

Quality-of-Service Mechanisms in IP-over-WDM Networks

Ayman Kaheel, Tamer Khattab, Amr Mohamed and Hussein Alnuweiri

Department of Electrical and Computer Engineering

The University of British Columbia

Vancouver, B.C., Canada V6T 1Z4

Email: aymank@ece.ubc.ca

ABSTRACT

Classical approaches to quality-of-service (QoS) provisioning in IP networks are difficult to apply in all-optical networks. This is mainly because there is no optical counterpart to the store-and-forward model that mandates the use of buffers for queuing packets during contention for bandwidth in electronic packet switches. Since plain IP assumes a best effort service model, there is a need to devise mechanisms for QoS provisioning in IP over wavelength-division-multiplexing, or IP-over-WDM, networks. Such mechanisms must consider the physical characteristics and limitations of the optical domain. This paper presents a classification and a survey of recent proposals for QoS provisioning and enforcement in IP-over-WDM networks. The different QoS proposals surveyed cover three major optical switching methods: wavelength routing, optical packet switching and optical burst switching.

1. INTRODUCTION

The proliferation of Internet Protocol (IP) technology coupled with the vast bandwidth offered by optical wavelength-division-multiplexing (WDM) technology are paving the way for IP-over-WDM to become the primary means for transporting data across large distances in the next generation Internet. WDM is an optical multiplexing technique that allows better exploration of the fiber capacity by simultaneously transmitting data packets over multiple frequencies, or wavelengths. The tremendous bandwidth offered by WDM is promising to reduce the cost of core network equipment and simplify bandwidth management. However, the problem of providing quality-of-service (QoS) guarantees for several advanced services, such as transport of real-time packet voice and video, remains largely unsolved for optical backbones. The QoS problem in optical WDM networks has several fundamental differences from QoS methods in electronic routers and switches. One major difference is the absence of the concept of “packet queues” in WDM devices, beyond the number of packets that can be buffered (while in-flight) in Fiber Delay Lines (FDLs). FDLs are long fiber lines used to delay the optical signal for a particular period of time. As an alternative to queuing, optical networks use additional signaling to reserve bandwidth on a path ahead of the arrival of optically switched data.

Over the past decade, a significant amount of work has been dedicated to the issue of providing QoS in non-WDM IP networks. Basic IP assumes a best-effort service model. In this model, the network allocates bandwidth to all active users as best as it can, but does not make any explicit commitment as to bandwidth, delay or actual delivery. This service model is not adequate for many real-time applications that normally require assurances on the maximum delay of transmitting a packet through the network connecting the end points. A number of enhancements have been proposed to enable offering different levels of QoS in IP networks. This work has culminated in the proposal of the Integrated Services (Intserv) [1] and the Differentiated Services (Diffserv) [2] architectures by the IETF. Intserv achieves QoS guarantees through end-to-end resource (bandwidth) reservation for packet flows and performing per-flow scheduling in all intermediate routers or switches. Diffserv, on the other hand, defines a number of per-hop behaviors that enable providing relative QoS advantage for different classes of traffic aggregates. Both schemes require sources to shape their traffic as a precondition for providing end-to-end QoS guarantees.

Since Internet traffic will eventually be aggregated and carried over the core networks, it is imperative to address end-to-end QoS issues in WDM networks. However, previous QoS methods proposed for IP networks are difficult to apply in WDM networks mainly due to the fact that these approaches are based on the store-and-forward model and mandate the use of buffers for contention resolution. Currently there is no optical memory and the use of electronic memory in an optical switch necessitates optical-to-electrical (O/E) and electrical-to-optical (E/O) conversions within the switch. Using O/E and E/O converters limits the speed of the optical switch. In addition, switches that utilize O/E and E/O converters lose the advantage of being bit rate transparent. Furthermore, these converters increase the cost of the optical switch significantly. The only means currently for providing a limited buffering capability in optical switches is the use of FDLs. However, FDLs cannot provide the full buffering capability required by the classical QoS approaches. In addition to FDLs, the wavelength domain provides a further opportunity for contention resolution based on the number of wavelengths available and the wavelength assignment method.

This paper classifies and surveys different approaches that have been proposed for implementing service differentiation in WDM networks with different switching techniques. It is our aim in this survey to present general mechanisms for providing QoS in WDM networks, and give examples of proposals that implement and enhance these mechanisms. We present an overview of the different switching techniques employed in optical networks in Section 2. Then we provide a survey of the different mechanisms for QoS in WDM networks in Section 3, and provide concluding remarks in Section 4.

