An Absolute QoS Framework for Loss Guarantees in Optical Burst-Switched Networks

Minh Hoang Phùng, Kee Chaing Chua, Gurusamy Mohan, *Member, IEEE*, Mehul Motani, and David Tung Chong Wong, *Senior Member, IEEE*

***Abstract*—In order to meet the requirements of real-time applications, Optical Burst Switched backbone networks need to provide quantitative edge-to-edge loss guarantees to traffic flows. For this purpose, there have been several proposals based on the relative differentiation quality of service (QoS) model. However, this model has an inherent difficulty in communicating information about internal network states to the edge in a timely manner for making admission control decisions. In this paper, we propose an absolute QoS framework to overcome this difficulty. The key idea is to offer quantitative loss guarantees at each hop using a differentiation mechanism and an admission control mechanism. The edge-to-edge loss requirement is then translated into a series of small per-node loss probabilities that are allocated to the intermediate core nodes. The framework includes a preemptive differentiation scheme, a node-based admission control scheme and an edge-to-edge reservation scheme. The schemes are analyzed and evaluated through simulation. It is shown that the framework can effectively offer quantitative edge-to-edge loss guarantees under various traffic conditions.**

***Index Terms*—Absolute quality of service (QoS), computer networks, edge-to-edge QoS provisioning, optical burst switching, wavelength division multiplexing.**

I. Introduction

The bandwidth demand in the Internet is growing rapidly due to emerging multimedia applications such as video on demand, video conferencing and Internet telephony. According to TeleGeography, the average international Internet traffic growth rate is nearly 80% per annum between 2003 and 2005. To satisfy this bandwidth demand, Dense Wavelength Division Multiplexing (DWDM) has emerged as a core transmission technology for the next-generation Internet backbone. It provides enormous bandwidth with its ability to support hundreds of gigabit per second wavelength channels in a single fiber. However, this makes the mismatch between electronic processor speed and the optical transmission rate more acute. Recently, all-optical WDM networks have received much attention because they can bypass the opto-electronic-opto bottleneck in electronic routers. Optical Burst Switching (OBS) [1]–[3] is a promising optical networking technology that can potentially provide high wavelength utilization. In OBS, a burst is preceded in time by a header packet that is sent on a dedicated control wavelength. The time separation between the burst and its header packet (called the offset time) allows optical core nodes to process the header before the data burst arrives. This out-of-band signalling paradigm enables OBS to use flexible electronic control while still having an all-optical data path.

An important issue in the next-generation Optical Internet is how to support quality of service (QoS) at the optical layer due to the proliferation of multimedia and mission critical Internet applications that require stringent QoS guarantees. In OBS networks, there is only very limited buffering capability in the form of fiber delay lines (FDLs) (see [4] and references therein). Therefore, burst loss probability is a primary QoS metric of interest. There are currently two main QoS models in OBS: *relative QoS* and *absolute QoS*. In the relative QoS model, the QoS performance of a traffic class is defined relative to those of other classes. For instance, the loss rate of a higher priority class is guaranteed to be better or at least no worse than that of a lower priority class. To date, several relative QoS schemes have been proposed for OBS networks. They include offset-based schemes [5], [6], segmentation-based schemes [7], [8], dropping-based schemes [9]–[11] and preemptive schemes [12]–[15]. A common feature is that they require no admission control mechanism. On the other hand, the absolute QoS model gives quantitative QoS guarantees for each traffic class, e.g., having an end-to-end packet loss probability no greater than 1%. Absolute QoS schemes have been proposed for Optical Packet Switching (OPS) and OBS networks in [16]–[19].

From a user's point of view, the absolute QoS model is preferred. This is because an end user usually has a specific quantitative QoS requirement depending on the applications in use and would prefer to receive it independent of the network load. Although it is possible for an Internet Service Provider to emulate this absolute QoS behavior in a relative QoS environment by continuously adjusting the user's subscribed QoS class, it may not be feasible to do this for all users under heavy load conditions. Besides, the required performance and accounting management functions would be very complex. By placing absolute QoS mechanisms at core nodes, which have the most updated information about their own traffic conditions, the absolute QoS model can offer more accurate and robust quantitative QoS guarantees in simpler ways than the relative QoS model.

Paper approved by A. Pattavina, the Editor for Network Architecture of the IEEE Communications Society. Manuscript received July 6, 2005; revised January 20, 2006, and June 19, 2006. This paper was presented in part at the IASTED International Conference on Optical Communication Systems and Networks, Banff, Alberta, Canada, July 2004, and at the IEEE GLOBECOM, Dallas, TX, November 2004.

M. H. Phùng is with the Optical Networks Laboratory, Department of Electrical and Computer Engineering, National University of Singapore, Singapore 117576 (e-mail: engp1627@nus.edu.sg).

K. C. Chua, G. Mohan, and M. Motani are with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore 117576.

D. T. C. Wong is with the Institute for Infocomm Research, Singapore 119613.

Digital Object Identifier 10.1109/TCOMM.2007.898846

In this paper, we propose a novel absolute QoS framework to provide quantitative edge-to-edge (e2e) loss guarantees to burst flows in an OBS network. It defines a limited number of per-hop QoS classes and assigns each class a loss threshold. The framework employs a preemptive differentiation mechanism and an admission control mechanism at each output link of a core node. The differentiation mechanism shifts burst loss from classes in danger of breaching their thresholds to other classes while the admission control mechanism limits the link's offered load to a certain level. They work together to guarantee the loss threshold for each class at the link. Using these classes as building blocks, the framework employs a signalling and reservation mechanism to assign each burst flow to a certain class at each intermediate link such that the flow's e2e loss probability request is satisfied. The framework is shown to achieve reliable bounds on e2e loss probabilities under all traffic loads.

