

# Optical Burst Switching: A Viable Solution for Terabit IP Backbone

Sanjeev Verma, Hemant Chaskar, and Rayadurgam Ravikanth

## Abstract

Wavelength-division multiplexing has emerged as an important physical layer technology. Optical transmission provides a physical layer capable of carrying bits at the speed of the order of a gigabit per second. Optical burst switching is proposed to overcome the shortcomings of conventional WDM deployment, such as lack of fine bandwidth granularity in wavelength routing and electronic speed bottlenecks in SONET/SDH. In this article, we describe an architecture for IP network over the OBS WDM transmission core. The use of MPLS-type technique for forwarding data bursts and the inclusion of a medium access control layer between the optical (WDM) and IP layers are the key ingredients of the proposed architecture. In particular, the architecture is based on provisioning MPLS paths, also called label switched paths, of desired quality of service through the OBS WDM transmission core. The MAC layer performs various OBS-specific functions, such as burst assembly, burst scheduling, and offset setting/traffic shaping. While burst assembly and burst scheduling are relatively straightforward, we point out that the offset setting strategy has significant impact on the performance of IP network operating over OBS WDM core. We describe a shaping scheme to set the offset, an important system parameter for OBS, between the successive data bursts of a given data stream (label switched path) and their associated control packets. This scheme results in robust operation of the network and also facilitates traffic engineering. Guidelines are provided for implementing various IP QoS mechanisms in the optical backbone using OBS.

There has been explosive growth in Internet Protocol (IP) traffic in the last few years. It has triggered a lot of research activities in devising new high-speed transmission and switching technologies [1–5]. Wavelength-division multiplexing (WDM) [1] has emerged as a core transmission technology for the next-generation IP backbone network with its ability to support a number of high-speed (gigabit) channels in a single fiber. This provides enormous bandwidth at the physical layer. There is a need to develop framework and protocols at higher layers to efficiently use the raw bandwidth available at the optical (WDM) layer.

Presently, WDM is mainly deployed in the backbone of major long distance carriers as point-to-point links with a synchronous optical network (SONET) as a standard interface to higher layers in the protocol stack. This necessitates optical-to-electrical (O/E) and electrical-to-optical (E/O) conversions at every node, and hence fails to take advantage of the wavelength routing capability provided by WDM technology. Also, electronic multiplexing layers — IP, asynchronous transfer mode (ATM), frame relay — introduce further bandwidth inefficiencies. Although there has been a dramatic increase in the speed of electronic devices in the recent past, it is not likely to catch up with the transmission

speed available at the optical layer. This necessitates a renewed effort to minimize or eliminate electronic processing to fully exploit the bandwidth potential provided by WDM technology [6].

One possibility is to have an all-optical backbone using optical packet switching technology. However, this technology is yet to mature and needs to overcome a number of technological challenges. Here we discuss optical burst switching (OBS) as a viable transmission technology for the next-generation optical backbone, and provide a framework to deploy IP over WDM using OBS. Salient features of this framework are multiprotocol label switching (MPLS)-type forwarding and inclusion of a medium access control (MAC) layer between the optical and IP layers. A number of design issues and solutions for the effective implementation of OBS to carry IP traffic are presented.

## Switching Technologies for WDM

A number of switching technologies are available for transport of IP traffic over WDM. One approach is to use traditional circuit switching, which can be implemented in two ways. One way [7–8] is to eliminate the IP layer altogether

from the optical backbone by setting up optical paths between ingress-egress router pairs using the wavelength routing capability of the optical layer. However, this is justified only if there is a large volume of traffic between the ingress-egress router pair, since wavelength is a scarce resource. This is because the majority of resource allocation in the wavelength routing approach is coarse. The other way (currently used) is to use additional electronic layers (SONET/ATM/frame relay) between the IP and optical layers to provide flexible and granular access to bandwidth. This approach suffers from bandwidth inefficiencies and involves electronic processing at the intermediate nodes. An alternative to circuit switching is to use optical packet switching technology [9–10] in the backbone. This will provide the most flexible and efficient use of raw bandwidth available at the optical layer. Presently, this technology is not mature and has to overcome a number of technological constraints, such as limited optical buffering (implemented presently using optical delay lines) that can store only a few packets and difficulty in switching optical pulses using optics as control.

