Sub-Wavelength Solutions for Next-Generation Optical Networks

Nizar Bouabdallah, INRIA

ABSTRACT

The recent trend in optical networks is switching packets directly in the optical domain. The aim is to benefit from both packet flexibility and optical transparency. In this article, we review current optical architectures that try to reconcile these two requirements. We discuss the challenges encountered in these new architectures and their respective niches. To meet the requirements of next-generation high-speed optical networks, we also propose a new solution based on the distribution of the aggregation process in the network. The feasibility of this scheme and the benefit that it provides over existing solutions are analyzed in this article.

INTRODUCTION

Over the last decade, networks have been witnessing perpetual growth in data traffic. This growth, driven primarily by the proliferation of the Internet, has created a rising demand for robust networks, with increasingly high-link capacity and node throughput. Due to new and the incumbent challenges, operators are progressively migrating toward all-optical networks, taking advantage of the tremendous transmission capacity offered by optical technology. In these networks, all-optical cross connects (OXC) are used to switch an incoming optical signal on a wavelength channel from an input port to an output port; thus a connection (lightpath) is routed from its source to its destination in the optical domain, optically bypassing all the intermediate nodes. Hence, all-optical networks are the opposite of opaque networks, where lightwave channels are detected at each node, then electronically processed, switched, and reassigned to a new outgoing wavelength when needed.

Realizing connections in an all-optical wavelength routed network involves the establishment of point-to-point (P-to-P) lightpaths between every edge node pair. In view of this, the all-optical wavelength routing approach, also called the optical circuit-switched (OCS) approach, presents two obvious advantages. The first advantage stems from the fact that the optical bypass eliminates the requirement for expansive optical-electrical-optical (OEO) conversion at intermediate nodes. The second advantage is due to all-optical routing, which is transparent with regard to the bit rate and the format of the optical signal.

In spite of the aforementioned advantages, all-optical wavelength routing still presents two drawbacks. The first one is related to the large number of wavelengths required within a large network when routing is performed at a wavelength granularity. In fact, when full connectivity is required, the number of wavelengths required is expected to have scalability issues. For instance, if P-to-P lightpaths must be established between every edge node pair in a network presenting N edge nodes, then $O(N^2)$ lightpaths are required. The second drawback of wavelength routing is the rigid routing granularity entailed by such an approach. This granularity is large, which could lead to severe bandwidth waste, especially when only a portion of wavelength capacity is used. For operators, an efficient use of network resources is always a concern. In wavelength routed networks, this efficiency is possible only when there is enough traffic between pair nodes to fill the entire capacity of wavelengths.

On the other hand, an opaque network has the advantage of being able to use the link bandwidth efficiently. Nonetheless, as nodes do not have optical bypass, this results in a maximum transceiver cost. The major advantage of such an electronic packet-switching scheme is its bandwidth efficiency due to statistical multiplexing. Therefore, currently much research is focusing on bringing the packet switching concept into the optical domain. The aim is to benefit from both optical transparency and sub-wavelength multiplexing gain. However, optical packet switching (OPS) is not yet ready and is hampered by major technology limitations due to the issues related to the fine switching granularity adopted (optical packet) at high bit rate [1]. In this regard, OPS seems to be a long term solution that is still many years away. Meanwhile, several solutions were proposed in the literature as alternatives to the immature OPS technology [2–8]. The challenges to these solutions, their efficiencies, and their respective niches are discussed in this article.

To alleviate the aforementioned problems, we also propose a new solution, which combines the advantage of the optical bypass in transparent wavelength routed networks with statistical

0163-6804/07/$20.00 © 2007 IEEE

multiplexing gain. In this technique, a lightpath that remains entirely in the optical domain is shared by the source node and all the intermediate nodes up to the destination. So, in essence, a single lightpath is used to establish a multipoint-to-point (MP-to-P) connection. We refer to this technique as the distributed aggregation (DA) scheme [9].

The rest of the article is organized as follows. The following section discusses in more detail prior research related to this work. Then, a detailed description of our proposed DA approach is outlined. We first investigate the node architecture required to support such traffic-aggregation features within wavelength-division multiplexing (WDM) optical networks. Moreover, we emphasize the medium access control (MAC) context, including a description of the associated fairness control mechanism. Then we assess the benefits introduced by our proposal with respect to existing solutions. To achieve this, all underlying network blocking probabilities are compared. Finally, some conclusions are drawn.

CURRENT TECHNOLOGIES

As stated before, widely deployed all-optical wavelength routed networks currently are no longer consistent with the packet-switching philosophy of the Internet. Specifically, in next-generation networks, packet-based data traffic of bursty nature will become prevalent. Hence, the lack of grooming may lead to critical underutilization of resources. Consequently, the architecture of next-generation networks must evolve to enable tackling these new challenges. In this regard, two major enabling factors are identified as crucial for the evolutionary process of next-generation network architecture: packet switching and optical transparency.

Indeed, packet switching constitutes the ingredient required for building bandwidth-efficient and flexible networks [1]. Asynchronous transmission, which is more suitable for bursty traffic in comparison with slotted WDM networks, must be addressed jointly in the future. Moreover, handling a finer granularity than a full wavelength, must be achieved while trying to preserve the optical transparency property to avoid extra and expensive electronic conversions.