2. OPTICAL SWITCHING TECHNIQUES

Three major switching techniques have been proposed in the literature for transporting IP traffic over WDM-based optical networks. Accordingly, IP-over-WDM networks can be classified as wavelength routing (WR) networks, optical packet switching (OPS) networks and optical burst switching (OBS) networks.

Wavelength Routing Networks

In WR networks an all-optical wavelength path is established between edges of the network. This optical path is called a lightpath and is created by reserving a dedicated wavelength channel on every link along the path as shown in Figure 1. After data is transferred the lightpath is released. WR networks consist of optical cross-connect (OXC) devices connected by point-to-point fiber links in an arbitrary topology. OXC devices are capable of differentiating data streams based on the input port from which a data stream arrives and its wavelength [13]. As a result, data transmitted between lightpath endpoints require no processing, no E/O conversion and no buffering at intermediate nodes. However, being a form of circuit switching networks, WR networks do not use statistical sharing of resources and therefore, provides lower bandwidth utilization.

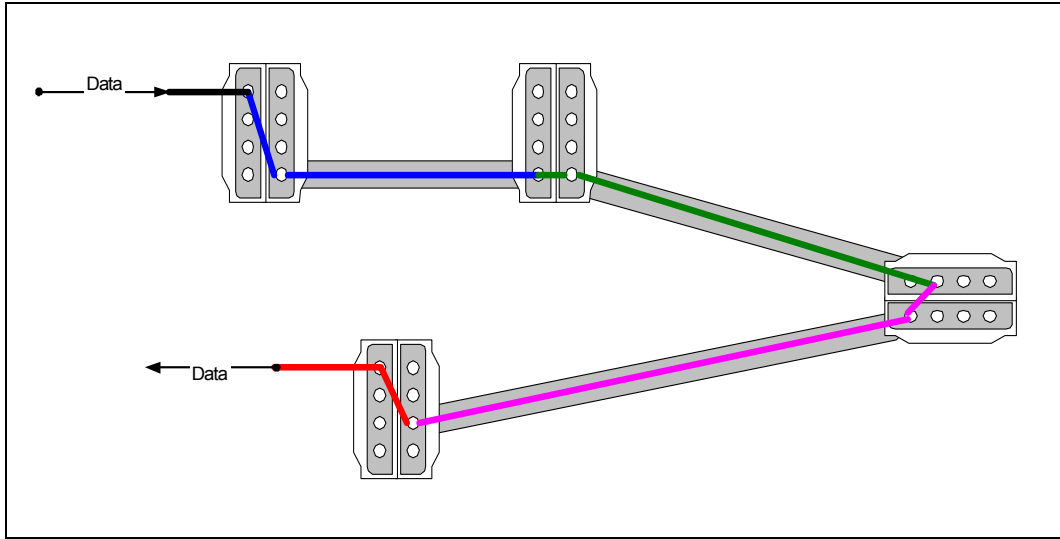


Figure 1 – Lightpath establishment

Optical Packet Switching Networks

In packet switching networks, IP traffic is processed and switched at every IP router on a packet-by-packet basis. An IP packet contains a payload and header. The packet header contains the information required for routing the packet while the payload carries the actual data. The future and ultimate goal of OPS networks is to process the packet header entirely in the optical domain. With the current technology, it is not possible to do such processing in the optical domain. A solution for this problem is to process the header in the electronic domain and keep the payload in the optical domain. Nevertheless, many technical challenges remain to be addressed for this solution to become viable. The main advantage of OPS is that it can increase the network's bandwidth utilization by utilizing statistical multiplexing for bandwidth sharing.

Optical Burst Switching Networks

OBS networks combine the advantages of both WR networks and OPS networks. As in WR networks, there is no need for buffering and electronic processing for data at intermediate nodes. At the same time, OBS increases the network utilization by reserving the channel for a limited time period. The basic switching entity in OBS is a burst. A

burst is a train of packets moving together from one ingress node to one egress node and switched together at intermediate nodes. A number of approaches exist for burst forming, such as, the *containerization with aggregation-timeout* (CAT) technique proposed by the authors [15]. A burst consists of two parts header and data. The header is called the control burst (CB) and is transmitted separately from the data, which is called the data burst (DB). The CB is transmitted first to reserve the bandwidth along the path for the corresponding DB. Then it is followed by the data burst, which travels over the same path reserved by the control burst.

Several signaling protocols have been proposed for OBS [14]. In this paper, we explain one of these protocols, called Just-Enough-Time (JET) protocol.

In JET, the CB is sent first on a control channel and then followed by the DB on a data channel with a time delay equal to the burst offset time (T_o). When the CB reaches a node, it reserves a wavelength on the outgoing link for a duration equal to the burst length starting from the arrival time of the DB.