A viable alternative is to use OBS [11–15] in the backbone. OBS allows switching of data channels entirely in the optical domain by doing resource allocation in the electronic domain. In OBS, a control packet precedes every data burst. The control packet and the corresponding data burst are launched at the source at time instants separated by an *offset*. The control packet contains information required to route the data burst through the optical transmission core, the length of the burst, and the offset value, and it is sent over an out-of-band control channel (wavelength). It is processed electronically at each of the intermediate nodes (optical cross-connects) to make routing decisions (outgoing interface and wavelength), and the optical cross-connects are configured to switch the data burst which is expected to arrive after a time interval given by the offset field in the control packet. The impending data burst is then switched entirely in the optical domain, thereby removing the electronic bottleneck in the end-to-end data path. This results in *subwavelength allocation* in the sense that, at an outgoing interface, wavelength is allocated only for the duration of a data burst and can be statistically shared by data bursts belonging to different connections. A comparison of the characteristics of different switching technologies is given in Table 1.

A number of design and architectural issues need to be addressed before the OBS technique can be efficiently

Switching technology	Granularity	Utilization	Complexity
Circuit (wavelength) switching	Coarse	Poor	Low
Optical packet switching	Fine	High	High, not mature
Optical burst switching	Moderate	Moderate	Moderate

■ Table 1. A comparison of switching technologies for WDM.

deployed in IP networks. In the next section we give an architectural framework to support IP over WDM using OBS as the transmission mechanism, and discuss various design issues and solutions for them.

### A Framework for IP over WDM Using OBS

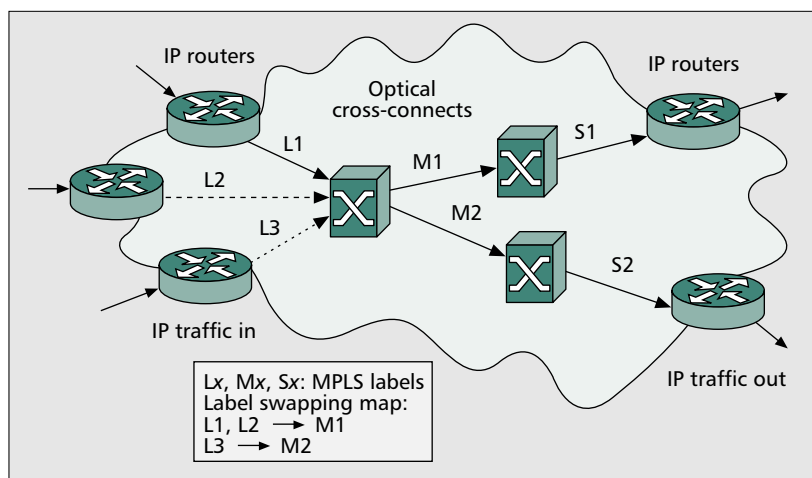
We envision an IP network operating over the optical backbone using OBS as the transmission mechanism as shown in Fig. 1. Incoming IP packets are assembled into data bursts at the ingress IP router and transported across the optical core to the destination egress IP router. There is either no buffering or very limited optical buffering (implemented using fiber delay lines) at optical cross-connects in the optical backbone. In the next section we discuss how MPLS-type techniques can be used with OBS to forward data bursts through the optical core. Then we describe the design issues at the OBS-MAC layer and control functions at optical cross-connects.

#### Optical Burst Switching Using MPLS

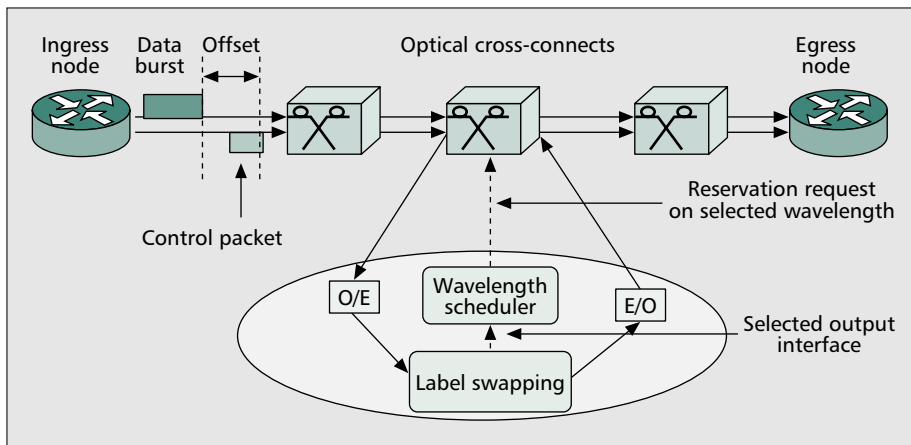
Semi-permanent data pipes can be set up between different ingress-egress router pairs using an MPLS-type technique. MPLS uses labels to make forwarding decisions at the network nodes, in contrast to the traditional destination-based hop-by-hop forwarding in IP networks. In MPLS, the space of all possible forwarding options is partitioned into *forwarding equivalence classes* (FECs). For example, all the packets destined for a given egress and having the same quality of service (QoS) may belong to the same FEC. The packets are labeled at the ingress depending on the FEC to which they belong. Each of the intermediate nodes uses the label of the incoming packet to determine its next hop, and also performs *label swapping* (i.e., replaces the incoming label with the new outgoing label that identifies the respective FEC for the downstream node). Such a label-based forwarding technique reduces the processing overhead required for routing at the intermediate nodes, thereby improving their packet forwarding performance and scalability. Also, the label swapping process used by MPLS creates multipoint-to-point packet forwarding trees, in contrast to a routing mesh in conventional networks based on a similar paradigm such as ATM networks.

