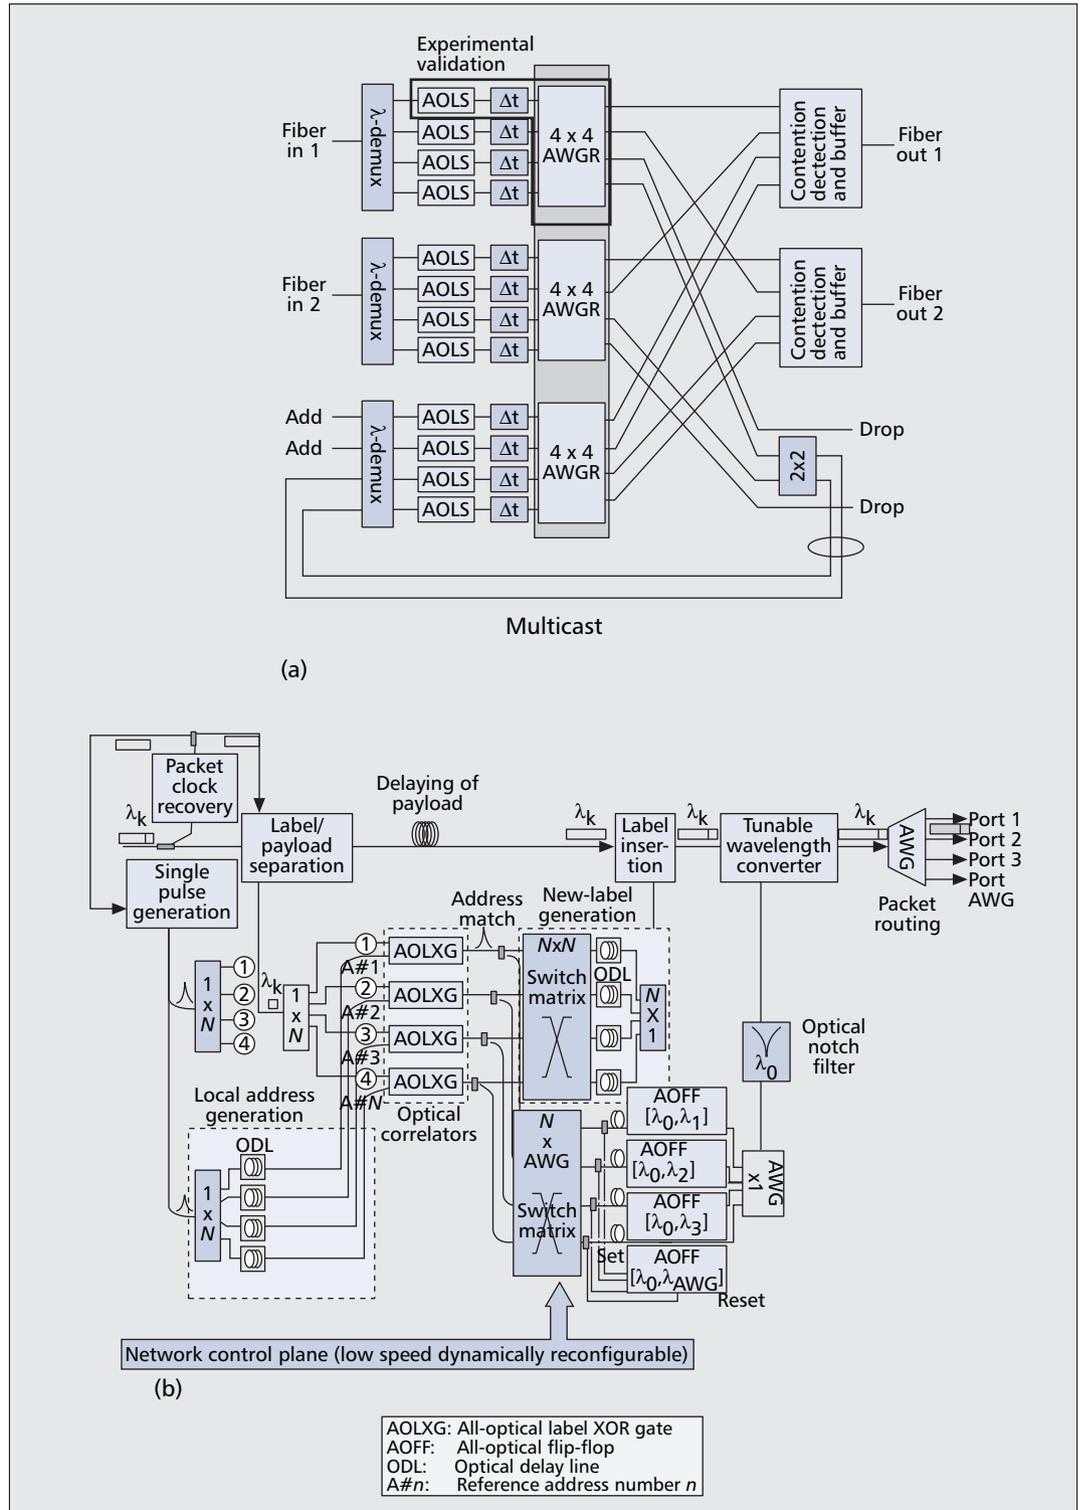
■ Figure 2. *a) AOLS LASAGNE node; b) the AOLS block architecture.*

client protocol data units (PDUs) in optical packets. To enable the receiving node to extract the correct client PDUs (insertion/extraction is supposed to be performed electronically), some mapping information is also inserted (edge header).