To cope with these requirements, many interesting solutions have been proposed in the literature [2–8]. These solutions fall into two categories: OPS-based or OCS-based solutions. In the following, we describe these new technologies, pointing out how they reconcile the optical transparency and sub-wavelength grooming concepts.

OPTICAL PACKET SWITCHING-BASED SOLUTIONS

Much research is now focused on bringing the packet-switching concept into the optical domain. The motivation is to take advantage of both packet flexibility and optical transparency. However, OPS is hampered by major technology bottlenecks, such as the lack of optical processing logic, optical memory, cost-effective fast switching, synchronization technology, and so on [1]. To meet these technology bottlenecks, we identified two promising solutions: photonic slot routing (PSR) [2] and optical burst switching (OBS) [3].

Photonic Slot Routing — One approach toward realizing optical packet switching is PSR [2]. In PSR, time is slotted, and data is transmitted in the form of photonic slots that are fixed in length and span across all wavelengths of a given fiber link. Each wavelength in the photonic slot contains a single packet, and all the packets in the photonic slot are destined to the same node. By requiring the packets to have the same destination, the photonic slot may be routed as a single integrated unit without the need for demultiplexing wavelengths at intermediate nodes. The basic aim of this approach is the use of wavelength-insensitive components at each node, resulting in less complexity, faster routing, and lower cost compared to the classical OPS concept.

However, the PSR approach still requires high-speed configurable optical packet switches. Indeed, the data processing and switching is done on a slot basis. Moreover, implementing PSR in a mesh environment is an even more challenging key issue, as it involves maintaining the synchronization of photonic slots at each node level.

Optical Burst Switching — Another way to relax the very short switching time required by OPS switches is OBS [3]. A burst is an aggregation of packets with the same requirements (same egress node, same class of service, etc). By manipulating such relatively large entities, the processing on the network can be reduced compared to classical OPS. Therefore, a burst can be considered as a hybrid approach between coarse-grained circuit switching and fine-grained packet switching. The bursts are aggregated by edge nodes and transmitted all-optically towards their destination. Each burst is preceded in time (offset time) by its header to reserve resources at each switch in the core network. The offset time must be large enough to configure the switches prior to the arrival of the burst.

In this regard, the edge routers require complex interfaces to implement burst assembly, disassembly, and queue fairness algorithms. Thus, the access unit design may become challenging at high data rates. Furthermore, under this basic concept, the number of dropped bursts by intermediate nodes, due to contentions among bursts going to the same outgoing port may be significant. The effects of dropping bursts can be detrimental to the network as each burst is an amalgamation of packets.

OPTICAL CIRCUIT SWITCHED-BASED SOLUTIONS

OPS, in particular PSR, appears as a long term solution. This technology could become practical in the future. Meanwhile, the trend is to improve the efficiency of existing mature all-optical networks. In this regard, recently, there has been emphasis on circuit-switched all-optical networks, where the goal in this context is shifted toward the improvement of the optical resources

A burst is an aggregation of packets with the same requirements (same egress node, same class of service, etc). By manipulating such relatively large entities, the processing on the network can be reduced compared to classical OPS.

usage of such networks by means of new traffic aggregation schemes, rather than toward the attempt to realize optical packet switching.

TWIN may lead to fiber link underutilization due to the lack of wavelength reuse, because a particular wavelength, wherever the link that belongs to, can only be used to transmit to a specific egress node.

The Multihop Approach — The key idea behind multihop (MH) networks is to enable electronic processing at some of the intermediate nodes of the all-optical circuit-switched network to increase its grooming capacity [4]. Accordingly, a packet may undergo electronic processing at some of the intermediate nodes before reaching its final destination. Hence, lightpaths can be seen as chains of physical channels through which packets are moved from one router to another toward their destinations. At intermediate nodes, the transit lightpaths are switched transparently through an OXC that does not process transit data. Instead, incoming lightpaths destined to the current node are terminated and converted to the electronic domain, so that packets can be extracted, processed, and possibly retransmitted on outgoing lightpaths if the current node is not the final destination of the data.

Although significant cost may be introduced by this electronic processing operation at intermediate nodes, it enables a better use of the network resources and can reduce the total network cost [4]. In this case, the main challenge is to identify the optimal logical topology that minimizes the total network cost while accommodating all the traffic requests. It has been demonstrated that the identification of the optimal logical topology is computationally intractable for large-size networks [10]. In view of this, several heuristic approaches were proposed in the literature [4].

The Super-Lightpath Approach — To increase the grooming capacity of classical all-optical circuit-switched networks, the super-lightpath concept [5] transforms the lightpath from a P-to-P pipe to a P-to-MP pipe. In other words, the source node of a super-lightpath does not limit its transmission to the end node of that lightpath; instead, it can transmit its traffic to all the intermediate nodes along the route. This enables the super-lightpath to carry multiple connections, resulting in better wavelength utilization.