Another important capability MPLS provides is *constraint-based routing*. The ingress node can establish an explicit route through the network. Rather than inefficiently carrying the explicit route in each packet, MPLS allows the explicit route to be carried only at the time the *label switched path* (LSP) is set up. The subsequent packets traversing this path are forwarded using packet labels. Constraint-based routing is potentially useful for traffic engineering.

MPLS can play a major role in traffic engineering and improving the throughput performance of an OBS-based network, as seen below. Each cross-connect in the optical backbone will have label swapping information about the precomputed routes in its *label infor-*



■ Figure 1. An IP-over-OBS WDM transmission backbone using MPLS.



■ Figure 2. Functions at the optical cross-connect supporting OBS and MPLS.

tion base (LIB). An LIB can be set up using standard techniques such as routing protocols with traffic engineering extensions to distribute information about the optical domain (available bandwidth per wavelength, number of wavelengths per fiber) and Constraint-Based Routing Label Distribution Protocol (CR-LDP) or Resource Reservation Protocol (RSVP) to distribute labels. Whenever an ingress router has a data burst to transmit, it refers to its LIB to determine the appropriate label. This label is included in the control packet that precedes this data burst. When the control packet arrives at any of the intermediate nodes, the following actions take place:

- The label in the control packet is used to point to the data-burst forwarding information in the LIB such as the output interface and any priority or QoS information.
- The cross-connect is set up to switch the data burst corresponding to that control packet in the all-optical domain. For this, information in the control packet about the length and offset of the data burst is used in addition to the forwarding information derived from the LIB. In particular, the latter is used to determine the mapping from the incoming fiber and wavelength to the outgoing fiber and wavelength. In order to be able to forward successive data bursts of the same connection (LSP) on different wavelengths in a given fiber, we propose that the label only specify incoming-fiber-to-outgoing-fiber mapping, while the information about the wavelength be appended to the outgoing label at every hop. The LIB may furnish other QoS information as well. Examples include defining a subset of candidate wavelengths on the outgoing fiber, determining the eligibility of that data burst to use wavelength conversion,<sup>1</sup> stating whether (in case of contention) the control packet is allowed to preempt some reservation already acquired by the control packet of low priority data burst, and so on.
- The control packet then undergoes label swapping (and wavelength information appending) and is forwarded on the dedicated control channel of the outgoing fiber as indicated by the LIB.

A functional diagram of an optical cross-connect supporting OBS and using MPLS-type forwarding is shown in Fig. 2.

<sup>1</sup> Note that wavelength converters are expensive, and hence, typically only a limited number of wavelength converters are provided at the cross-connect. This may necessitate restricted access to wavelength converters, say, limited only to premium class connections.