3. QOS IN IP-OVER-WDM NETWORKS

Several approaches have been proposed for implementing service differentiation in optical networks. Early approaches proposed smart queue management to guarantee different packet loss probabilities to different packet streams. Examples of these algorithms are threshold dropping and priority scheduling. Nevertheless, in this paper we are more interested in presenting approaches that exploit the unique characteristics of the optical domain.

3.1. QoS in WR networks

We present here a general framework for providing differentiated service in WR networks. This framework extends the Differentiated Optical Services (DoS) model presented in [3]. We consider other QoS proposals for WR networks in the context of DoS.

The DoS model considers the unique optical characteristics of lightpaths. A lightpath is uniquely identified by a set of optical *parameters* such as bit error rate (BER), delay, jitter, etc. and *behaviors* including protection, monitoring, and security capabilities. These optical parameters and behaviors provide the basis for measuring the quality of optical service available over a given path. The purpose of such measurements is to define classes of optical services equivalent to the IP QoS classes. The DoS framework consists of six components.

1. **Service Classes:** A DoS service class is qualified by a set of parameters that characterize the quality and impairments of the optical signal carried over a lightpath. These parameters are either specified in quantitative terms, such as delay, average BER, jitter, and bandwidth, or based on functional capabilities such as monitoring, protection and security.

2. **Routing and Wavelength Assignment Algorithm:** In order to establish a lightpath, a dedicated wavelength has to be reserved throughout the lightpath route. An algorithm used for selecting routes and wavelengths to establish lightpaths is known as a *routing and wavelength assignment* (RWA) algorithm. In order to provide QoS in WR networks, it is mandatory to use a RWA algorithm that considers the QoS characteristics of different wavelength channels. An example for RWA algorithms is presented in [4]. The underlying idea behind this RWA algorithm is to employ adaptive weight functions that characterize the properties of different wavelength channels (such as delay, capacity, etc.).

3. **Lightpaths Groups:** Lightpaths in the network are classified into groups that reflect the unique qualities of the optical transmission, such that each group corresponds to a DoS service.

4. **Traffic Classifier:** Traffic flows are classified into one of the supported classes by the network. Classification is done at the network ingress.

5. **Lightpaths Allocation Algorithm:** A number of algorithms have been proposed in literature for allocating lightpaths to different service classes [6] [7]. We discuss these algorithms in following section.

6. **Admission Control:** Similar to the bandwidth broker entity in differentiated services architecture, an entity called optical resource allocator is required in WDM networks to handle the dynamic provisioning of lightpaths [3]. The optical resource allocator keeps track of the resources, such as the number of wavelengths, links, cross-connects, and amplifiers, available for each lightpath, and evaluates the lightpath characteristics (BER computation) and functional capabilities (protection, monitoring and security). The optical resource allocator is also responsible for initiating end-to-end call setup along the chain of optical resource allocators representing the different domains traversed by the lightpath.

All of the above mentioned components are implemented in the edge devices and/or the optical resource allocator, Figure 2 shows a WR network with edge devices and optical resource allocator and interior OXC devices. The interior OXC devices are only required to configure the switching core to setup the required lightpaths.

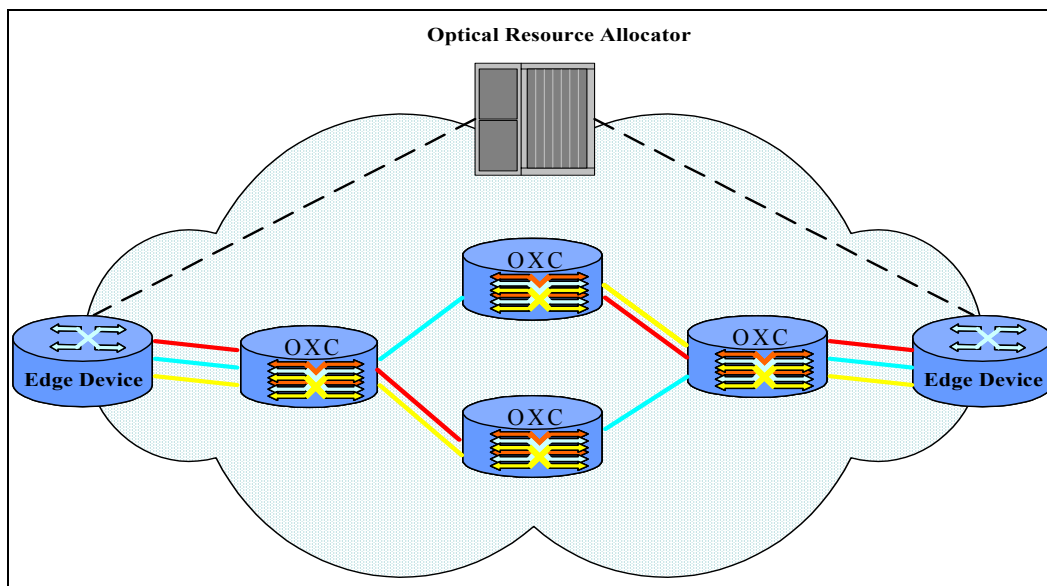


Figure 2 - WR network

Lightpath Allocation (LA) Algorithms

In general LA algorithms partition the available lightpaths into different subsets. Each subset is assigned to a service class. LA approaches differ from each other in the way

lightpaths subsets are allocated to service classes. This allocation can be static, static with borrowing or dynamic.