The most obvious approach is to define two types of nodes:
• **Core nodes**, which are all-optical devices that implement only core OPS functionality
• **Edge nodes**, consisting of a core node with electronic client cards that perform the mapping, stuffing, and edge header generation

An important function performed at the network edge is the interworking between client protocols and OPS. Besides the mapping of client PDUs in OPS packets, this function also includes the translation of address information. The procedure creates a translation table recording the relationship between incoming client signals and OPS LSPs.

THE LASAGNE NODE ARCHITECTURE

The LASAGNE AOLS node [3] was proposed as a candidate solution for an AOPS network and can be used as a core router in the scenario presented in Fig. 1a. The LASAGNE node concept is shown in Fig. 2a. At the entrance of the AOLS block (Fig. 2b), the optical label and payload are separated. The extracted optical label is fed to a bank of all-optical logic exclusive-OR (XOR) gates (AOLXGs) [3] and is compared to a local address word. This local address is created by one pulse that is sent into an optical delay line (ODL), comprising a set of interconnected fiber delay lines, couplers, and splitters. After this comparison, a high-intensity pulse appears at the output of the AOLXG correlator that corresponds to the matching address. This pulse feeds a control block that drives a wavelength converter and generates a new optical label (which is generated again in an ODL). Two switches provide the flexibility to assign different outgoing labels and wavelengths to the same incoming label. The switch settings are controlled by the control plane. When the payload and new label are combined, the packet is sent to an arrayed waveguide grating (AWG) that forwards the packet to the corresponding output.

## LASAGNE NODE DIMENSIONING AND COST

This section presents a dimensioning study of the all-optical LASAGNE node described previously. Additionally, it provides a techno-economic comparison of the LASAGNE node architecture to other packet switching alternatives.

THE ALL-OPTICAL NODE DIMENSIONS

From the description of the AOLS block (Fig. 2b), it is clear that for each different label entering the node, an incoming ODL, a correlator,
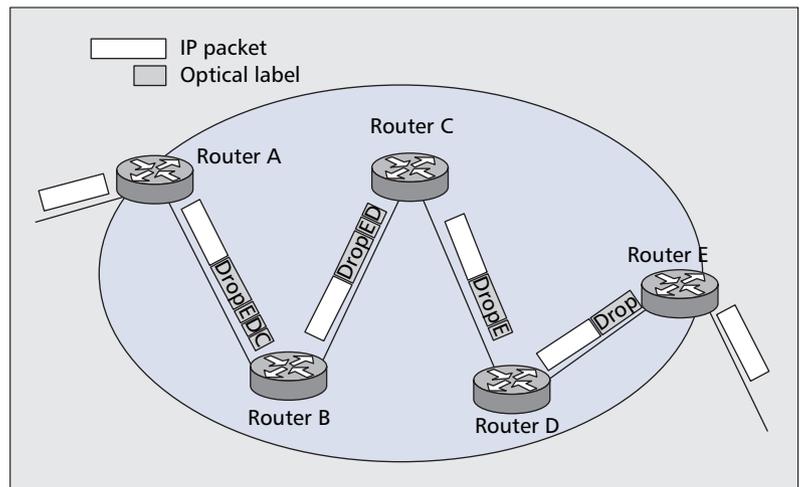
■ **Figure 3.** *Label stripping strategy.*
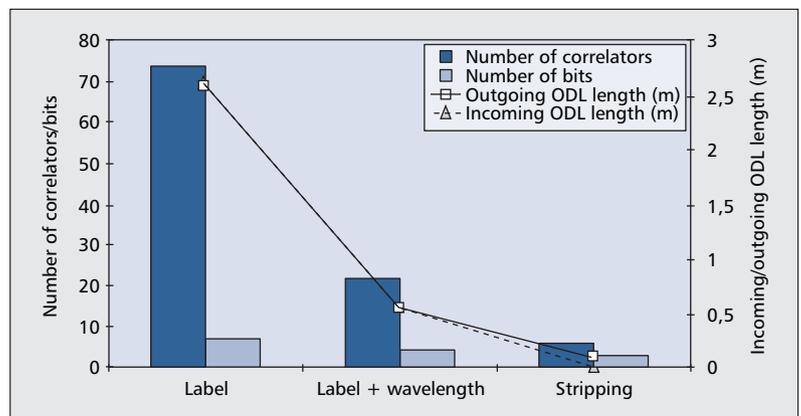
■ **Figure 4.** *Number of components needed per wavelength for each label switching strategy.*

and an outgoing ODL must be installed. The dimensions of the AOLS block thus relate directly to the number of LSPs passing through the node. This section describes how AOLS blocks can be dimensioned economically.