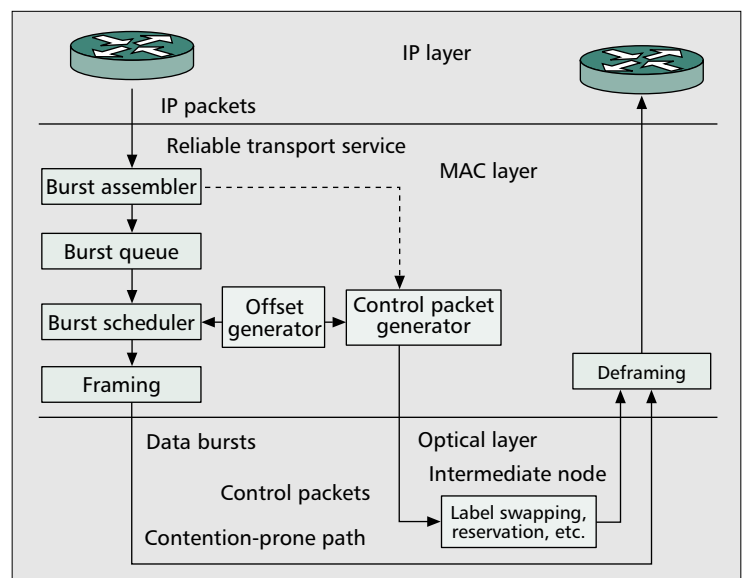
### The OBS MAC Layer

Apart from the architectural issues discussed above, there are a number of OBS-specific design issues that must be addressed when implementing the OBS scheme. A MAC layer is required between the IP and optical layers to implement these functions. We require that the optical layer using OBS appear to the IP layer as a reliable transmission medium that guarantees certain (low) burst blocking probability. Figure 3 illustrates the functional blocks needed at the OBS MAC layer. The key functions the OBS MAC layer must perform at the ingress router are:

- Assemble IP packets into bursts
- When the burst is at the head of the burst queue, determine the offset value to be used for this burst and launch a control packet that contains information about this offset, the length of the burst, and routing information (label)
- Frame the burst after the offset time has elapsed and send the burst into the optical layer

At the egress router, the OBS-MAC layer simply deframes the bursts and extracts IP packets from them.

An important OBS MAC design issue is determination of the offset between the control packet and the corresponding data burst. Of course, it should be large enough to allow for the processing of control packets at optical cross-connects in the LSP to minimize or eliminate optical buffering. In addition, an offset determination algorithm can be developed to reduce the contention among the data bursts arriving from different ingress routers at a node in the optical layer. The performance of a typical offset setting technique can be measured in terms of burst blocking probability. A fixed offset setting scheme called *just enough time* (JET) is proposed in [11] that gives better QoS to high-priority traffic by assigning longer offset values to its bursts. However, we believe that this scheme is not robust in a distributed environment due to contention at intermediate nodes among bursts arriving from geographically dispersed sources. Figure 4 illustrates a case in which control packets originating from two ingress routers, A



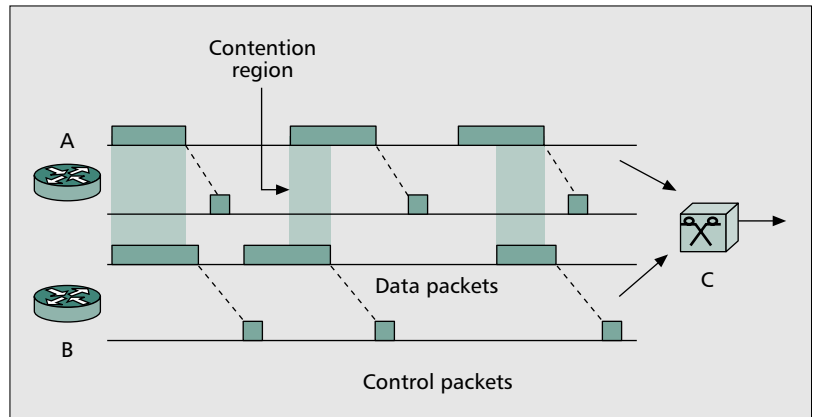
■ Figure 3. The MAC interface between the IP and OBS WDM layers.

and B, are nearly synchronized. If a fixed offset mechanism is deployed, an intermediate node C can only honor (assuming no buffering) the reservation request of A or B. This leads to a high burst blocking rate.

One solution would be to randomize the offset generation process. For example, it has been shown in [14, 15] that by using a statistical shaping scheme to determine offsets at the ingress router, LSPs can be engineered for the desired performance level in terms of burst blocking probability. Figure 5 depicts this scheme. This scheme is akin to a leaky bucket regulator in which tokens are generated by some random process, say Poisson process, and there is no buffering for the tokens. As soon as a burst is ready (at the head of the the burst queue), a control packet is launched with an offset value equal to the time difference between the current time and the arrival time of the next token according to the underlying random process. Note that an arriving token is lost if there is no fresh data burst available. Also, a token cannot be grabbed by a fresh data burst if the transmission time of a previous data burst overlaps the arrival time of the token. In that case, the next arriving token after the completion of the previous data burst is used for offset computation. For example, we see in Fig. 5 that a token arriving at time  $T_3(\omega)$  cannot be grabbed by the data burst arriving at  $a_3$  since  $T_3(\omega)$  overlaps the transmission time of the previous data burst. This approach has a number of desirable properties:

- It regulates the average rate at which data bursts are released into the OBS WDM layer over any LSP.
- Randomized generation of tokens prevents any systematic synchronization among the token streams of different optical LSPs. This is necessary because if the token generators of two LSPs traversing the same output interface of some node in the OBS WDM core happen to be synchronized, the data bursts in these LSPs will almost always collide with each other at that interface, causing excessive loss. The random generators for tokens at different routers can be initialized with, say, IP-address-specific seeds so that they never get in systematic synchronization.
- It has also been shown in [14, 15] that the above offset setting strategy imposes a certain statistical domination property on the stream of data bursts of an LSP. This property holds for the data stream of that LSP at every node it traverses between the ingress-egress pair. This turns out to be a useful handle for traffic engineering and QoS provisioning purposes.

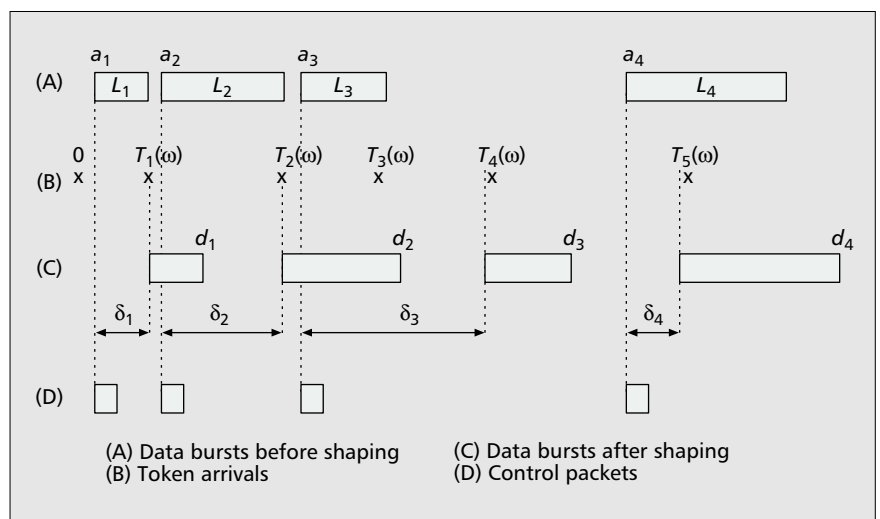
Burst assembly, which involves assembling IP packets into data bursts, is yet another important function to be performed in the OBS MAC layer at the ingress router. Here, burst size is an important design parameter, as explained below. The electronic processing speed of the control channel will limit the number of control packets and hence data bursts that can be transported per unit time across the optical channel. The electronic processing speed of the control channel is improved through the use of MPLS-type forwarding, as mentioned before. Thus, it is mainly limited by the speed of O/E and E/O conversions. Note that every packet passing through the



■ Figure 4. Synchronization among flow aggregates.

cross-connect causes one data burst to be forwarded through that node in the all-optical domain. Hence, roughly, if the ratio of the length data burst to that of control packet is  $\alpha$ , data transmission can take place at a rate that is  $\alpha$  times that of the electronic speed. However, the burst size should not be too large because it determines the delay encountered by IP packets, as described in the following paragraph. This is an important issue for real-time applications requiring tight end-to-end delay. Furthermore, some kind of priority scheduling scheme may be used to offer different QoS to different IP packet streams. Separate queues can then be provided for the data bursts at the ingress router depending on the QoS class of the burst.

Figure 3 depicts the flow of IP packets through the OBS MAC and optical layers. The delay encountered by IP packets consists of a fixed component of propagation delay through the optical layer and a variable component during their flow through the OBS MAC layer. First, IP packets incur delay during the burst assembly process. This is the time that elapses between the instant an IP packet is offered to the burst assembly block and the time the burst, encapsulating this packet, is offered to the burst queue. This delay is mainly determined by the maximum burst size as given below, and can be upper bounded by timer-based burst assembly. Second, IP packets incur queuing delay in the burst queue. This delay depends on packet arrival statistics and the nature of the burst scheduling algorithm used. Finally, there is additional offset delay which can be either fixed (deterministic offset generation) or variable (statistical



■ Figure 5. Data burst shaping.