In case of static allocation, a fixed subset of lightpaths is assigned to each of the service classes. The number of lightpaths in each of the subsets depends on the service class i.e. higher service classes are allocated more lightpaths.

When borrowing is allowed, different priority classes can borrow lightpaths from each other according to certain criteria. Different approaches for such borrowing are described in [6]. An example approach is to allow lower priority traffic to borrow lightpaths from higher priority traffic. However, borrowing in the reverse direction is not allowed because lightpaths originally assigned to lower priority traffic may not satisfy the QoS requirements of higher priority classes.

In dynamic approaches, the network starts with no reserved lightpaths for service classes. The available pool of lightpaths can then be assigned dynamically to any of the available service classes, under the assumption that all lightpaths have similar characteristics. One approach for dynamic lightpath allocation is to use proportional differentiation [7]. In the proportional differentiation model, one can quantitatively adjust the service differentiation of a particular QoS metric to be proportional to the differentiation factors that a network service provider sets beforehand. In other words, If q_i is the QoS metric of interest and s_i is the differentiation factors for class i , then using the proportional differentiation model we should have $\frac{q_i}{q_j} = \frac{s_i}{s_j}$ where $i, j = 1, \dots, N$ for all pairs of service

classes. For example, assume that q_1, q_2 are the packet loss probability for classes 1 and 2 respectively. If s_1 is 1 while s_2 is 2, then the packet loss probability of class 2 should be twice that of the packet loss probability of class 1. Consequently, the number of lightpaths assigned to class 1 must be twice the number assigned to class 2.

3.2. QoS in Optical Packet Switching Networks

The idea underlying most proposals for optical packet switching is to decouple the data-path from the control path. This way, routing and forwarding functions are performed

using electronic chips after an optical-to-electrical conversion of the packet header, while the payload is switched transparently in the optical domain without any conversion. Up till now, there have been very few proposals for providing service differentiation in OPS networks. This is expected considering that OPS is a fairly new switching technique and still has many problems that remain to be solved.

In any packet switching scenario, contention may arise when more packets are to be forwarded to the same output link at the same time. In general, QoS techniques in OPS networks aim at providing service differentiation when contention occurs by using wavelengths and FDLs assignment algorithms. The authors in [5] have presented two algorithms for service differentiation in optical packet switches. We give an overview of these algorithms as general techniques for providing QoS in OPS networks.

- **Wavelength Allocation (WA):** This technique divides the available wavelengths into disjoint subsets and assigns each subset to a different priority level such that higher priority levels get a larger share of the available wavelengths. Different wavelength allocation algorithms are possible, which are similar to lightpath allocation algorithms presented in Section 3.1. WA techniques use the wavelength domain only for service differentiation and do not utilize FDL buffers.

- **Combined Wavelength Allocation and Threshold Dropping (WATD):** In addition to WA, this technique uses threshold dropping to differentiate between different priority classes. When the FDLs buffer occupancy is above a certain threshold, lower priority packets are discarded. By using a different dropping threshold for each priority level, different classes of service can be provided. This technique exploits both wavelength domain (WA) and time domain (FDLs) to provide service differentiation and hence it has more computational complexity than the buffer-less WA technique.

Although the techniques presented here seem simple, their implementations in OPS networks can be complex because of the required synchronization between the packet header and the packet payload. This process requires the packet payload to be delayed until the header is fully processed and the packet is classified, after which the packet is assigned a wavelength. This is done on a packet-by-packet basis, which limits the switching speed. Moreover, since packets in FDLs cannot be randomly accessed as in the

case of electronic buffers, new elaborate techniques are required to access individual variable-sized packets stored in FDLs.

3.3. QoS in Optical Burst Switching Networks

In this section we survey approaches for QoS provisioning in OBS networks. Providing QoS in OBS networks requires a signaling (reservation) protocol that supports QoS. In addition, a burst-scheduling algorithm is needed in the network core burst switches.