Label Switching Strategies — The routing function includes the decision about the outgoing port and wavelength to which a packet will be forwarded and the new label it will carry. This decision is based on an internal routing table, the current label of the packet, the incoming wavelength, and the incoming fiber port, or a combination of these.

We propose and evaluate three label switching strategies. The first strategy corresponds to the case where the wavelength domain is used to resolve contention. Contention occurs when two or more packets compete for the same wavelength on the same outgoing fiber port. If we can find appropriate wavelengths on which to place the contending packets, it is possible to resolve contention without buffering. In this case the wavelength on which a packet arrives is not fixed; thus, the wavelength information cannot be used for the routing function. We refer to this as the *label swapping* strategy. We refer to the case where both the label and wavelength information are used for the routing

function as the *label wavelength swapping* strategy. In these strategies, the incoming label (with or without information on the wavelength) is used to swap the label into a new label and decide the outgoing fiber port (and possibly the fixed outgoing wavelength). The third switching strategy (Fig. 3) switches a packet through the network based on an *end-to-end label*. This label is a concatenation of multiple local labels. In each intermediate node, the AOLS block strips off the first bits (this is the corresponding local label) of the end-to-end label and decides the output port to which the packet should be forwarded. We refer to this as the *stripping* strategy. At this point, we must note that all examined switching strategies use labels in a general way in accordance with the GMPLS framework.

A Node Dimensioning Point of View — To assess the resource requirements of AOLS blocks, we investigate the impact of the aforementioned labeling strategies on the European Network and the traffic-demand scenarios proposed in [4]. We assume the wavelengths have a bit rate of 40 Gb/s. Our results (Fig. 4) are presented in terms of resources required per wavelength. The dimension indicators that define the physical dimensions and implementation difficulty of the node are the following:
- *Number of correlators*: the number of correlators required to distinguish the different possible incoming labels
- *Length of incoming/outgoing ODLs*: the total number of ODLs required to provide the local addresses/outgoing labels multiplied by the delay line length of one bit and the number of label bits

Number of ODLs and Correlators — Figure 4 presents the resources required for the wavelength with the maximum number of correlators and that with the minimum number of correlators. The advantages obtained by using the wavelength as a routing information parameter can be seen clearly. The *label wavelength* strategy reduces the number of correlators to less than half the corresponding number required by the label strategy. For example, consider the case of a total of four transient LSPs at a node with two ports and two wavelengths per port. When the wavelength is not used as a routing information parameter, each LSP receives a different label, and four different labels are required. When the wavelength is used as a routing information parameter, only two labels are required. When more labels are required, these labels must be longer, resulting in more complex components. This is indicated by the length of incoming/outgoing ODLs in Fig. 4. Nevertheless, the stripping strategy that uses end-to-end labels is by far the best among all switching strategies. The possible incoming labels are restricted to the number of outgoing fiber ports in the node. Moreover, the length of the incoming ODLs is considerably smaller because the length of the labels is shorter, and there is no need for outgoing ODLs because the label does not need replacement.

TECHNO-ECONOMIC STUDY ON THE COMPONENT COST

The preceding study shows that all-optical nodes are very hardware-consuming, and prudence is required when labeling the packets and deciding the switching strategy to use. This section describes a cost comparison between all-optical packet switching nodes and packet switching nodes with electronic header processing. The cost function of an optical component is modeled by the number of fiber-to-chip couplings (FCCs) of the component. This model is based on the hypothesis that packaging, and more specifically the number of interconnections to the outer world, dominates the cost of optical components.

The total cost of a node is formed by the sum of the individual component costs. To model the cost of the electronics, we must consider the cost of one OEO conversion. We assume that multiplying the bit rate by four induces a 2.5 times OEO cost increase. To sum both costs, we multiply the electronic cost by P. P is the parameter that represents the cost ratio between the cost for an FCC and the cost for an OEO ($P = cost_{OEO}/cost_{FCC}$ and $P = 1$ for a header bit rate of 155 Mb/s).

A comparison of the cost estimates shows that the cost of an all-optical node depends on the number of bits in the label in an exponential way. A label bit length reduction reduces not only the node hardware requirements but also its cost. Because label stripping uses smaller labels, the cost for a label stripping node is less than that of a label swapping node. In Fig. 5 the cost of the all-optical node designs are compared to that of a packet switching node with electronic header processing and also that of a fully electronic node (IP over WDM). Figure 5 shows that for an increasing header bit rate, the intersection value for the parameter P increases, or the cost difference between the all-optical node and the other packet switching nodes decreases.

THE LASAGNE CONTROL PLANE

LABEL SWITCHED PATH SETUP MECHANISM
Packets are forwarded in the LASAGNE AOLS network over established LSPs in accordance with the GMPLS framework. The control plane employs a two-way signaling protocol where labels are distributed in a *downstream-on-demand* mode [5], and the path of the LSP is defined by the source node (source routing). The LSPs are established either prior to data transmission (control-driven) or upon detection of a certain flow of data (data-driven). Path selection uses network utilization information communicated to the source via the bandwidth monitoring mechanism described later.