The rest of the paper is organized as follows. In Section II, we present an overview of the proposed framework. The preemptive differentiation scheme and the node-based admission control scheme are presented in detail in Section III. This section includes an analysis of the differentiation scheme. The e2e signalling and reservation scheme is presented next in Section IV. In Section V, we review two existing absolute QoS proposals and compare them with our framework. The differentiation scheme and the whole system are evaluated through simulation in Section VI. Finally, concluding remarks are given in Section VII.

II. Overview of the Proposed QoS Framework

Our proposed QoS framework is applied to a burst flow with a required maximum edge-to-edge (e2e) loss probability over a predetermined path. It attempts to reserve resources over the path so that the required e2e loss probability is guaranteed. If the reservation process fails for that particular path, a routing protocol with route-pinning capability in a larger traffic engineering framework may be used to select another path and the reservation process is applied again. As such, the framework should be considered as part of a final solution to the traffic engineering and QoS provisioning problem in OBS networks.

The key idea of the proposed framework is to define a limited number of per-hop absolute QoS classes[1] and enforce their loss thresholds at each link. The network then divides the required e2e loss probability of the flow into a series of small loss probabilities and maps them to the available thresholds at the intermediate links on the path. When each intermediate node guarantees that the actual loss probability at its link is below the allocated loss probability, the overall e2e loss guarantee is fulfilled.

The proposed QoS framework includes two mechanisms to enforce per-hop thresholds, i.e., a preemptive differentiation mechanism and an admission control mechanism. The differentiation mechanism allows bursts from classes that are in danger of breaching their thresholds to preempt bursts from other classes. Thus, burst loss is shifted among the classes based on the differences between the thresholds and the measured loss probabilities of the classes. The admission control mechanism limits the link's offered load to an acceptable level and thereby makes it feasible to keep the loss probabilities of all classes under their respective thresholds.

For the mapping of classes over an e2e path, we assume a label switching architecture such as Multi-Protocol Label Switching (MPLS) [20] to be present in the OBS network. In this architecture, each burst header carries a label to identify the burst flow or Label Switched Path (LSP) that it belongs to. When a header arrives at a core node, the node uses the header's label to look up the associated routing and QoS information from its Label Information Base (LIB). The old label is also swapped with a new one. Label information is downloaded to the node in advance by a Label Distribution Protocol. Such label switching architecture enables an LSP to be mapped to different QoS classes at different links.

An e2e signalling and reservation mechanism is responsible for probing the path of a new LSP and mapping it to a class at each intermediate link. When the LSP setup process begins, a reservation message that contains the requested bandwidth and the required e2e loss probability of the LSP is sent along the LSP's path toward the egress node. The message polls intermediate nodes on their available capacity and conveys the information to the egress node. Based on this information, the egress node decides whether the LSP's request can be accommodated. If the result is positive, an array of QoS classes whose elements correspond to the links along the path is allocated to the LSP. The class allocation is calculated such that the resulting e2e loss probability is less than that required by the LSP. It is then signalled to the intermediate core nodes using a returned acknowledgement message.

Finally, existing LSPs are policed for conformance to their reservation at ingress nodes. When the traffic of an LSP exceeds its reserved traffic profile, its generated bursts are marked as out of profile. Such bursts receive best-effort service inside the network.

III. Preemptive Differentiation Scheme

A. Description

In this section, we describe a preemptive differentiation scheme for absolute QoS that we proposed in [18]. To quantify the risk of breaching the threshold of a local QoS class, we introduce a metric called the *distance to threshold*, which is defined as the difference between the predefined loss threshold and the measured loss probability. A differentiation scheme in our absolute QoS framework must meet the following requirements:

- *Inter-class requirement*: It must ensure that as the offered load to a link increases, the distances to thresholds of all classes present at the link converge to zero. This implies that burst loss from classes that are in danger of breaching their thresholds is shifted to other classes by the differentiation scheme.
- *Intra-class requirement*: It must ensure that bursts belonging to the same class experience the same loss probability at a particular link regardless of their offsets and burst lengths. In OBS networks, it is well-known that burst lengths and offsets have significant impacts on burst loss probability. Hence, without intervention from the

[1] In the rest of this paper, the term "class" refers to per-hop QoS class unless otherwise specified.

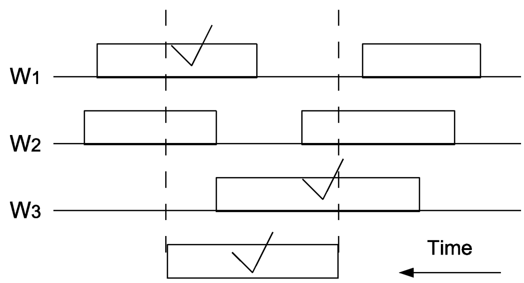

Fig. 1. Construction of the contention list.

differentiation scheme, some flows with unfavorable burst characteristics may experience loss probabilities above the threshold even though the overall loss probability of the class is still below the threshold.

Having set out the requirements and their rationales, we will now describe the proposed differentiation scheme. For scalability reasons, the scheme uses class-level differentiation. The scheme only requires a core node to keep per-class information, which includes the predefined loss threshold, the amount of admitted traffic and the current average loss probability. The average loss probability is continuously updated using an exponentially weighted averaging algorithm.

The scheme works as follows. When a burst header arrives at a node and fails to reserve an output wavelength, the node constructs a *contention list* that contains the incoming burst reservation and scheduled burst reservations that overlap (or contend) with the incoming one. Only one scheduled reservation on each wavelength is included if its preemption helps to schedule the new reservation. This is to prevent the case wherein several bursts may be preempted by one single higher priority burst and improve bandwidth utilization. The process is illustrated in Fig. 1 where the ticked reservations on wavelengths W_1, W_2, and W_3 are included in the contention list. The node then selects one reservation from the list to drop according to some criteria described later. If the dropped reservation is a scheduled one then the incoming reservation will be scheduled in its place. In that case, we say that the incoming reservation *preempts* the scheduled reservation.