3.3.1. Signaling Protocols with QoS Support

In [8] a protocol for supporting quality of service in optical burst switching is proposed. This protocol is based on the JET protocol and is called prioritized JET (pJET). This protocol uses offset time as a way to provide different classes of service in buffer-less optical networks. Assume we have two classes of service, class 1 which is the high priority service class and class 0 which is the best effort service class. In order for class 1 to have higher priority for bandwidth reservation, an additional offset time, denoted t_{offset} , is given to this class. The value of t_{offset} is constant and considerably larger than the original JET offset time T_0 . In addition, it needs to be larger than the maximum burst length over all bursts in the class 0. With that much offset time, the blocking probability of bursts in class 1 becomes independent of the offered load in class 0, and only a function of the offered load in class 1. On the other hand, the offered load in both classes will determine the blocking probability of class 0. Figure 3 illustrates why a class 1 request that is assigned t_{offset} obtains a higher priority for reservation than a class 0 request. Let t_{ai} and t_{si} be the arrival time and the service-start time for a class i request (denoted by $req(i)$), where $i = 0, 1$ in our example, and let l_i be the burst requested by class i . Consider the following two situations where conflicts among the two classes are possible. In the first case as illustrated in Figure 3(a), $req(1)$ comes first and reserves bandwidth, and $req(0)$ comes afterwards.

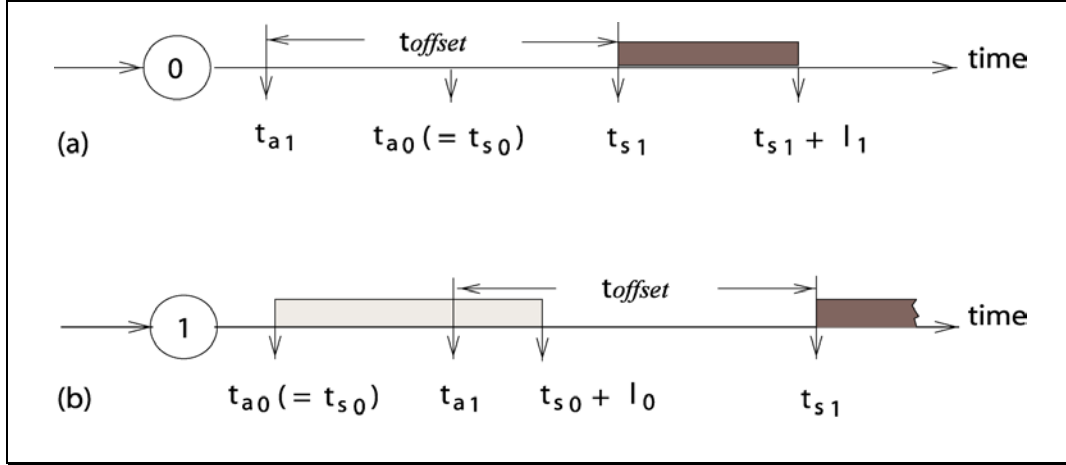


Figure 3 – Offset time for guaranteed service

Clearly, $req(1)$ will succeed, but $req(0)$ will be blocked if $t_{a0} < t_{s1}$ and $t_{a0} + l_0 > t_{s1}$, or if $t_{a0} < t_{s1} + l_1$. In the second case, $req(0)$ arrives, followed by $req(1)$ as shown in Figure 3(b). It can be seen that as long as t_{offset} for class 1 is larger than maximum burst length in class 0, any request from class 1 can avoid being blocked by a request from class 0. Authors in [8] also gave a simple analytical model to evaluate the blocking probability as a function of t_{offset} and concluded that to provide almost 100% isolation between classes 0 and 1, it is sufficient to have t_{offset} equal to $5L_0$, where L_0 is the average burst size in class 0. The main problem with this algorithm is that it introduces a significant amount of delay for high priority traffic.

3.3.2. Scheduling in OBS

When a control burst arrives at node, a wavelength channel-scheduling algorithm is used to determine the wavelength channel (and also FDLs if available) on an outgoing link for the corresponding data burst. The information required by the scheduler such as the burst's arrival time and its duration are obtained from the control burst. The scheduler keeps track of the availability of the time slots on every wavelength channel. If FDLs are available at the node, the scheduler selects one or more FDLs to delay the data burst, if necessary. A wavelength channel is said to be unscheduled at time t when no burst is

using the channel at or after time t . A channel is said to be unused for the duration of voids between successive bursts and after the last burst assigned to the channel.

Several issues affect the performance of the OBS scheduler. First, it must select wavelength channels and FDLs in an efficient way to reduce burst dropping probability. In addition, it must be simple enough to be able to handle large number of bursts in a very high-speed environment. Furthermore, the scheduler must not lead to "early data burst arrival" situation, in which case, the data burst will arrive before the control burst has been processed.