RESOURCE RESERVATION
In addition to the label distribution mechanism, the LASAGNE LSP setup protocol also includes bandwidth reservation features. Traffic in the LASAGNE network is categorized as follows:
- If the exact size of the data to be transferred is known, the time duration of the LSP can be computed, and thus the LSP can be
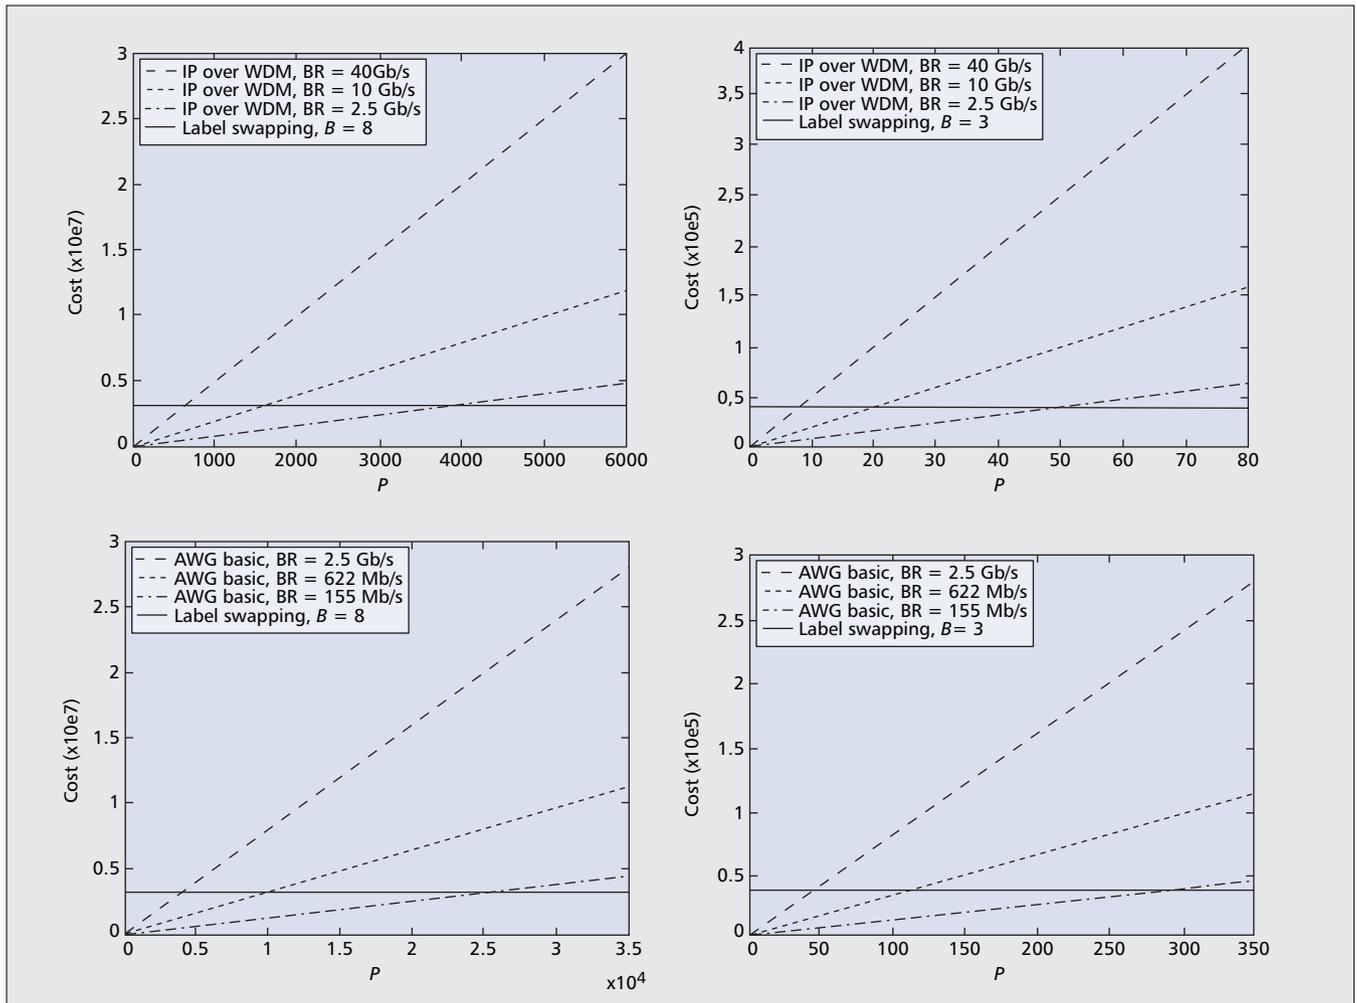
■ Figure 5. *Cost comparison for different packet switching node architectures and different values of the header bit rate (BR). The nodes have 4 fibers, 32 wavelengths/fiber, 8-bit label length for label swapping, and 3-bit label length for label stripping.*

implicitly released when it expires. This type of traffic includes bulk transfer services and the transfer of data bursts. In the latter case we assume that several packets belonging to the same forwarding equivalence class (FEC) can be aggregated at an ingress node into an AOLS super-packet before being released into the core network at a constant rate or even as a single burst. We refer to this type of traffic as type 1 traffic.