When preemption happens, a special NOTIFY packet will be immediately generated and sent on the control channel to the downstream nodes to inform them of the preemption. The downstream nodes then remove the burst reservation corresponding to the preempted burst. Although one NOTIFY packet is required for every preemption, the number of preemptions is at most equal to the number of lost bursts. In other words, the rate of NOTIFY packets is bounded by the burst loss rate, which is usually kept very small in operational networks. Therefore, the additional overhead by the transmission of NOTIFY packets is not significant.

There are two criteria for selecting a burst reservation from the contention list to drop. The first criterion is that the selected reservation belongs to the class with the largest distance to threshold in the contention list. This criterion ensures that all the distances to thresholds of the classes present at the node are kept equal, thereby satisfying the first requirement above. The second criterion is applied when there are more than one reservation belonging to the class with the largest distance to threshold. In that case, only one of them is selected for dropping. Let the length of the ith reservation be l_i, $(1 \le i \le N)$, where N is the number of reservations belonging to the class with the largest distance to threshold in the contention list. The probability of it being dropped is

$$p_i^d = \frac{\frac{1}{l_i}}{\sum_{j=1}^{N} \frac{1}{l_j}}. \quad (1)$$

This is because the probability that a reservation is involved in a contention is roughly proportional to its length, assuming Poisson burst arrivals. So p_i^d is explicitly formulated to compensate for that burst length selection effect. In addition, the selection is independent of burst offsets. That is, although a large-offset burst is less likely to encounter contention when its header first arrives, it is as likely to be preempted as other bursts in subsequent contention with shorter-offset bursts. Therefore, the second requirement is achieved.

The above description assumes that no FDL buffer is present. It can be trivially extended to work with FDL buffers by repeating the preemption procedure for each FDL and the new reservation interval.

B. Analysis

In this section, we analyze the overall loss probability for the preemptive differentiation scheme. Both the lower and upper bounds and an approximate formula for the loss probability will be derived. Depending on the application's requirement, one can choose the most suitable formula to use.

The following assumptions are used in the analysis. Firstly, for the sake of tractability, only one QoS class is assumed to be active, i.e., having traffic. The simulation in Fig. 3 indicates that the results obtained are also applicable to the case with multiple classes. Secondly, burst arrivals follow a Poisson process with mean rate λ. A number of studies [21] indicate that this is the case for traffic in core networks at very short time scale. Thirdly, the incoming traffic consists of a number of traffic components with the ith component having a constant burst length $1/\mu_i$ and arrival rate λ_i. This assumption results from the fact that size-triggered burst assembly is a popular method to assemble bursts. When the triggering size is large compared to the average packet size, which is typically the case, this method produces burst lengths with a very narrow dynamic range, which can be considered constant. Finally, we assume that no FDL buffer is present and the offset difference among incoming bursts is minimal.

The lower bound on loss probability is easily derived by observing that preemption itself does not change the total number of lost bursts in the system. Thus, it is determined using the following Erlang's loss formula for an $M|G|k|k$ queueing model

$$P_l = B(k, \rho) = \frac{\frac{r^k}{k!}}{\sum_{i=0}^{k} \frac{r^i}{i!}} \quad (2)$$

where k is the number of wavelengths per link; ρ is the total offered load and $r = k\rho$. The Erlang B formula is used here due to its simplicity and ease of computation, especially when

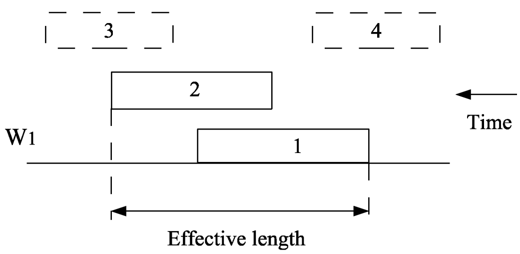

Fig. 2. Illustration of a preemption scenario.

k is large. It assumes that incoming bursts arrive from an infinite number of wavelengths. Such assumption is reasonably approximated when the number of wavelengths per link k is large, which is the case for core networks. When k is small, a more accurate model such as the Engset traffic model [22] should be used. However, the latter model does not scale well for large k's.

Although preemption does not directly affect the number of lost bursts, it affects the probability that a burst whose header arrives later is successfully scheduled. Depending on the reservation intervals of later bursts, the preemption may have detrimental or beneficial effects. Consider a preemption scenario as illustrated in Fig. 2 where burst 1 is preempted by burst 2. Let bursts 3 and 4 be two bursts whose headers arrive after the preemption. For burst 3, the preemption is detrimental because had there been no preemption, burst 3 would be successfully scheduled. On the other hand, the preemption is beneficial to burst 4. However, for that to happen, burst 4 has to have a considerably shorter offset than other bursts, which is unlikely due to our assumption that the offset difference among bursts is minimal. For other preemption scenarios, it can also be demonstrated that a considerable offset difference is required for a preemption to have beneficial effects. Therefore, it can be argued that preemption generally worsens the schedulability of later bursts.

To quantify that effect, we observe that from the perspective of burst 3, the preemption is equivalent to dropping burst 2 and extending the effective length of burst 1 as in Fig. 2. Therefore, it increases the time that the system spends with all k wavelengths occupied. The upper bound on burst loss probability is derived by assuming that the loss probability is also increased by the same proportion. Let $\delta = (l' - l)/l$ where l' is the new effective length and l is the actual length of the preempted burst. The upper bound on loss probability is then given as

$$P_u = \begin{cases} (1+\delta\rho)B(k,\rho), & \text{if } B(k,\rho) < \frac{1}{1+\delta\rho} \\ 1, & \text{otherwise.} \end{cases} \quad (3)$$

An approximate formula for the loss probability can be derived based on (3) by observing that the increase in effective length of a preempted burst will increase the overall loss probability only if another incoming burst contends with it again during the extended duration. The probability that this does not happen is

$$p = \sum_{i=0}^{\infty} \frac{e^{-\delta r}(\delta r)^i}{i!} \left(\frac{k}{k+1}\right)^i = e^{-\frac{\delta r}{k+1}}. \quad (4)$$