To split the wavelength bandwidth among more than one traffic flow, the super-lightpath technique uses a simple optical time-division multiplexing (OTDM) method. Accordingly, each bit in a given position of the fixed-size TDM frame, called a bit slot, identifies a particular subchannel. Using a bit interleaver, the transmitter multiplexes sub-channels into the frame and transmits the resulting stream into one lightpath. With regard to reception, each intermediate node splits the transit signal, synchronizes its receiver to a particular bit slot, and only receives data in that particular sub-channel.

The super-lightpath technique presents many advantages. First, it reduces the number of transmitters per node because each transmitter is used to send data to more than one receiver, due to the OTDM approach. Moreover, it improves the lightpath utilization. But the main concern with this method is related to the limited length of the distributing super-lightpath (in term of traversed nodes) due to power limitation, because a significant portion of the passing-through optical signal is tapped at each receiving intermediate node.

The TWIN (Time-Domain Wavelength Interleaved Networking) Approach — Unlike the super-lightpath concept that uses a P-to-MP approach to improve the traffic grooming capacity of traditional OCS networks, the TWIN technique adopts a MP-to-P approach [6]. Indeed, TWIN makes use of optical MP-to-P trees that are overlaid on top of the physical topology. In TWIN, a particular wavelength is attributed to each egress node to receive its data. In doing so, sources that have information to transmit to a particular destination, tune their transmitters to the associated wavelength. As such, the optical signals from various sources to a particular destination may be merged at the intermediate nodes. Thus, TWIN technology requires special OXC that are able to merge incoming signals of the same wavelength to the same outgoing wavelength.

Despite the complex scheduling algorithms entailed by such an approach, the MP-to-P concept itself is interesting. It avoids the above-mentioned limitations on the length of lightpaths introduced by the P-to-MP approach (i.e., super-lightpath), since no splitting operations are performed in this case.

Nevertheless, the MP-to-P concept as described in TWIN suffers from scalability issues. Indeed, the assignment of multiple wavelengths to each egress node (according to the volume of its destined traffic) puts serious stress on the number of wavelength channels required on each fiber link. Moreover, TWIN may lead to fiber link underutilization due to the lack of wavelength reuse, because a particular wavelength, wherever the link to which it belongs, can only be used to transmit to a specific egress node.

The Optical Light-Trails Approach — The light-trail (LT) is a OCS-based technology that aims to minimize active switching, maximize wavelength utilization, and offer protocol and bit-rate transparency [7, 8]. So far, to achieve this target, we have presented a P-to-P approach (i.e., MH), a P-to-MP approach (i.e., super-lightpath), and a MP-to-P approach (i.e., TWIN). Instead, the LT solution is a MP-to-MP approach. A LT is therefore a lightpath where intermediate nodes can both receive and transmit data on the pass-through channel.

The basic operation of this scheme follows. Each intermediate node *i* of the LT taps a sufficient amount of optical power from the incoming signal, using a splitter, to recover its corresponding packets sent by the upstream nodes. On the other side, with regard to transmission, the original transit signal is coupled with the local signal, by means of a coupler, before it continues its path to serve the remaining downstream nodes of the LT.

The main concern with this method is the design of a MAC protocol that avoids collisions between transit and local inserted packets. A simple MAC protocol based on in-band signaling was suggested in the original LT proposal [7]. Accordingly, each intermediate node *i* wishing to transmit a packet, first sends a beacon signal to order downstream nodes to stop their activities on the shared medium. Then, following a guard band, it transmits its data packet. Note that, node *i* may receive a beacon signal from upstream nodes during its transmission of a beacon signal or a data packet. In this case, it preempts instantaneously its transmission and the truncated packet is lost.

The previous concerns may have a negative impact on the performance of the LT approach. Indeed, the MAC scheme may result in low resource utilization due to the guard band, extra signaling packets, and wasted truncated packets. In this regard, other research is focusing now on the elaboration of more efficient MAC schemes adapted to the LT technology [11]. Furthermore, the fact that a significant portion of the signal is tapped at each intermediate node causes a severe constraint on the LT length. Finally, we note that packets received by an intermediate node are not removed from the LT, which prevents the bandwidth reutilization by downstream nodes. This feature becomes interesting only when dealing with multicast applications.