• If the size of the data to be transmitted is not known but the rate at which the data will be sent is known and constant, the SETUP packet communicates to all nodes across the path the requested bandwidth and establishes the LSP for an unspecified duration. We refer to this type of traffic as type 2 traffic.

• If the size of the data to be transmitted is unknown and the data rate changes with time but has a known average value, the SETUP packet communicates to all nodes across the path the requested average bandwidth and establishes the LSP. This is used as an active open session, where the LSP is periodically refreshed, and capacity can be renegotiated via new SETUP messages. We refer to this type of traffic as type 3 traffic.

An intermediate node rejects a new LSP set-up request when the sum of the new bandwidth request plus the reserved bandwidth of established LSPs exceeds (a fraction or the whole of) the link capacity. Data of the first type can be transmitted as packets at a certain constant rate, and the LSP and bandwidth can be *hard reserved* exactly for the time duration they actually are going to be used, increasing resource utilization efficiency. Traffic of the second type also hard reserves the LSP and the bandwidth but releases the reserved bandwidth when the LSP is explicitly torn down. To ensure a low loss rate for type 1 and 2 traffic, the contention resolution block of the LASAGNE node (Fig. 2a) resolves packet-level contentions. However, buffer size requirements for these types of traffic are small, because the role of the buffer basically is to synchronize incoming constant rate flows that never exceed the link capacity (bandwidth hard reservations). Finally, for traffic of the third type, the bandwidth requested during LSP set up is only an estimated average rate that is expected to be utilized. The traffic rate of an LSP may change due to the dynamic nature of the packet arrival process. The average traffic rate estimate is used by intermediate nodes for book-keeping and admission control purposes. In this case, the requested capacity is *soft reserved* during the LSP

set-up process, meaning that it is used by multiple LSPs of type 3, on a demand basis, through statistical multiplexing. This is in contrast to *hard reservations* of bandwidth made for constant rate LSPs (type 1 or 2), where a specific fraction of the capacity is dedicated to them, and they have absolute priority over its use.

The necessity to use small optical buffers in the LASAGNE AOLS node is in contrast to the notion that large buffers lead to good statistical multiplexing gains and low packet losses. However, recent work has shown that a buffer of only 20 packets should be sufficient in typical core networks if appropriate rate control is used at the ingress nodes [6].

For the LASAGNE AOLS network scenario, the main idea is that bandwidth reservations are used to handle long-range variations in the traffic rate, so that the total long-term average rate of the traffic passing through a link does not exceed its capacity. The LSP set-up mechanism blocks requests when the sum of the newly requested rate and the total existing bandwidth reservations exceeds the wavelength capacity. In this way, the network level provides flow control, and the contention resolution block at AOLS nodes handles packet-level contentions or short-time scale traffic variations with a minimum size of employed optical buffers. However, in this case, packets still may be dropped at intermediate nodes due to an overflow of the limited buffer space. Figure 6a shows the way the resource reservation mechanism is applied in combination with the LSP set-up process.

The resource reservation mechanism requires each node to keep a record of the bandwidth reservations made on its outgoing links. We let $r_{ij}(t)$ denote the capacity reserved on wavelength i of link j by LSP sessions with known or unknown duration, at time t relative to the present time. The function $r_{ij}(t)$ is called the *utilization profile* of the wavelength i on link j and is a *piece-wise* constant function (Fig. 6b) that can be stored as a linked list of numbers.

For type 1 traffic, which resembles the optical burst-switching (OBS) paradigm, the requested capacity is hard-reserved in the core by applying a two-way reservation scheme, called the Efficient Burst Reservation Protocol (EBRP) [7], simultaneously with the LSP set-up process. EBRP employs *delayed* and *in-advance* reservation mechanisms to efficiently utilize the following available resources:

- **Delayed reservations**: Outgoing capacity is reserved only for the requested duration, and it is released after the data traverses the node. EBRP negotiates during the downstream set-up phase the reservation duration (RD), which may exceed the required holding time. In that case, strict timed requirements are relaxed and restored during the acknowledgment phase.
- **In-advance reservations**: If the capacity is not available at the time requested, it may be reserved at a time in the future. If the first time at which the requested capacity becomes available exceeds the maximum delay requirements of the session, the bandwidth request is rejected.

SETUP packets traversing a longer path, or SETUP packets requesting bandwidth for a longer duration, will have a higher risk of not finding the appropriate resources. To increase the probability of successfully reserving the appropriate resources and establishing the path, EBRP introduces a reservation duration (RD) parameter that is chosen dynamically for every request. We have studied various functions that can be used for selecting the initial value of the **RD** like: $RD(T_{data}, h) = k \cdot t = T_{data}^m \cdot h_n, k \geq 1$, where k, m, and n are constant parameters, h is the number of hops on the path to be followed, and T_{data} is the transmission duration (known for type-1 traffic).