From (3) and (4), the loss probability is given as

$$P = P_u - e^{-\frac{\delta r}{k+1}} \delta\rho B(k,\rho). \quad (5)$$

We will now derive δ. Suppose the incoming traffic has N_c traffic components with N_c different burst lengths. Let a and b denote the component indices of the incoming burst and the preempted burst, respectively. The probability of a particular combination (a, b) is given by the formula

$$P(a,b) = \frac{\lambda_a}{\sum_{i=1}^{N_c}\lambda_i} \cdot \frac{\rho_b}{\sum_{j=1}^{N_c}\rho_j} \cdot \frac{\mu_b}{\sum_{k=1}^{N_c}\mu_k} \quad (1 \le a, b \le N_c). \quad (6)$$

The first and second factors are the probabilities that an incoming burst and a scheduled burst belong to components a and b, respectively. The third factor accounts for the length selective mechanism of the preemption scheme. For a preemption situation (a, b), the effective length is increased by $(1/\mu_a) - (1/2\mu_b)$. Therefore, it follows that

$$\delta = \sum_{a=1}^{N_c}\sum_{b=1}^{N_c} P(a,b)\left(\frac{\mu_b}{\mu_a} - \frac{1}{2}\right). \quad (7)$$

C. Local Admission Control at a Link

Since the distances to thresholds of the classes at a node are kept equal by the differentiation scheme, the admission control routine only needs to keep the average of these greater than zero. In other words, it needs to keep the overall loss probability smaller than the weighted average threshold. Suppose there are M QoS classes at the node and let T_i and B_i be the predefined threshold and the total reserved bandwidth of the ith class, respectively. The weighted average threshold is calculated as

$$T = \frac{\sum_{i=1}^{M} T_i B_i}{\sum_{j=1}^{M} B_j}. \quad (8)$$

The overall loss probability P can be calculated using (5) in the previous section. Alternatively, the upper bound given by (3) may be used for better protection against threshold violation. In case that the analytical formulas cannot be used, e.g., due to non-Poisson traffic, an empirical graph of the overall loss probability versus the total offered load may be used.

A reservation request will contain the amount of bandwidth to be reserved b_0 and the QoS class c to accommodate b_0 in. When a request arrives, the admission control routine substitutes B_c with $B'_c = B_c + b_0$ and recalculates the weighted average threshold T' and the overall loss probability P' as above. If $P' \le T'$, it implies that new overall loss probability P' is still within the total loss "allowance" indicated by T'. Therefore, the request is admitted. Otherwise, if $P' > T'$, the request is rejected.

D. Per-Hop QoS Class Definition

Since per-hop QoS classes are used as building blocks by the network to construct quantitative e2e loss guarantees, their definition is an important part of configuring the system. Usually, the number of classes M, which is directly related to the complexity of a core node's QoS differentiation block, is fixed. Hence, in this process, one only decides on where to place the available thresholds, namely the lowest and highest loss thresholds T_l and T_h and those between them.

Consider an OBS network in which LSPs have a maximum path length of H hops and a required e2e loss guarantee between P_l and P_h (not counting best-effort and out-of-profile traffic). The case requiring the lowest loss threshold T_l occurs when an LSP over the longest H-hop path requires P_l. Thus, T_l can be calculated as follows:

$$T_l = 1 - (1 - P_l)^{1/H}.$$

Similarly, the highest threshold is $T_h = P_h$ for the case when a one-hop LSP requires P_h.

When considering how to place the remaining thresholds between T_l and T_h, it is noted that since the potential required e2e loss probability P_0 is continuous and the threshold values are discrete, the e2e loss bound $P_{\rm e2e}$ offered by the network will almost always be more stringent than P_0. This "discretization error" reduces the maximum amount of traffic that can be admitted. Therefore, the thresholds need to be spaced so that this discretization error is minimized. A simple and effective way to do this is to distribute the thresholds evenly on the logarithmic scale. That is, they are assigned the values $T_l, \gamma T_l, \gamma^2 T_l, \ldots, \gamma^{M-1} T_l$, where $\gamma = (T_h/T_l)^{1/(M-1)}$.

IV. Edge-to-Edge Signalling and Reservation

A. Overview

Edge-to-edge signalling and reservation mechanisms, as the name implies, are responsible for coordinating the QoS reservation setup and teardown for LSPs over the e2e paths. During the reservation process of an LSP, the signalling mechanism polls all the intermediate core nodes about the remaining capacity on the output links and conveys the information to the egress node. Using that information as the input, the egress node produces a class allocation that maps the LSP to an appropriate class for each link on the path. The signalling mechanism then distributes the class allocation to the core nodes. As a simple illustration, let us consider an LSP with an e2e loss requirement of 5% that needs to be established over a four-hop path and the second hop is near congestion. Suppose the lowest threshold is $T_l = 0.05\%$ and the ratio between two adjacent thresholds is $\gamma = 2$. The network allocates the LSP a threshold of $\gamma^6 T_l = 3.2\%$ for the second hop and $\gamma^3 T_l = 0.4\%$ for the other hops to reflect the fact that the second node is congested. The resulting guaranteed upper bound on e2e threshold will roughly be 4.4%, satisfying the LSP's requirement.

The QoS requirements of an LSP consists of its minimum required bandwidth and its maximum e2e loss probability. As new IP flows join an LSP or existing IP flows terminate, a reservation or teardown process needs to be carried out for the LSP. The reservation scenarios for an LSP can be categorized as follows.