Burst size (bytes)	Loss rate		Total throughput (Mb/s)		Throughput fairness	
	Shaping	No shaping	Shaping	No shaping	Shaping	No shaping
576	$0.60 \times 10^{-4}$	$7.40 \times 10^{-2}$	58.63	9.20	4.8	9.2
2304	$0.80 \times 10^{-4}$	$1.63 \times 10^{-2}$	57.20	34.00	3.0	23.0
4608	$0.62 \times 10^{-4}$	$1.51 \times 10^{-2}$	60.00	48.68	3.0	43.0

■ Table 2. Performance of statistical offset setting (burst blocking probability target =  $10^{-4}$ ).

offset generation). Hence, the total delay  $T_{MAC}$  incurred at the OBS MAC layer by packets comprises  $T_{BA}$ ,  $T_{BQ}$ , and  $T_{BO}$ , which represent the delays encountered during burst assembly, burst queuing, and maintaining the offset between the control packet and the burst, respectively. The average delay at the OBS MAC layer for the first packet in a burst can be given as

$$T_{MAC} = \frac{\bar{S}_B}{R_{IP}} + T_{BQ} + T_{BO}, \quad (1)$$

where  $\bar{S}_B$  is the maximum burst size and  $R_{IP}$  the average arrival rate of the IP traffic.  $T_{BO}$  is constant for a fixed offset setting scheme, and is determined by the subscribed rate of the LSP (inversely proportional to the subscribed rate) for a statistical shaping scheme.

#### OBS-Specific Control Functions at Cross-Connects

In addition to processing of the control packet, there are other control functions to be performed at the optical cross-connects in the transmission core. Once the output interface is determined using the LIB database, it is necessary to select the outgoing wavelength for the data burst. The goal is to share the bandwidth available across the parallel optical channels. The wavelength conversion capability available at the optical cross-connect is an important factor here. It can play a significant role in improving the utilization of available wavelengths in the network or reducing the blocking of data burst requests. Without wavelength conversion, the same wavelength needs to be assigned to an impending data burst at both incoming and outgoing interfaces. It may cause increased blocking of data burst requests. Full wavelength conversion is desirable, but difficult to implement all-optically. However, limited wavelength conversion can be implemented more easily. By cleverly designing the wavelength assignment and scheduling algorithm, it is possible to come out with an efficient design even with limited conversion. Similar issues, albeit in the context of wavelength routing, are addressed in [16] and elsewhere.

#### IP QoS Mechanisms

There are a number of ways to implement IP QoS procedures in the optical backbone using OBS. First, we observe from an earlier discussion (for details see [14, 15]) that it is possible to engineer an LSP to a certain low burst blocking probability by using statistical shaping in the OBS MAC layer and employing admission control during LSP setup. To elaborate, we note that by virtue of the statistical shaping scheme, the statistics of the arrival of bursts over an LSP is known and invariant even after the LSP traverses a number of nodes, due to inherently bufferless forwarding of data bursts. Hence, at the time of setting up a new LSP, every output interface in its path can predict whether it can maintain low burst blocking probability even after admitting this LSP. Also, by properly engineering the OBS MAC layer, desired delay characteristics can be obtained along the LSP. Hence, the LSP appears to the IP layer as a reliable data pipe with a certain provisioned data

rate, low burst blocking rate, and certain transfer delay. Then, standard IP QoS mechanisms such as differentiated services (DiffServ) can be used at the IP layer. Another IP QoS dimension is that the service provider may partition the backbone network to provide different levels of reliability along LSPs. Standard QoS routing schemes can then be

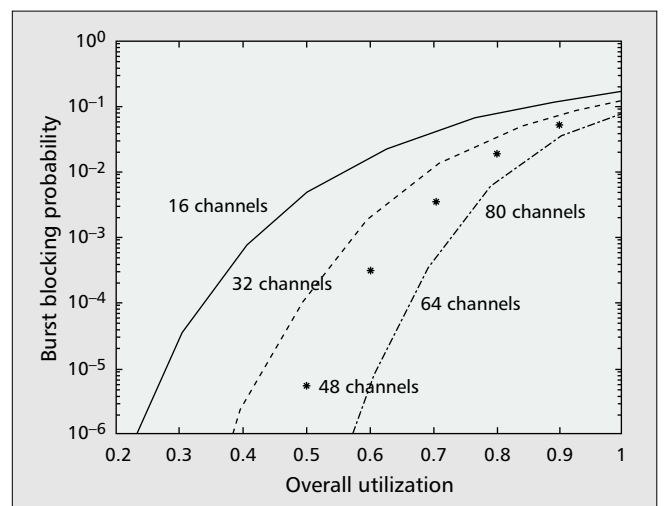
used in access networks to route packets to appropriate edge routers depending on their QoS needs (in terms of loss rate).