MONITORING MECHANISM

If delayed and in-advance reservations are to be used for reserving bandwidth in AOLS networks without temporal buffering capabilities in the traversing nodes, new traffic monitoring mechanisms and routing algorithms must be employed. For example, in the LASAGNE network concept, path selection based on the shortest path idea is not appropriate, because it would lead to low resource utilization and most importantly, to high-packet drop ratios in the overloaded intermediate nodes or to long queuing delays at the edge nodes. The utilization profile shown in Fig. 6b is a step-wise function, with discontinuities at the points at which reservations begin or end and is updated dynamically with the admission of each new LSP. Since the LASAGNE signaling protocol employs source routing, it is essential for the monitoring mechanism to have access to the utilization profiles of the LASAGNE network links. In [8] we present a way to enhance the Simple Network Management Protocol (SNMP) to monitor and communicate the available bandwidth of the LASAGNE network links. The utilization profiles of the nodes collected through the enhanced SNMP protocol can be used in sophisticated routing algorithms to avoid congestion at the core and balance the load of the connections.

OPPORTUNITY FOR AOLS TECHNOLOGY IN AN OPERATOR NETWORK

From a technological point of view, AOLS is very promising, because it enables the routing of packets independently of bit rate and packet length. Moreover, it helps reduce the mismatch between the forwarding speed of the router and the transmission speed of the fiber. However, to bring AOLS to market, network operators must be convinced of its advantages and the opportunities it creates to better serve their clients. The present section addresses the latency requirements that different applications impose. These requirements are very important from a network operator's perspective because if fulfilled, they enable the operator to deliver (new) guaranteed services to different types of user groups.

SERVICES DESCRIPTION

The interest in new services that are delivered by a telecommunication network is still growing worldwide. We focus on three classes of user services:
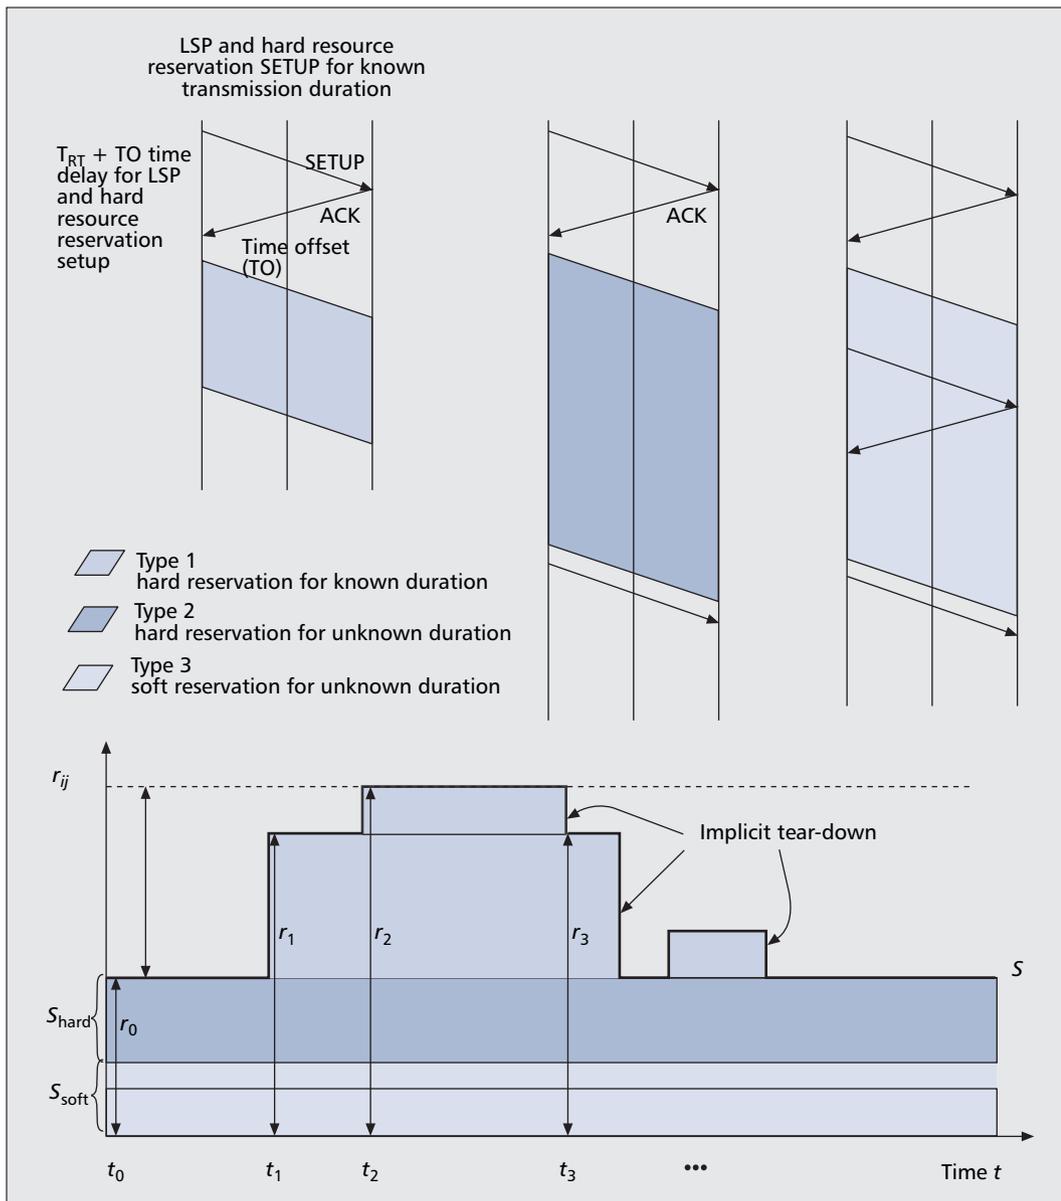
■ **Figure 6.** *a) Resource reservation mechanism in combination with the LSP setup process; b) utilization profile* $r_{ij}(t)$ *of wavelength* i *of link* j *at each time instance.*