1) A new LSP is to be established with a specified minimum bandwidth requirement and a maximum e2e loss probability. This happens when some IP flow requests arrive at the ingress node and cannot be fitted into any of the existing LSPs.
2) An existing LSP needs to increase its reserved bandwidth by a specified amount. This happens when some incoming IP flows have e2e loss requirements compatible with that of the LSP.
3) An existing LSP needs to decrease its reserved bandwidth by a specified amount. This happens when some existing IP flows within the LSP terminate.
4) An existing LSP terminates because all of its existing IP flows terminate.

The detailed reservation process for the first scenario is as follows. The ingress node sends a reservation message towards the egress node over the path that the LSP will take. The message contains a requested bandwidth b_0 and a required e2e loss probability P_0. When a core node receives the message, its admission control routine checks each class using the method described in Section III-C to see if the requested bandwidth can be accommodated in that class. The check starts from the lowest index class, which corresponds to the lowest threshold, and moves up. The node stops at the first satisfactory class and records in the message the class index c and a parameter κ calculated as follows:

$$\kappa = \left\lfloor log_{\gamma}\left(\frac{T'}{P'}\right)\right\rfloor \tag{9}$$

where T' and P' are as described in Section III-C and γ is the ratio between the thresholds of two adjacent classes. These parameters will be used by the egress node for the final admission control and class allocation. The message is then passed downstream. The node also locks in the requested bandwidth by setting the total reserved bandwidth B_c of class c as $B_c(\text{new}) = B_c(\text{old}) + b_0$ so that the LSP will not be affected by later reservation messages. On the other hand, if all the classes have been checked unsuccessfully, the request is rejected and an error message is sent back to the ingress node. Upon receiving the error message, the upstream nodes release the bandwidth locked up earlier.

The final admission control decision for the LSP is made at the egress node. The received reservation message contains two arrays $\boldsymbol{c}$ and $\boldsymbol{\kappa}$ for the intermediate links of the path. Assuming burst blocking at each link is independent, the lowest possible e2e loss probability $P^0_{\rm e2e}$ given as

$$P^0_{\rm e2e} = 1 - \prod_{i=1}^{n}\left(1 - p^0_i\right) \tag{10}$$

where p^0_i is the threshold of the class offered by the ith node (as indicated in array $\boldsymbol{c}$). If $P^0_{\rm e2e} \le P_0$, the request is admitted. The egress node then allocates each core node one of the predefined classes in which to support the LSP such that

$$\begin{cases} p_i \ge p^0_i \\ P_{\rm e2e} = 1 - \prod\limits_{i=1}^{n}(1 - p_i) \le P_0 \end{cases} \tag{11}$$

where p_i is the threshold of the class allocated to the LSP at the ith node and $P_{\rm e2e}$ is the corresponding e2e loss probability. The class allocation algorithm will be described in the next section. This class allocation is signalled back to the intermediate core nodes using a returned acknowledgement message that contains the old index array $\boldsymbol{c}$ and an allocated index array $\boldsymbol{c_a}$. Upon receiving the acknowledgement message, a core node moves the reserved bandwidth of the LSP from class c to class c_a. The new LSP is allowed to start only after the ingress node has received the successful acknowledgement message. If $P^0_{\rm e2e} > P_0$, the request is rejected and an error message is sent back to the ingress node. The intermediate core nodes will release the locked bandwidth upon receiving the error message.

A brief examination of the overhead of the signalling mechanism is in order. We observe that the signalling message always contains two arrays: $\boldsymbol{c}$ and $\boldsymbol{\kappa}$ in the forward trip and $\boldsymbol{c}$ and $\boldsymbol{c_a}$ in the return trip. Among them, only $\boldsymbol{\kappa}$ contains floating point numbers. The remaining are integer arrays. The number of elements of each array is equal to the number of intermediate core nodes on the path of the LSP. Besides, other variables contained in each message are the requested bandwidth b_0, the required e2e loss probability P_0, the label of the LSP and a flag to indicate various reservation statuses. Suppose $n = 10$ and the sizes of an integer and a floating point number are 4 byte and 16 bytes, respectively, the maximum size of each signalling message is $10 \times (4 + 16) + 16 + 16 + 4 + 1 = 237$ bytes. Since only two messages are required to set up each LSP, the signalling overhead is not significant.

The reservation process for the second scenario is relatively simpler. In this case, the ingress node sends out a reservation message containing the requested bandwidth b_0 and the LSP's label. Since there is already a QoS class associated with the LSP at each of the core nodes, a core node on the path only needs to check if b_0 can be supported in the registered class. If the outcome is positive, the node locks in b_0 and passes the reservation message on. Otherwise, an error message is sent back and the upstream nodes release the bandwidth locked previously. If the reservation message reaches the egress node, a successful acknowledgement message is returned to the ingress node and the LSP is allowed to increase its operating bandwidth.

In the last two scenarios, the reservation processes are similar. The ingress node sends out a message carrying the amount of bandwidth with a flag to indicate that it is to be released and the LSP's label. The released bandwidth is equal to the reserved bandwidth of the LSP if the LSP terminates. At intermediate core nodes, the total reserved bandwidth of the class associated with the LSP is decreased by that amount. No admission control check is necessary. Since the core nodes do not keep track of bandwidth reservation by individual LSPs, the processing at core nodes is identical for both scenarios. It should be noted that when an LSP terminates, there is a separate signalling process to remove the LSP's information from core nodes' LIBs. However, it is not considered part of our QoS framework.

B. *Dynamic Class Allocation*

When the egress node has determined that the LSP request is admissible using (10), it uses a dynamic class allocation algorithm to find the bottleneck link and allocate the class with the highest possible threshold to it while still satisfying (11). This shifts some of the loss guarantee burden from the bottleneck link to other lightly loaded links. Since the remaining capacity of the path is determined by the bottleneck link, the algorithm will maximize the path's remaining capacity and allows more future QoS traffic to be admitted.