#### Network Utilization vs. Burst Blocking Probability

The network deploying OBS needs to be engineered to operate at a high level of utilization, with low burst blocking probability. As mentioned before, optical buffering technology is still immature, and only limited buffering is available in the optical domain using optical delay lines. Availability of a large number (such as 160 channels in a single fiber) of parallel optical channels can improve the utilization of a burst switched network even with limited or even no buffering. For example, if we assume that data bursts are arriving at a bottleneck node with Poisson rate  $\lambda$  (this can be enforced via the shaping scheme of [14, 15]) and the number of channels is  $c$ , then with no buffering the burst blocking probability can be calculated by using the well-known Erlang loss formula,

$$P_{\text{Poisson}}(\lambda) = \frac{(\lambda/\mu)^c / c!}{\sum_{i=0}^c (\lambda/\mu)^i / i!} \quad (2)$$

Figure 6 shows the burst blocking probability as a function of the offered load for an interface supporting 16, 32, 48, 64, and 80 optical channels (wavelengths). We observe that for the target blocking probability of  $10^{-4}$ , it is possible to operate a burst switched network with 60 percent utilization if the underlying DWDM technology supports 64 optical channels/fiber. This shows that it is possible to operate a burst switched network at high utilization if the number of channels per fiber is sufficiently large, even without any buffering.



■ Figure 6. Burst blocking probability.

## Performance Analysis

In this section we report some simulation results from [14, 15] to compare the performance of statistical and fixed offset setting schemes. More information about simulation experiments and numerical results can be found in [14, 15]. In these simulations, an output interface of a cross-connect with three OC-12 (622 Mb/s) wavelengths is considered. This interface is traversed by a number of LSPs, each carrying traffic from a number of TCP and UDP flows. Independent Poisson processes were used to generate random offsets (statistical shaping) at the ingress of these LSPs. The results of the simulations are given in Table 2. They show that the proposed offset setting/traffic shaping scheme maintains the burst blocking probability below the designed value. When shaping is not used (fixed offsets), there is no control over the burst blocking probability. We also observe that fixed offset setting causes high blocking of data bursts due to synchronization among TCP flows. The impact of high blocking probability is seen in reduced total throughput and worse fairness among TCP flows. Fairness is measured as the ratio of throughput of the flow which gets maximum to that which receives minimum throughput.

## Conclusion

WDM has emerged as a core transmission technology for the Internet backbone to meet the growing bandwidth demand of IP traffic. It provides enormous capacity and configurability at the optical layer. In this article we discuss optical burst switching as a currently viable switching technology to exploit this bandwidth potential. We provide a framework to deploy IP over WDM using OBS. Using this framework, label-switched paths of desired data rates can be set up across an optical backbone network. Furthermore, these LSPs can be engineered to provide desired delay and loss characteristics. The reliability of LSPs and traffic engineering in the optical backbone are particularly enabled by the statistical offset setting scheme. We also provide generic guidelines to support IP QoS mechanisms such as DiffServ and QoS routing in the context of this framework.

## References

- [1] C. Brackett, "Dense Wavelength Division Multiplexing Networks: Principles and Applications," *IEEE JSAC*, vol. 8, no. 6, Aug. 1990, pp. 948–64.
- [2] P. Newman *et al.*, "IP Switching and Gigabit Routers," *IEEE Commun. Mag.*, vol. 35, no. 1, Jan. 1997, pp. 64–69.
- [3] E. Varvarigos and V. Sharma, "The Ready-To-Go Virtual Circuit Protocol: A Loss-Free Protocol for Multigigabit Networks Using FIFO Buffers," *IEEE/ACM Trans. Net.*, vol. 5, no. 5, Oct. 1997, pp. 705–18.
- [4] A. Viswanathan *et al.*, "Evolution of Multiprotocol Label Switching," *IEEE Commun. Mag.*, vol. 36, no. 5, May 1998, pp. 165–73.
- [5] J. Manchester *et al.*, "IP over SONET," *IEEE Commun. Mag.*, vol. 36, no. 5, May 1998, pp. 136–42.
- [6] Eytan Modiano, "WDM-Based Packet Networks," *IEEE Commun. Mag.*, vol. 37, no. 3, Mar. 1999, pp. 130–35.