- **Multimedia applications**: This class encompasses all services related to voice, video, and data transfer, both for residential and business end-users.
- **Storage applications**: Broadband technologies in the metro and core area create the capability of transferring massive amounts of data as requested by storage area networks (SAN), where huge data centers are interconnected for data mirroring, data back up, and disaster recovery [9].
- **Grid applications**: Resources (computational power, storage devices, and memory) are distributed over a wide area. A growing interest in these applications leads to high availability and bandwidth requirements.

Moreover, some of the currently existing services are supposed to survive in the long term, migrating to other technologies, and/or evolving.

REQUIREMENTS

For the different applications listed previously, the following network requirements may be established [10]:
- Real time requirements: delay/jitter (i.e., latency/latency variation) and maximum distance
- Data integrity requirements: bit error rate (BER), packet loss rate (PLR)
- Availability requirements: protection, redundancy

Note that because of the multiservice nature of an OPS network, the requirements that are usually of overriding importance are the following.

Latency requirements: This includes the end-to-end delay experienced by a data unit in one communication direction and also its variability (jitter). Typically, the most stringent requirements are counted in storage applications that

A techno-economic study showed that all-optical packet-switching nodes are still quite expensive compared to other packet-switching approaches, but on the other hand, we showed that operators may see opportunities to introduce all-optical packet-switching in their networks.

involve real-time, acknowledged data transfer. We refer to them as hardly real time (HRT):
• Delay: THRT = 3–10 ms
• Jitter: DTHRT = 1 ms

In FC, the maximum distance depends on the data rate and is limited to a few hundred meters. To overcome this and to extend the SAN distances, the framing procedure at the edge of the network must be adapted accordingly.

L2/3 data integrity requirements: Packet loss applies to those applications that employ L3 or L2 network services. The most demanding values come from grid applications, which require a PLR of at most, 0.01 percent.

L1 data integrity requirements: For L1 network services, the data integrity parameter that usually is considered is the BER. A typical value is 10-12.

Availability requirements: The availability of a service is defined as the percentage of time the service is available to the user. Four/five nines (0.9999/0.99999) typically are used in telecommunications.

Maximum recovery time: This is the duration of the recovery procedure when a protection mechanism is used. High availability requirements lead to a requirement to design both the network and the node with strong resilience solutions.

LATENCY REQUIREMENTS AND OPS NETWORKS

The total delay experienced by client traffic when transported through an OPS network is given by the sum of individual delays induced by the following different processes:
• **Packetization delay (T_{pack}):** is the time spent at the edge of the network to build an OPS packet from client PDUs. It comprises a serialization delay and an assembly delay.
• **Latency of OPS core nodes:** In the case of contention-free routing, this delay is only a propagation delay. In the case of contention, the delay depends on the contention resolution strategy. The maximum delay a packet experiences is then given by the length of the longest optical delay line.
• **Depacketization delay:** is the time required to extract client PDUs from the OPS payload. Because of the high bit rate of the optical packet, depacketization is realized electronically. This implies a fixed delay.

The previously listed delays are of the order of tens of microseconds and are actually negligible compared to the propagation delays (1 ms every 200 km) of a core wide area network (WAN). The delay is thus comparable to that of OCS. Delay variation between client PDUs occurs because of two reasons: assembly time variation and difference in the number of optical delay lines on the packet path. Concerning latency requirements, AOPS is compatible with both real time and hardly real time applications. In fact, the most demanding latency requirements (max delay 10 ms, max jitter 1 ms) are two orders of magnitude higher than the values expected for AOPS networks.

SUMMARY

This article described the design of a high-performance network architecture based on the LASAGNE all-optical node architecture. Our studies of dimensioning and economics showed that the proposed node architecture is very resource consuming, and thus we also examined an alternative label switching approach, the label-stripping strategy. For the LASAGNE network, we also described a suitable control protocol to distribute the labels in accordance with the GMPLS framework. The LASAGNE LSP set-up protocol also includes bandwidth reservation features to provide a mechanism for flow control that handles different types of traffic expected to be carried by an AOLS network. A techno-economic study showed that all-optical packet-switching nodes are still quite expensive compared to other packet-switching approaches, but on the other hand, we showed that operators may see opportunities to introduce all-optical packet-switching in their networks.