A number of wavelength channel-scheduling algorithms have been proposed in the literature [9][10][11][12]. The following sections describe these algorithms.

3.3.2.1. First Fit Unscheduled Channel (FFUC) Algorithm

For each of the outgoing wavelength channels, the FFUC algorithm keeps track of the unscheduled time. Whenever a control burst arrives, the FFUC algorithm searches all wavelength channels in a fixed order and assigns the burst to the first channel that has unscheduled time less than the data burst arrival time. This algorithm's main advantage is its computational simplicity. Its main drawback is that it results in high dropping probability, as the algorithm does not consider voids between bursts scheduling.

3.3.2.2. Latest Available Unscheduled Channel (LAUC) Algorithm

The basic idea of LAUC algorithm is to increase channels utilization by minimizing voids created between bursts. This is accomplished by selecting the latest available unscheduled data channel for each arriving data burst. For example, in Figure 4 wavelength 1 and 2 are unscheduled at time t_a , and wavelength 1 will be selected to carry the new data burst arriving at t_a . In which case, the void on wavelength 1 will be smaller than the void that would have been created if wavelength 2 were selected. Therefore, LAUC yields better burst dropping performance than FFUC algorithm while it does not add any computation overhead. However, since it does not take advantage of voids

between bursts, as was the case for the FFUC, it still leads to relatively high dropping probability.

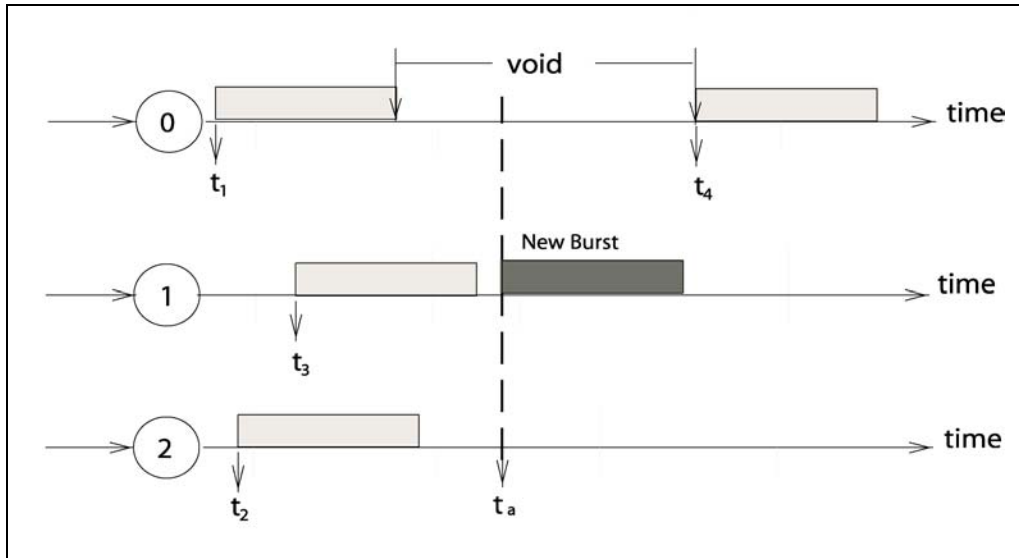


Figure 4 Illustration of LAUC algorithm

3.3.2.3. LAUC with Void Filling (LAUC-VF) Algorithm

The void/gap between the two data bursts in wavelength 1 of Figure 4 is unused channel capacity. The LAUC-VF algorithm is similar to LAUC algorithm except that voids can be filled by new arriving bursts. The basic idea of this algorithm is to minimize voids by selecting the latest available *unused* data channel for each arriving data burst. Given the arrival time t_a of a data burst with duration L to the optical switch, the scheduler first finds the outgoing data channels that are available for the time period of $(t_a; t_a + L)$. If there is at least one such data channel, the scheduler selects the latest available data channel, i.e., the channel having the smallest gap between t_a and the end of the last data burst just before t_a . Figure 5 shows an illustration of LAUC-VF algorithm. A new burst arrives at time t_a . At time t_a wavelengths 1 and 3 are ineligible because the void on channel 1 are too small for the new burst, while channel 3 are busy. The LAUC-VF algorithm chooses channel 2 since this will produce the smallest gap.

Since the voids are used effectively, the LAUC-VF algorithm yields better performance in terms of burst dropping probability than FFUC and LAUC algorithms. On the other

hand, the algorithm is more complex than FFUC and LAUC algorithms because it keeps track of two variables instead of one.