For this purpose, the egress node has at its disposal two arrays $\boldsymbol{c}$ and $\boldsymbol{\kappa}$ recorded in the reservation message by the core nodes. As described in the last section, $\boldsymbol{c}[i]$ is the lowest index class (with the lowest threshold) that can accommodate the requested bandwidth at the ith node. As long as $\boldsymbol{c}[i] > 0$, it is an accurate indicator of how much capacity is available at link i since it cannot be decreased further without making P' exceed T'. However, when $\boldsymbol{c}[i] = 0$, we do not know how much lower P' is compared to T' based on $\boldsymbol{c}[i]$ alone. Hence, $\boldsymbol{\kappa}[i]$ given by (9) is introduced to enable the core node to convey that information to the egress node. We observe that when $\boldsymbol{c}[i] > 0$, $\gamma^{-1}T' < P' \leq T'$. Therefore, $\boldsymbol{\kappa}[i] > 0$ only if $\boldsymbol{c}[i] = 0$. Thus, $\boldsymbol{c} - \boldsymbol{\kappa}$ indicates the remaining capacity at the intermediate links in all cases. The higher $\boldsymbol{c}[i] - \boldsymbol{\kappa}[i]$, the lower the remaining capacity at link i and *vice versa*.

Based on the above observation, the class allocation algorithm is detailed in Algorithm 1. In executing this algorithm, negative class indices in $\boldsymbol{c_a}$ are counted as zero. In the first two lines, the algorithm sets $\boldsymbol{c_a}$ such that the maximum element is $M - 1$ and the differences among the elements are the same as in the array $\boldsymbol{c} - \boldsymbol{\kappa}$. Next, it repeatedly decrements all the elements of $\boldsymbol{c_a}$ until $P_{\rm e2e} \leq P_0$. Finally, the elements of $\boldsymbol{c_a}$ are incremented one by one until just before $P_{\rm e2e} > P_0$ in order to push $P_{\rm e2e}$ as close as possible to P_0 without exceeding it.

Algorithm 1: Class allocation algorithm

(1) $d_{\max} \leftarrow \mathrm{Max}(\boldsymbol{c} - \boldsymbol{\kappa})$

(2) $\boldsymbol{c_a} \leftarrow (\boldsymbol{c} - \boldsymbol{\kappa}) - d_{\max} + M - 1$

(3) **while** $P_{\rm e2e} > P_0$

(4) $\quad \boldsymbol{c_a} \leftarrow \boldsymbol{c_a} - 1$

(5) Increment elements of $\boldsymbol{c_a}$ one by one until just before $P_{\rm e2e} > P_0$

As an illustrative example, consider an OBS network that has eight predefined QoS classes with indices $\{0, 1, \ldots, 7\}$. The lowest threshold is $T_l = 0.05\%$ and the ratio between two adjacent thresholds is $\gamma = 2$. These settings provide a sufficiently small lowest threshold T_l and sufficiently fine threshold increment to satisfy most typical e2e loss requirements. An LSP with an e2e loss requirement of 1% is to be set up over a three-hop path. Its required bandwidth is assumed to be very small compared to the link capacity. The utilization levels at the intermediate links are $\{0.3, 0.6, 0.35\}$. Suppose the received message at the egress node contains $\boldsymbol{c} = \{0, 0, 0\}$ and $\boldsymbol{\kappa} = \{50, 2, 35\}$. The values of $\boldsymbol{\kappa}$ indicate that although all links are relatively lightly loaded, the second link has least available capacity. Therefore, the class allocation algorithm should give it the highest threshold. Going through the algorithm, we have $\boldsymbol{c_a} = \{-41, 7, -26\}$ on line (3) and $\boldsymbol{c_a} = \{-44, 7, -29\}$ on line (5). The final result is $\boldsymbol{c_a} = \{1, 4, 1\}$ corresponding to thresholds of $\{0.1\%, 0.8\%, 0.1\%\}$. It shows that the algorithm successfully allocates the maximum possible class index to the bottleneck node.

Most of the computations in the above algorithm are array manipulation and the calculation of $P_{\rm e2e}$ according to (10). Since their computational complexity is determined by number of intermediate core nodes, which is small, the computational complexity of the algorithm is not significant.

V. Previous Work and Discussion

To date, two absolute QoS schemes designed to provide quantitative e2e loss probabilities in OPS and OBS networks have been proposed in [16], [17]. In these schemes, e2e QoS classes with absolute loss thresholds are defined first. New flows will choose a suitable e2e class in which to transmit. For each e2e threshold $P_0^{\rm end}$, core nodes are given equal per-hop thresholds, which are calculated as

$$p_0^i = 1 - \left(1 - P_0^{\rm end}\right)^{1/H} \qquad (12)$$

where H is the maximum hop length of a network path. The scheme in [17] further improves on this by grouping possible network paths into clusters based on their hop lengths and uses the maximum hop length H_c in a cluster instead of H in (12). A core node uses one of the differentiation mechanisms described below to shift burst loss among the per-hop classes to attempt to keep all classes below their thresholds. No admission control mechanism is provided in either scheme except that paper [17] briefly mentions that it is to be done only at edge nodes.

Various absolute QoS differentiation mechanisms are proposed in the above two papers. In [16], only one traffic class requiring absolute QoS is considered. The threshold of the class is maintained by adjusting the probability that a burst belonging to that class is allowed to preempt non-priority bursts. In [17], three absolute differentiation schemes are presented. The early dropping scheme uses intentional dropping of low priority bursts as the means to shift burst loss away from higher priority classes when needed. In the wavelength grouping scheme, each per-hop class is allocated a certain number of wavelengths on which to schedule its bursts. The number of allocated wavelengths of a class is based on the loss threshold and the offered load of that class. Finally, in the integrated scheme, low priority bursts that would have been dropped under the early dropping scheme are scheduled on a limited number of wavelengths while other bursts are scheduled on all wavelengths.