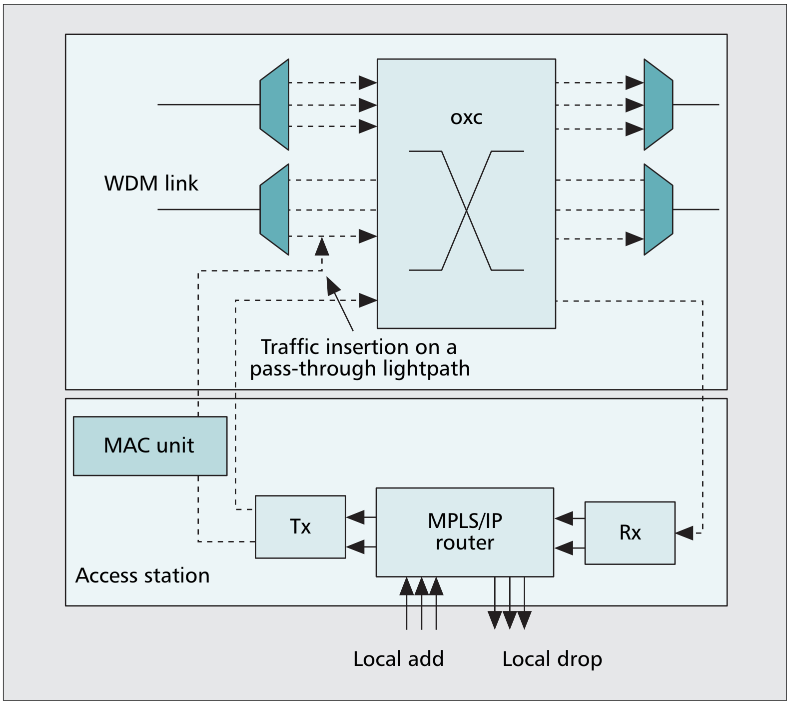

Figure 1. *Node architecture.*

THE DISTRIBUTED AGGREGATION APPROACH

As explained before, methods based on multiple node reception, such as super-lightpath and LT, suffer from power limitations due to splitting requirements. Moreover, the multiple-nodes reception feature in an LT approach is effective only when dealing with multicast applications due to the lack of bandwidth reutilization of the shared lightpath. In view of this, the MP-to-P strategy seems to be the best choice to improve the grooming capacity of the classical all-optical lightpaths. In this context, TWIN technology is a good candidate. However, this technique suffers from inherent scalability and lack of wavelength reuse issues. To relieve this, we propose a new MP-to-P OCS-based solution, called the distributed aggregation (DA) scheme [9].

The key idea underlying our proposed scheme is to allow sharing of a lightpath among several access nodes. Instead of limiting the access to the lightpath capacity at the ingress point, each node along the path can fill the lightpath on the fly according to its availability. In this case, a lightpath can be shared by multiple connections traveling toward a common destination (i.e., MP-to-P lightpaths). Wavelength routing is performed in a similar way as in all-optical networks, that is, signals remain in the optical domain from end to end and are optically switched by intermediate nodes. Since the lightpath remains transparent at intermediate nodes, a MAC protocol is required to avoid collisions between transient optical packets and local ones injected into the lightpath [12]. Moreover, additional control mechanisms must be introduced to alleviate fairness problems, which are pronounced in the case of shared medium networks [13]. In the following, we provide a detailed description of the proposed control mechanisms, as well as the node architecture required to support the DA feature.

THE NODE ARCHITECTURE

A typical node in a WDM network is shown in Fig. 1. It consists of an OXC part and an access station part. While the OXC performs wavelength routing and wavelength multiplexing/demultiplexing, the access station performs local traffic adding/dropping functionality. Each OXC is connected to the access station, typically a multiprotocol label switching/Internet Protocol (MPLS/IP) router, which can be the source or the destination of a traffic flow. Each access station is equipped with a certain number of transmitters and receivers (transceivers). Traffic that originated at the access station is transmitted as an optical signal on one wavelength channel by virtue of a transmitter. Considering the DA, the access station can be either the origin of a lightpath or an intermediate node using an already established lightpath. In the latter case, the injected traffic by an intermediate node should have the same destination as that of the traversing lightpath. In this context, a MAC unit is required to avoid collisions between transient packets and local ones. On the other side, the traffic destined to the access station is directed by the OXC to the access station, where it is converted from an optical signal to electronic data by means of a receiver.

Aggregating low-speed connections to high-capacity lightpaths is done by the MPLS/IP router according to the MAC unit decision. Usually, the potential disadvantage of such a model is that the processing speed of the MPLS/IP router may not be fast enough compared to the vast amount of bandwidth provided by the opti-

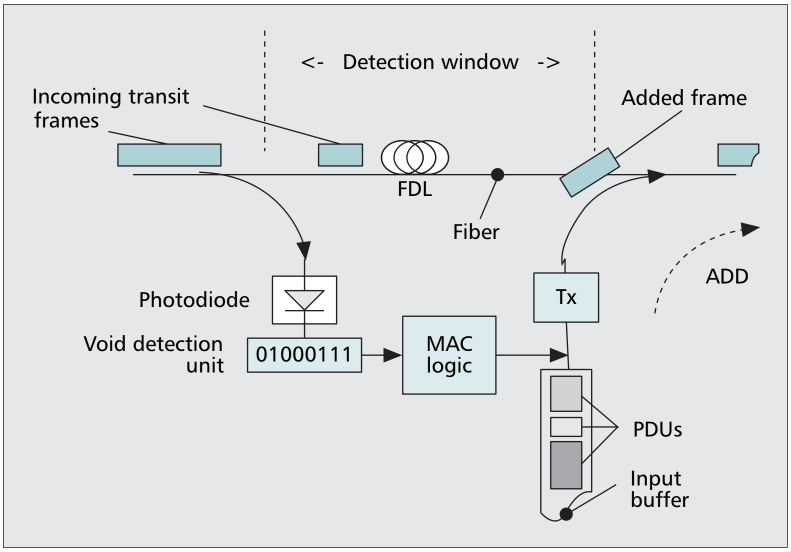

Figure 2. *Schema of the void-detection-based MAC.*

cal fiber link. However, our scheme alleviates this issue because each MPLS/IP router processes only its local traffic. In other words, the transit traffic traveling through a WDM node remains at the optical layer, and it is not processed by the intermediate access nodes.