REFERENCES

[1] E. Mannie, ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," RFC 3945, Oct. 2004.
[2] D. J. Blumenthal, "All-Optical Label Swapping for the Future Internet," *Optics and Photonics News*, Special Issue on Fiber Optics, Mar. 2002, pp. 22–25.
[3] F. Ramos *et al.*, "IST-LASAGNE: Towards All-Optical Label Swapping Employing Optical Logic Gates and Optical Flip-flops," *IEEE/OSA J. Lightwave Tech.*, Nov. 2005.
[4] S. De Maesschalck *et al.*, "Pan-European Optical Transport Networks: An Availability-Based Comparison," *Photonic Network Commun.*, vol. 5, no. 3, May 2003, pp. 203–25.
[5] G. Bernstein *et al.*, *Optical Network Control: Architecture, Protocols and Standards*, Addison-Wesley, 2003.
[6] M. Enachescu *et al.*, "Part III: Routers with Very Small Buffers," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 35, no. 2, Jul. 2005.
[7] K. Christodoulopoulos *et al.*, "Relaxing Delayed Reservations: An approach for QoS Provision and Service Differentiation in Optical Burst Switching Networks," *BROADNET*, San Jose, CA, 2006.
[8] K. Manousakis *et al.*, "A Bandwidth Monitoring Mechanism: Enhancing SNMP to Record Timed Resource Reservations," *J. Network and Sys. Mgmt.*, vol. 14, no. 4, Dec. 2006, pp. 583–97.
[9] R.Girardi *et al.*, "Optical Networking for Storage Services," *ECOC 2003*.
[10] R. Van Caenegem *et al.*, "Evolution Towards an All-Optically Switched Packet Network: The LASAGNE Viewpoint," *Proc. 11th Euro. Conf. Networks and Optical Commun.*, Berlin, Germany, July 2006.

BIOGRAPHIES

RUTH VAN CAENEGEM (Ruth.VanCaenegem@intec.ugent.be) received an M.Sc. degree in electrical engineering from Ghent University, Belgium, in 2003. Since August 2003 she has been working toward her Ph.D. on all-optical networks in the Department of Information Technology (INTEC) at the same university. She was involved in the European project IST-LASAGNE. Her research interests include optical network architectures and all-optical packet/burst switching node architectures.

DIDIER COLLE received M.Sc. and Ph.D. degrees in electrotechnical engineering from Ghent University in 1997 and 2002, respectively. In 2003 he was granted a postdoctoral scholarship. His research interest is in the design and planning of communication networks.

MARIO PICKAVET received M.Sc. and Ph.D. degrees in electrical engineering from Ghent University in 1996 and 1999, respectively. Since 2000 he has been a professor at Ghent University, where he teaches telecommunications networks and algorithm design.

PIET DEMEESTER received a Ph.D. degree from Ghent University (INTEC) in 1988. He became a professor at Ghent University in 1993, where he is responsible for research on communication networks. He has been involved in several European projects.

KONSTANTINOS CHRISTODOULOPOULOS received a Diploma in electrical and computer engineering from the National Technical University of Athens (NTUA), Greece, in 2002 and an M.Sc. degree in advanced computing from Imperial College London, United Kingdom, in 2004. He is currently working toward a Ph.D. degree in the Computer Engineering and Informatics Department of the University of Patras, Greece.

KYRIAKOS VLACHOS received his Ph.D. in electrical and computer engineering from the National University of Athens in 2001. He is a faculty member with the Computer Engineering and Informatics Department of the University of Patras, Greece. From 1997 to 2001 he was a senior research associate in the Photonics Communications Research Laboratory, and in April 2001 he joined Bell Laboratories, Lucent Technologies. In 2003 he joined the National Regulation Authority of Telecommunication and Postal Service of Greece (EETT), where he served as a scientific advisor.

LEONTIOS STAMPOULIDIS received a Diploma in electrical and computer engineering from the University of Patras in 2002. Since 2002 he has been with the Photonics Communications Research Laboratory, NTUA. His research interests are in optical burst switching architectures and all-optical network subsystems, including the design and implementation of all-optical buffering and contention resolution systems.

EMMANOUEL (MANOS) VARVARIGOS received his Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology in 1992. In 2000 he became a professor with the Department of Computer Engineering and Informatics at the University of Patras, where he heads the Communication Networks Laboratory. He has held faculty positions at the University of California, Santa Barbara (1992–1998) and Delft University of Technology, The Netherlands (1998–2000).

DIEGO ROCCATO graduated in physics from the University of Pisa in 1984. In 1986 he joined CSELT (now Telecom Italia Lab) to work on optical transmission, in particular, soliton propagation and nonlinear optics. He was later involved in the study of no-dig techniques and the development of a ground penetrating radar to detect buried obstacles before deploying new underground cables. In recent years, he has been involved in the development of software to dimension the Telecom Italia optical network. He is an author of technical articles and holds patents in the previously mentioned fields.

RUTH VILAR received a telecommunication degree from the Universidad Politecnica de Valencia, Spain, in 2004. Currently, she works at FIBERNET on disaster recovery solutions. She has been involved in European projects such as IST-LASAGNE. Her research interests include optical network architecture, MPLS/GMPLS protocols, all-optical label swapping, optical packet switching, optical networking and architectures for disaster recovery solutions, and investigations of optical performance monitoring.