Our proposed absolute QoS framework differs from the above two schemes in several important ways. Firstly, unlike the above two schemes, our framework provides elaborate admission control mechanisms for robust e2e loss guarantees under all network loads. Secondly, our proposal defines per-hop QoS classes first and use these as building blocks to provide e2e loss guarantees. Together with a label switching architecture, this feature enables the network to tailor the threshold allocation based on traffic loading at intermediate links as described in Section IV-B instead of uniform threshold allocation. Thus, less stringent thresholds can be allocated to bottleneck links, which allows more QoS traffic to be admitted. Moreover, it also reduces the number of thresholds required at a core node compared to the path clustering proposal. With path clustering, the number of thresholds required is $M = N \times Cl$, where N is the number of e2e loss guarantees and Cl is the number of clusters. On the other hand, a small number of predefined per-hop class in our proposal can be arbitrarily combined to generate a large number of e2e loss guarantees. Thirdly, our proposed absolute differentiation algorithm includes a mechanism to compensate for the burst length selection effect in order to achieve uniform loss probability within a class, which ensures that the loss probabilities of individual LSPs are the same as the overall loss probability of the class. This is missing from the existing proposals. Finally, admission control is done at core nodes in our proposal. This solves issues concerning inaccurate traffic information, which are inherent in edge-based admission control schemes, and minimizes the signalling overhead between edge and core nodes.

VI. Experimental Study

A. Absolute QoS Differentiation

In this section, we evaluate the proposed absolute differentiation algorithm against the two criteria in Section III-A and verify the analytical results obtained in Section III-B through simulation at the node level.

The node in the simulation has an output link with 64 data wavelengths, each having a transmission rate of 10 Gb/s. We assume that the node has full wavelength conversion capability and no buffering. Bursts arrive at the link according to a Poisson process with rate λ. This Poisson traffic assumption is valid for core networks due to the aggregation effect of a large number of flows per link. The burst lengths are generated by a size-limited burst assembly algorithm with a size limit of 50 kB. Thus, the generated bursts have lengths between 50 kB and 51.5 kB, or between 40 μs and 41.2 μs.

In the first experiment, we wish to verify the accuracy of the analysis in Section III-B. It is expected that the simulation and analytical results are the same. For this purpose, we plot the overall loss probabilities of traffic with one QoS class, traffic with seven QoS classes and the analytical value against the overall loading. In the case with seven classes, the classes are configured with thresholds ranging from $T_l = 0.0005$ to $T_h = 0.032$ and the ratio between two adjacent thresholds is $\gamma = 2$. The traffic of the highest threshold class takes up 40% of the total traffic. For each of the remaining classes, their traffic takes up 10% of the total traffic.

From the plot in Fig. 3, we observe that all the three loss curves match one another very well. It shows that the analysis is accurate and its assumption is valid, i.e., the traffic mix does not affect the overall loss probabilities. The reason is that in our differentiation scheme, preemption potentially happens whenever there is a contention between bursts, regardless of whether they are of different classes or of the same class. Therefore, the number of lost bursts depends only on the number of burst contentions but not the traffic mix.

In the next experiment, the loss probabilities of individual classes are plotted against the overall loading in Fig. 4. For easy visualization, we use only two QoS classes. Class 0 has a threshold of 0.005 and takes up 20% of the overall traffic. Class 1 has a threshold of 0.01. We observe that as the loading increases, the loss probabilities of both classes approach their corresponding thresholds. It shows that the algorithm satisfies

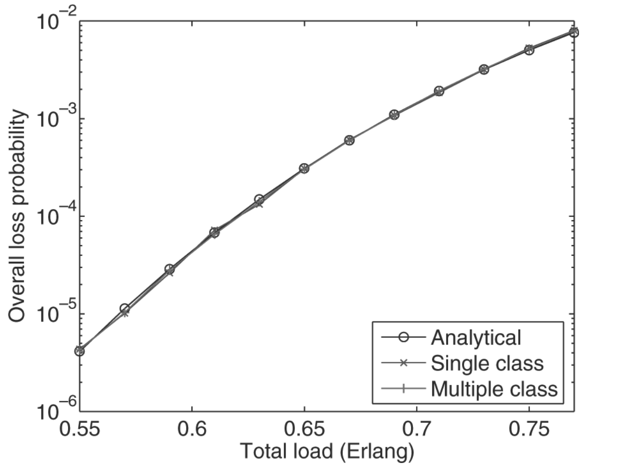

Fig. 3. Overall loss probabilities of various traffic scenarios versus overall loading.

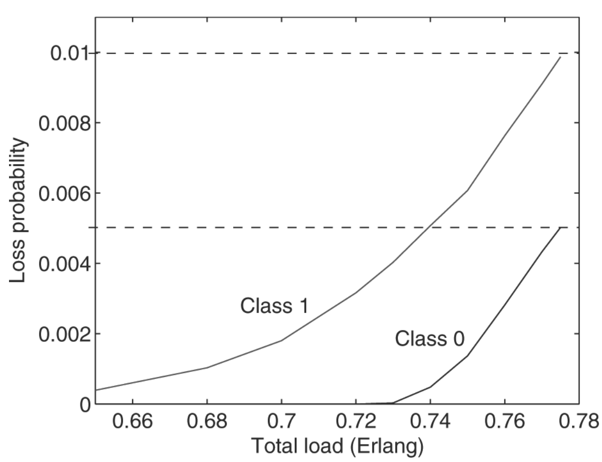

Fig. 4. Burst loss probabilities of individual classes versus overall loading.