The advantage of DA is that multiple connections with fractional demands can be multiplexed onto the same lightpath. As a result, the wasted bandwidth problem associated with pure wavelength-routed networks is alleviated. In addition, due to the sharing of lightpaths, the number of admissible connections in the network is increased. Furthermore, the destination node terminates fewer lightpaths as connections from different nodes to the same destination are aggregated onto the same lightpath. In view of this, fewer physical components, such as wavelengths and transceivers, are used, resulting in savings on equipment. Moreover, to provide connections between all access node pairs using MP-to-P lightpaths, a total number of $O(N)$ lightpaths is required as only one lightpath per individual egress node could be sufficient. Thus, we alleviate the scalability issue encountered in classical all-optical wavelength-routed (i.e., OCS) networks.

THE MAC PROTOCOL

To ensure collision-free insertion of intermediate node packets on the transit MP-to-P lightpath, we proposed a new MAC protocol based on the void detection principle [12]. The MAC protocol detects a gap between two transient packets on the optical channel, and then it tries to insert a local packet into the perceived gap. To do so, each access station must retain the transit traffic flow within the optical layer while monitoring the medium activity. So, as shown in Fig. 2, each node first uses an optical splitter to separate the incoming signal into two parts: the main transit signal and its copy used for control purposes. With regard to the control part, as in [14], low bit-rate photodiodes (ph) — typically 155 MHz — are used to monitor the activity of the transit wavelengths. After a free state of the medium is detected, the MAC unit measures the size of the progressing void. It is worthwhile to note that signal splitting is done to monitor the medium activity (i.e., to know whether the medium is idle or busy) rather than to recognize the transit stream as with super-lightpath and LT schemes. This requires simply tapping a small part of the transit signal. Hence, the power penalty is relatively negligible. In [15], it is demonstrated that one could cascade up to 10 nodes without a significant power penalty.

To use a detected void, a fiber delay line (FDL) is introduced on the transit path to delay the upstream flow by one maximum frame duration augmented by the MAC processing time. Therefore, the length of the FDL is slightly larger than the maximum transmission unit (MTU) size allowed within the network to provide the MAC unit with sufficient time to listen and to measure the medium occupancy. The access station begins injecting a packet to fill the void only if the null period is large enough (i.e., at least equal to the size of the packet to be inserted). Undelivered data remains buffered in the electronic memory of the access station until a sufficient void space is detected. This way, collision-free packet insertion on the transit lightpath is ensured. It is easy to see that such an access scheme relies only on passive components (couplers, FDL, ph) with relatively low cost. Therefore, the cost introduced by the MAC unit is therefore negligible compared to the transceiver cost.

RESOLVING FAIRNESS AND HEAD OF LINE BLOCKING ISSUES

Because the DA scheme relies on the lightpath sharing, efficient partition of the lightpath capacity among competing access nodes must be ensured; otherwise, head of line (HoL) blocking and fairness issues could arise with such a scheme.

In fact, the mismatch between the idle period distribution resulting from the upstream nodes' utilization of the medium and the packets' size distribution of the downstream nodes often leads to bandwidth waste, as well as fairness problems with regard to resource access. As soon as a packet of maximum size is at the head of the insertion buffer of an intermediate node, it blocks the node emission process until finding an adequate void: this is the well-known HoL blocking problem. Monitoring the distribution of voids on the medium reveals a low probability of finding regular and sufficiently large available room. The sharing process must be done intelligently to preserve a maximum of useful available bandwidth for downstream nodes. In this context, we demonstrated in [13] the limitations of the Token Bucket (TB) algorithm to resolve this issue. Certainly, due to the TB algorithm, the free bandwidth (stated in bit/s) allocated to each node is theoretically sufficient to handle its traffic. However, the main issue pertains to the inappropriate distribution of the free bandwidth. Hence, a basic rule is to avoid the random division of the optical resource. To achieve this aim, we proposed the traffic control architecture using the remote descriptors (TCARD) mechanism [13].

In TCARD, each transmitting station is equipped with anti-tokens that prevent the station from transmitting a packet during a gap in the optical packet stream. These anti-tokens permit some of the gaps to be unused, and therefore, they can be used by other downstream stations. The rate of generation of the anti-tokens at a station is set equal to the rate of the aggregate downstream transmission. Moreover, to avoid the HoL blocking problems, the reserved gaps must be big enough to enable downstream nodes the transmission of large packets. Hence, the key idea of TCARD is to force each node to preserve free bandwidth for its downstream neighbors in the form of rooms of size equal to the MTU size.