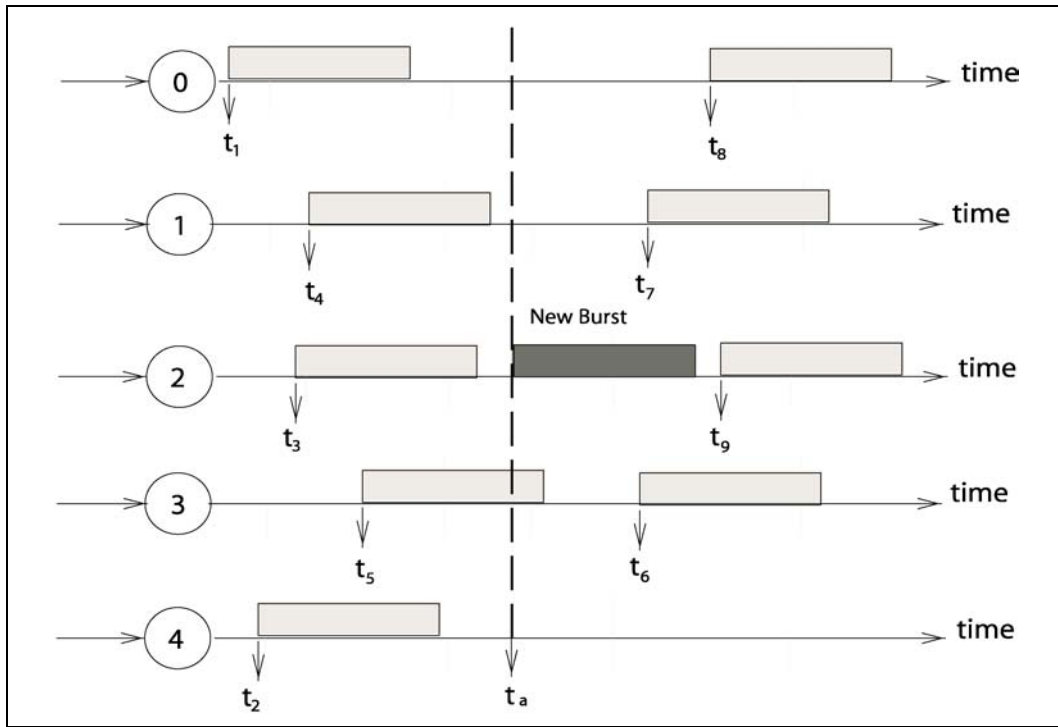


Figure 5 - Illustration of LAUC-VF algorithm

3.3.2.4. Generalized LAUC-VF (G-LAUC-VF) Algorithm

G-LAUC-VF algorithm generalizes the LAUC-VF algorithm to include QoS features.

At core switches, the scheduler associated with a link will schedule data bursts (DB) going to that output link. For each link, its scheduler maintains n queues $Q_1, Q_2 \dots Q_n$, with Q_i being used to store the control bursts (CB) of class i in FIFO order.

For each slot, the algorithm is executed once. Assuming that class i have higher priority than that of class j if $i < j$. The algorithm is as follows.

G-LAUC-VF Algorithm

```
for  $i = 1$  to  $n$  do  
    while  $Q_i$  has CB belonging to the current slot do  
         $CB_i = \text{dequeue}(Q_i)$ ;  
        Use LAUC-VF algorithm to schedule  $DB_i$  corresponding to  $CB_i$ ;  
    end  
end
```

This algorithm ensures that for those DBs whose CBs are in Q_i are scheduled before DBs whose CBs are in Q_j , if $i < j$.

4. CONCLUSION

In this paper we have surveyed different proposals for QoS provisioning in IP-over-WDM networks. We have presented general QoS mechanisms in wavelength routing networks, optical packet switching networks and optical burst switching networks. Proposals for these mechanisms are in different stages of maturity. QoS proposals for WR networks are the most mature proposals compared to those for OPS and OBS. This is due to the simplicity of the switching technique itself, and the fact that no optical buffers are needed to implement these proposals. On the other hand, proposals for QoS provisioning in OPS are still in early stages of research and many problems need to be addressed before these proposals become viable. However, QoS schemes in OBS networks are very promising since they are simple and require no buffering. It is evident from the research results we surveyed that overall and collectively much work is still needed before QoS mechanisms will be widely deployed in IP-over-WDM networks. This is mainly due to the technology restrictions imposed by the lack of optical memories and the limitations of the E/O and O/E conversion devices.