- [7] D. Awduche *et al.*, "Multi-Protocol Lambda Switching," IETF draft-awduche-mppls-te-optical-02.txt, Jan. 2001.
- [8] Nasir Ghani, "A Framework for IP-Over-WDM using MPLS," *Opt. Networks Mag.*, vol. 1, no. 2, April 2000, pp. 39–52.
- [9] Piero Gambini *et al.*, "Transparent Optical Packet Switching: Network Architecture and Demonstrators in the KEOPS Project," *IEEE JSAC*, vol. 16, no. 7, Sept. 1998, pp. 1245–59.
- [10] Yoshiaki Yamada *et al.*, "Optical Output Buffered ATM Switch Prototype Based on FRONTIERNET Architecture," *IEEE JSAC*, vol. 16, no. 7, Sept. 1998, pp. 1298–1307.
- [11] M. Yoo and C. Qio, "Optical Burst Switching (OBS)—A New Paradigm for an Optical Internet," *Int'l. J. High Speed Networks*, vol. 8, no. 1, 1999.
- [12] Jonathan Turner, "Terabit Burst Switching," *Int'l. J. High Speed Networks*, vol. 8, no. 1, 1999.
- [13] A. Ge, F. Callegati, and L. Tamil, "On Optical Burst Switching and Self-Similar Traffic," *IEEE Commun. Letts.*, vol. 4, no. 3, Mar. 2000, pp. 98–100.
- [14] H. Chaskar, S. Verma, and R. Ravikanth, "A Framework to Support IP Traffic Over WDM Using Optical Burst Switching," *IEEE/ACM/SPIE Opt. Network Wksp.*, Jan. 2000.
- [15] H. Chaskar, S. Verma, and R. Ravikanth, "Robust Transport of IP Traffic Over WDM Using Optical Burst Switching," Nokia internal rep., Jan. 2001.
- [16] S. Subramaniam, M. Azizoglu, and A. Samani, "All-Optical Networks with Sparse Wavelength Conversion," *IEEE/ACM Trans. Net.*, vol. 4, no. 4, Aug. 1996, pp. 544–57.

## Biographies

SANJEEV VERMA (sanjeev.verma@nokia.com) received his B.E. degree in electronics and telecommunications engineering from Calcutta University, India, and his M.Tech. degree in electrical engineering from the Indian Institute of Technology, Kanpur, in 1984 and 1988, respectively. From 1988 to 1992 he worked as an R&D engineer in the Signaling System No. 7 group at the Center for the Development of Telematics, Delhi, India. In 1998 he obtained his Ph.D. degree from the University of Toronto. From December 1997 to July 1999, he served as a member of technical staff at Lucent Technologies, Holmdel, New Jersey, where he was involved in the design and implementation of routing protocols in the PacketStar router project. Since August 1999 he has been with the Internet Multimedia Research group at Nokia Research Center, Boston, Massachusetts. His research interests are in QoS issues in high-speed networks, routing protocols, resource management, multicasting, and streaming applications.

HEMANT M. CHASKAR [M'99] (hemant.chaskar@nokia.com) received his Ph.D. degree in electrical engineering from the University of Illinois, Urbana-Champaign, in 1999. Prior to that, he received a Bachelor of Engineering degree from the University of Pune, India, in 1993, and a Master of Engineering degree from the Indian Institute of Science, Bangalore, in 1995, both in electronics and telecommunication engineering. From August 1995 to February 1999 he was a research assistant in the Coordinated Sciences Laboratory at the University of Illinois, Urbana-Champaign. He spent the summer of 1997 in the Performance Analysis Department at Lucent Technologies-Bell Labs, Holmdel, New Jersey. Currently he is with Nokia Research Center, Boston, Massachusetts. His research interests are in wireless communications, packet scheduling and congestion control in high-speed networks, and optical networks. He has worked in areas related to multimedia Internet with emphasis on differentiated services and MPLS, QoS on IP, cellular IP networking, and architecture and services in third generation wireless networks.

RAYADURGAM RAVIKANTH (rravikanth@axiowave.com) received his B.Tech. degree from the Indian Institute of Technology, Madras, in 1988 and M.S. and Ph.D. degrees from the University of Illinois at Urbana-Champaign in 1991 and 1996, all in electrical engineering. Since 1996 his interests have been focused in the area of Internetworking, specifically QoS on IP, MPLS, performance analysis, Internet traffic modeling, and IP over wireless networks. From July 1996 to July 2000 he worked at Nokia Research Center, Burlington, Massachusetts, conducting and managing research in the above areas. He is currently with Axiowave Networks, Marlborough, Massachusetts.