To illustrate the TCARD mechanism, we present a simple three-node MP-to-P lightpath example. The nodes share a common channel that runs at 1 Gb/s. We assume that the sustainable bit rate negotiated by each node and stipulated in its own service level specification is 0.3 Gb/s. We consider traffic of variable packet size where the MTU is equal to 1500 bytes. Considering the TCARD scheme, the first node must reserve 0.6 Gb/s on average of available bandwidth for the downstream nodes, that is, nodes 2 and 3. As explained before, the reserved bandwidth is representative of idle periods of 1500 bytes to comply with packets of maximum size. Thus the anti-tokens at node 1 are generated periodically at a rate equal to $(0.6\ 10^9)/(1500 \times 8)$ anti-tokens/s. Note that a reserved void can be exploited by a downstream node either to transmit a packet of maximum size or to emit a burst of smaller frames. Furthermore, similarly to the first node, the second node reserves 0.3 Gb/s of available bandwidth for the third node. The reserved bandwidth is also representative of voids of 1500 bytes.

PERFORMANCE EVALUATION

In the previous section, we described the MAC and fairness control mechanisms required by the DA scheme. The performances of such mechanisms are studied in depth in [13]. Rather, in this section, we evaluate the gain introduced by the DA scheme over classical approaches (all-optical networks, opaque networks, and MH networks). To accomplish this goal, we conduct a comparison study. Specifically, we assess the blocking probability of dynamically arriving connection requests when using the different approaches. To achieve this, we develop a new discrete-event simulation tool. We simulate the following schemes:

- All-optical (i.e., OCS) networks
- Opaque networks
- MH networks
- DA-enabled networks (i.e., MP-to-P approach)
- A hybrid variant of networks, a combination of the MH and MP-to-P approaches

The following assumptions are made in our simulations:

- The US backbone, shown in Fig. 3 is used. The network topology consists of 29 nodes and 43 links.
- Each link in the network represents two fibers that are used to communicate in opposite directions. Each fiber carries 32 wavelengths.
- Each node is equipped with 20 transceivers and 40 OXC interfaces.
- The shortest path adaptive routing is used.
- The first fit (FF) wavelength assignment approach is adopted.
- Connection requests arrive at each ingress node following the Poisson distribution, and their holding time is exponentially distributed. The total traffic load offered to the network by each node is $\rho = \lambda/\mu$, where λ and μ are the arrival and departure rates at each ingress node, respectively.
- The destination node of each arriving connection is randomly chosen among the N–1 remaining edge nodes of the network.
- The bandwidth requirement of each connection request λ_{sd} is randomly chosen in the interval [0,1], so at most one lightpath is required to carry any traffic request. We note that, in our simulations, we do not allow connection request traffic to be bifurcated among multiple lightpaths. Finally, each value of the blocking probability was calculated over multiple simulations to achieve very narrow 97.5 percent confidence intervals.

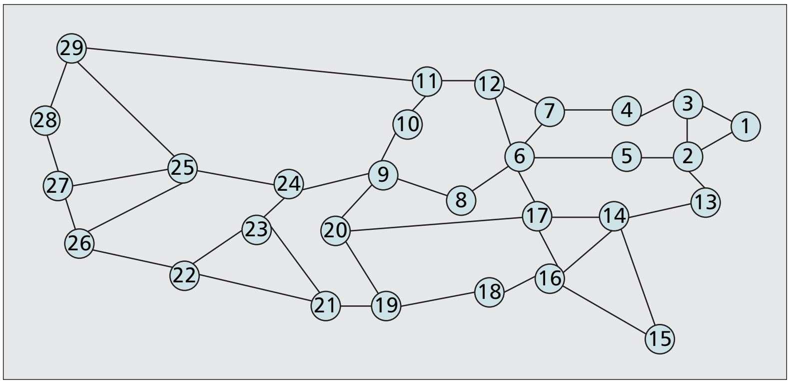

Figure 3. *The U.S. optical backbone.*

In our simulations, each arriving connection first tries to use the current virtual topology (i.e., already established lightpaths). If the current available bandwidth on existing lightpaths is not sufficient, the connection tries to establish new lightpaths subject to transceiver, OXC port, and wavelength constraints. Specifically, when the DA case is considered, the ingress node s of an arriving connection request going to destination d looks first for a pass-through lightpath traveling toward the same egress node d with sufficient available bandwidth. Otherwise, the node s tries to establish a new lightpath subject to resource availability. If there are not enough resources to satisfy the connection request, it is simply blocked. In the same way, when the MH approach is adopted, the source node s first tries to find an available root through the existing lightpaths. In this case, the connection may span multiple lightpaths before reaching its destination. If this root is not available, the connection tries to establish the missing lightpaths (end-to-end lightpath or missing segments along the root) to reach its destination.

Figure 4 plots the different blocking probabilities, which logically increase with the network load ρ. Figure 4 shows that the opaque strategy always leads to the maximum blocking probability. Mainly, this is due to the lack of available transceivers. Indeed, the total network capacity (in terms of transceiver equipment) is quickly exhausted, because nodes do not have optical pass-through. The all-optical circuit-switched strategy slightly relieves this issue due the optical transparency property. Even so, the blocking probability remains relatively high due to the great number of P-to-P lightpaths required in this case. These lightpaths are highly consuming in terms of OXC interfaces and wavelengths. The MH and DA schemes reduce significantly the blocking probability, because they alleviate the scalability issues of classical all-optical networks by increasing their grooming capacities.