REFERENCES

- [1] R. Braden et al., “Integrated services in the internet architecture: an overview,” RFC 1633, June 1994
- [2] S. Blake et al., “An architecture for differentiated services,” RFC 2475, Dec. 1998.
- [3] N. Golmie, T. D. Ndousse, and D.H. Su, “A differentiated optical services model for wdm networks,” *IEEE Communication Magazine*, pp. 68–73, February 2000.
- [4] Jukan, A., *QoS-based wavelength routing in multi-service WDM networks*, Springer, 2001.
- [5] F. Callegati, G. Corazza, and C. Raffaelli, “Exploitation of dwdm for optical packet switching with quality of service guarantees,” *IEEE Journal on Selected Areas in Communication*, vol. 20, pp. 190–201, January 2002.
- [6] N. M. Bhide, *Network protocols and algorithms for next generation optical wavelength division multiplexed networks*, Washington State University, 2001.
- [7] C. Yang, H. Mounir, and D.H.K. Tsang, “Proportional qos over wdm networks: Blocking probability,” *IEEE Symposium on Computers and Communications*, pp. 210–216, 2001.
- [8] M. Yoo and C. Qiao, “A new optical burst switching protocol for supporting quality of service,” in *SPIE Proceedings, All Optical Networking: Architecture, Control and Management Issue*, vol. 3531, pp. 396–405, November 1998.
- [9] Y. Xoing, M. Vandenhoute, and C. Cankaya, “Control architecture in optical burst-switched wdm networks,” *IEEE Journal on selected areas in communications*, vol. 18, pp. 1838–1851, October 2000.
- [10] J. Turner, “Terabit burst switching,” *Journal of High Speed Networks*, vol. 8, pp. 3–16, 1999.
- [11] M. Yang, S.Q. Zheng, and D. Verchere, “A qos supporting scheduling algorithm for optical burst switching in dwdm networks,” *Proceedings of GLOBECOM 2001*, pp. 86–91, 2001.
- [12] C. Murthy and M. Gurusamy, *WDM Optical Networks: Concepts, Design, and Algorithms*, Prentice Hall, 2002.
- [13] J.M.H. and H.T. Muftah, *Technologies and Architectures for Scalable Dynamic Dense WDM Networks*, *IEEE Communication Magazine* Feb. 2000.

- [14] C. Qiao and M. Yoo, "Choices, Features, and Issues in Optical Burst Switching," *Optical Network Magazine*, vol. 1, no. 2, pp. 36-44, April 2000.
- [15] T. Khattab, A. Mohamed, A. Kaheel, H. Alnuweiri, "Optical Packet Switching with Packet Aggregation," in *Proceedings of SoftCOM 2002*, Oct. 2002.

BIOGRAPHIES

Ayman Kaheel (aymank@ece.ubc.ca) received his B.S. degree in electrical engineering from Cairo University, Egypt in 1995, and his M.S. in computer science from the University of Louisville, Kentucky in 1998. He is currently a Ph.D. candidate in the Department of Electrical and Computer Engineering at the University of British Columbia. His research interests focus on issues in the WDM optical networks, including optical burst switching, QoS support and optical router architecture.

Tamer Khattab (tkhattab@ece.ubc.ca) received his B.Sc. and M.Sc. in Electrical Engineering from Cairo University (Egypt) in 1993 and 2000, respectively. He is currently studying towards his Ph.D. in Electrical Engineering at the University of British Columbia (Canada). Since 1994 Eng. Khattab has been working as a Research and Teaching Assistant at the National Telecommunication Institute (Egypt). Also, during the period 1994–1999 Eng. Khattab was involved with IBM Egypt in development projects for IBM corporate tools. Since Feb. 2000 Eng. Khattab has been involved with Alcatel Canada in the Network and Service Management R&D department.

Amr Mohamed (amrm@ece.ubc.ca) received his M.S. in electrical and computer engineering from the University of British Columbia, Vancouver, Canada, 2001. He received his B.S. from Cairo University in 1993. Since 1994 Amr has been working as a Research and Teaching Assistant at the National Telecommunication Institute, Cairo, Egypt. He currently works as a senior IT-Specialist in IBM Vancouver Innovation Centre, BC, Canada, since 1998. His main research interests include, QoS for WDM Networks, videoconferencing, video Broadcast/Multicast, and MPEG4 video streaming.

Hussein Alnuweiri (hussein@ece.ubc.ca) is an Associate Professor in the Department of Electrical and Computer Engineering at the University of British Columbia. He obtained

his Ph.D. in Computer Engineering from the University of Southern California, Los Angeles, in 1989. His main research interests are in the areas of advanced routing and traffic engineering algorithms and protocols in optical and packet networks including IP and MPLS/GMPLS router architecture, quality-of-service enforcement in IP networks, burst scheduling in IP-over-WDM networks and multimedia communications. In the past, he was an observing member at the ATM-Forum, and also a delegate of Canada to the ISO. Dr. Alnuweiri has authored or co-authored over 90 journal and conference papers, and holds two US patents.