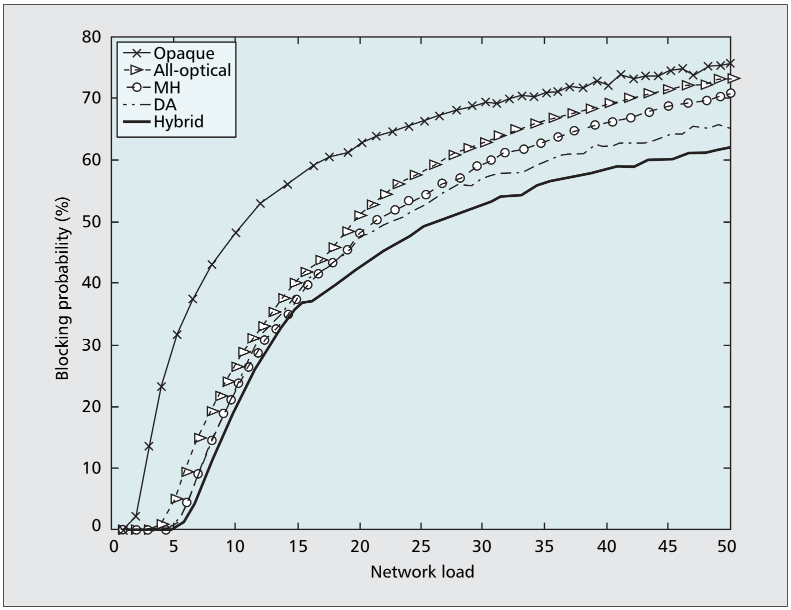

■ **Figure 4.** *Blocking probability evolution with the network load.*

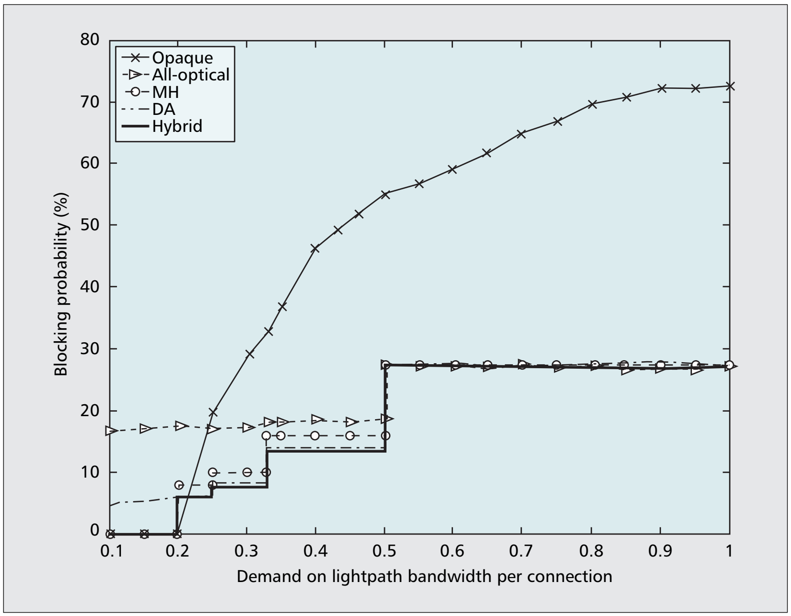

■ **Figure 5.** *Blocking probability evolution with the bandwidth requirement per connection.*

Moreover, the DA scheme outperforms the MH one, since it is less consuming in terms of transceivers and OXC interfaces. Indeed, the DA scheme improves the all-optical circuit-switched network grooming capacity while conserving its entire transparency as opposed to the MH approach, where electronic grooming is required at some network nodes to achieve the same target. This latter active operation enables an MH network to achieve component savings cost over classical all-optical and opaque networks, but requires additional equipment (OXC interfaces, transceivers) as compared to the passive DA insertion. Finally, we notice that the hybrid strategy, combining the MH and DA schemes, always achieves the best results.

Figure 5 plots the blocking probability as a function of the bandwidth requirement λ_{sd} of each connection request. In this case, we consider a uniform traffic matrix, that is, $\lambda_{sd} = \tau$ $\forall s$ and d, where τ ranges from 0 to 1. Moreover, the network load ρ is set equal to 10. This figure illustrates the general trade-off among different strategies. According to the value of τ, we obtain different optimal solutions. At one extreme, when each node transmits close to the wavelength capacity to each other node, the all-optical circuit-switched approach is the best solution, as the network is already well utilized without grooming. At the other extreme, when the total demand from each node is a small fraction of the wavelength capacity, the opaque strategy stands out as the best solution due to its grooming capability. In most cases, when the demand is moderate to normal, MH and DA schemes generally present the best solutions, with an advantage to the DA scheme. Finally, we underline that the hybrid solution enables achieving this trade-off, no matter what the value of τ is. It always leads to the minimal blocking probability. Therefore, this solution represents a sensible choice for next-generation networks. We note that MH, DA, hybrid, and classical all-optical strategies achieve almost the same results when $\tau > 1/2$. This is because we do not allow traffic belonging to the same connection request to be bifurcated among multiple lightpaths. In doing so, grooming multiple connections on the same lightpath is no longer possible when $\tau > 1/2$.

CONCLUSION

In this article, we reviewed the most prominent solutions that aim to achieve sub-wavelength grooming in all-optical networks. These architectures, which are alternatives to the immature optical packet-switching technology, try to reconcile two current opposite requirements: packet switching and optical transparency. To achieve this, we also proposed a new solution based on the distribution of the aggregation process in the optical network. In this approach, we tend to aggregate traffic traveling from different nodes to a common destination in the same lightpath channels. As a result, the number of transceivers and lightpaths are significantly reduced and the utilization percentage of the optical channels is improved as well.

To assess this gain, all underlying approaches were compared, using various simulation scenarios. Results showed that the distributed aggregation scheme always enables the lowest network blocking probability. In other words, the distributed aggregation scheme increases the total network throughput. This technique is proven to be extremely effective when the bandwidth requirements of connections between node pairs are fractions of the lightpath capacity. In view of this, the distributed aggregation scheme alleviates both the wasted bandwidth and the scalability issues encountered in the all-optical wavelength routed networks and preserves the benefits of optical bypass.

REFERENCES

[1] M. J. O'Mahony *et al.*, "The Application of Optical Packet Switching in Future Communication Networks," *IEEE Commun. Mag.*, Mar. 2001, pp. 128–35.

[2] H. Zhang, J. P. Jue, and B. Mukherjee, "Capacity Allocation and Contention Resolution in a Photonic Slot Routing All-Optical WDM Mesh Network," *IEEE/OSA J. Lightwave Tech.*, vol. 18, no. 12, Dec. 2000.

[3] Y. Chen, C. Qaio, and X. Yu, "Optical Burst Switching: a New Area in Optical Networking Research," *IEEE Network*, vol. 18, May 2004, pp. 16–23.

[4] K. Zhu and B. Mukherjee, "Traffic Grooming in an Optical WDM Mesh Network," *IEEE JSAC*, vol. 20, Jan. 2002, pp. 122–33.

[5] M. Mellia *et al.*, "Exploiting OTDM Technology in WDM Networks," *Proc. IEEE INFOCOM 2002*, New York, NY, June 2002, pp. 1822–83.

[6] I. Widjaja *et al.*, "Light Core and Intelligent Edge for a Flexible, Thin-Layered, and Cost-Effective Optical Transport Network," *IEEE Commun. Mag.*, May 2003, pp. S30–S36.

[7] A. Gumaste and I. Chlamtac, "Light-Trails: "A Novel Conceptual Framework for Conducting Optical Communications," *Wksp. High Perf. Switching and Routing*, June 2003, pp. 251–56.

[8] A. Gumaste, "Light-Trails and Light-Frame Architectures for Optical Networks," Ph.D. dissertation, Univ. TX Dallas, 2003; http://www.cotrion.com/light-trails

[9] N. Bouabdallah *et al.*, "Distributed Aggregation in All-Optical Wavelength Routed Networks," *Proc. IEEE ICC*, Paris, France, June 2004.

[10] B. Mukherjee, *Optical Communication Networks*, McGraw-Hill, 1997.

[11] S. Balasubramanian, A. Kamal, and A. K. Somani, "Medium Access Control Protocols for Light-Trail and Light-Bus Networks," *Proc. 8th IFIP Working Conf. Opt. Net. Design and Modeling*, Feb. 2004.

[12] N. Bouabdallah *et al.*, "Matching Fairness and Performance by Preventive Traffic Control in Optical Multiple Access Networks," *Proc. OptiComm '03*, Dallas, TX, Oct. 2003, pp. 424–29.

[13] N. Bouabdallah *et al.*, "Resolving the Fairness Issues in Bus-Based Optical Access Networks," *IEEE JSAC*, vol. 23, no. 8, Aug. 2005, pp. 1444–57.

[14] R. Gaudino *et al.*, "RINGO: A WDM Ring Optical Packet Network Demonstrator," *Proc. ECOC '01*, Amsterdam, The Netherlands, vol. 4, Sept. 2001, pp. 620–21.

[15] N. Le Sauze *et al.*, "A Novel, Low Cost Optical Packet Metropolitan Ring Architecture," *Proc. ECOC '01*, Amsterdam, The Netherlands, Sept. 2001.

BIOGRAPHY

NIZAR BOUABDALLAH (Nizar.bouabdallah@inria.fr) received his B.S. degree in telecommunications engineering from École Supérieur des Communications, Tunis, Tunisia, in 2001, and his M.S. and Ph.D. degrees in network and computer science from the University of Paris VI, France, in 2002 and 2004, respectively. He is currently a researcher at Institut National de Recherche en Informatique et en Automatique (INRIA). Since February 2007 he has been a visiting researcher at the School of Computer Science, University of Waterloo, Ontario, Canada. His research interests include optical networking, wireless and sensor networks, performance evaluation, network planning, and modeling, as well as control and management architectures. He joined Alcatel Research Laboratories, Marcoussis, France, in 2002 while working on his Ph.D. degree. In 2005 he was with North Carolina State University, Raleigh, as a postdoctoral fellow.

The distributed aggregation scheme alleviates both the wasted bandwidth and the scalability issues encountered in the all-optical wavelength routed networks and preserves the benefits of optical bypass.