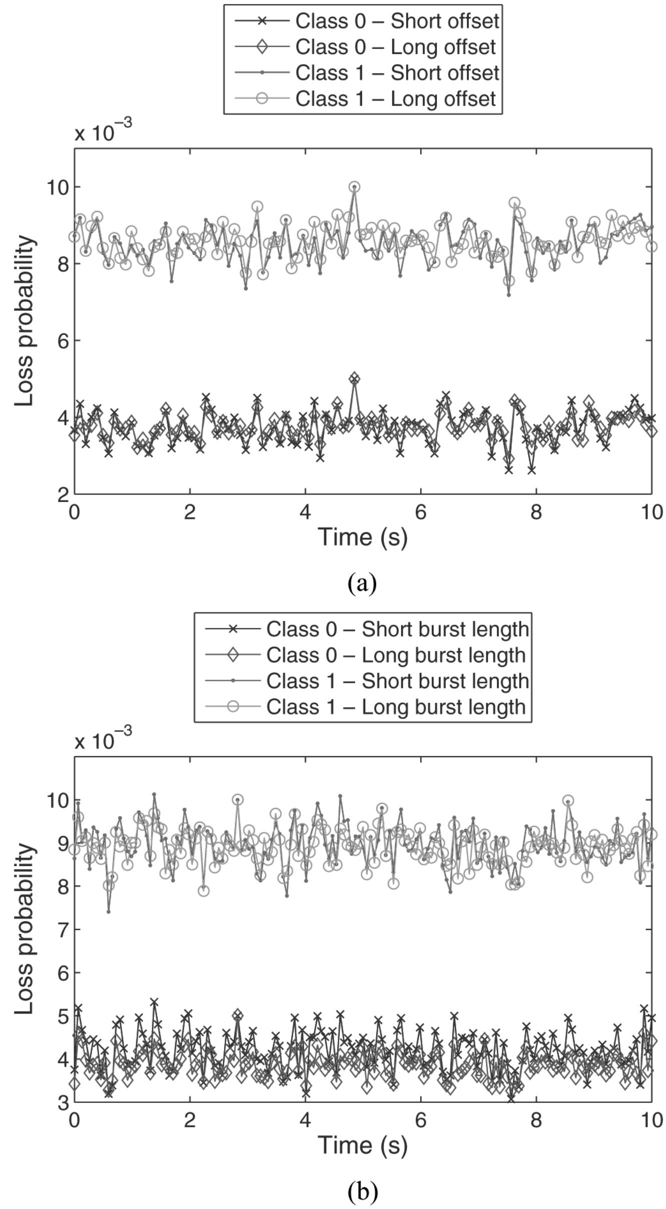

Fig. 5. Burst loss probability versus time for traffic components with different characteristics. (a) Different offsets. (b) Different burst lengths.

the first criterion set out in Section III-A. In addition, the distances to thresholds are always kept equal except when the loss probability of class 0 becomes zero. Although this feature is not required by the model, it is introduced to keep the minimum of the distances to thresholds as large as possible, thereby reducing the possibility that a sudden increase in incoming traffic could take the loss probability of a class beyond its threshold.

In Fig. 5, the loss performance of traffic components with different traffic characteristics within a class is investigated. The overall loading is 0.77 and the class configuration is the same as in the previous experiment. Each class has two traffic components in equal proportion. The plot in Fig. 5(a) is for the situation where the traffic components have different offsets. The offset difference is 40 μs, which is approximately one burst length. In Fig. 5(b), each class has two traffic components with different burst lengths. The size limits for the two components in the burst assembly algorithms are set at 50 kB and 100 kB, respectively. These settings would cause major differences in loss performance of the traffic components in a normal OBS system. Nevertheless, both figures show that the loss performance of different components within a class follow each other very closely despite the difference in their burst characteristics. It can be concluded that the proposed differentiation scheme can achieve uniform loss performance for individual flows within the same class as required by the second criterion in Section III-A.

B. Edge-to-Edge Reservation

In this section, we present the simulation results for the whole framework over an end-to-end path. The topology in Fig. 6, which is a simplified topology for the U.S. backbone network, is used. It consists of 24 nodes and 43 links. Fixed shortest path routing is used and the maximum path length is 6. For simplicity, the propagation delays between adjacent nodes are assumed to have a fixed value of 5 ms. The links are bi-directional, each implemented by two uni-directional links in opposite directions. Link parameters and burst characteristics are the same as in the previous section. Seven per-hop QoS classes are defined at each

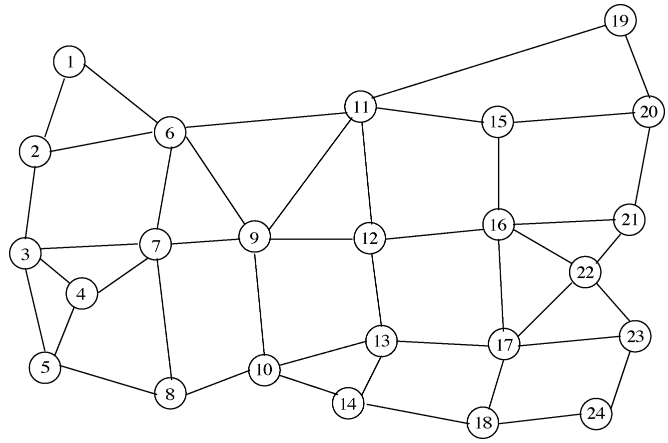

Fig. 6. Simulation topology.

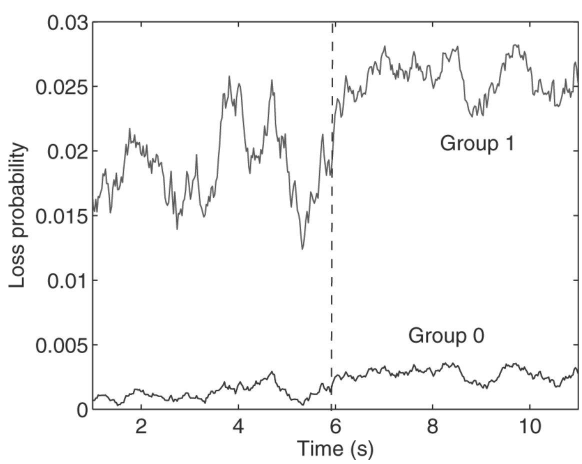

Fig. 7. E2e loss probability versus time.

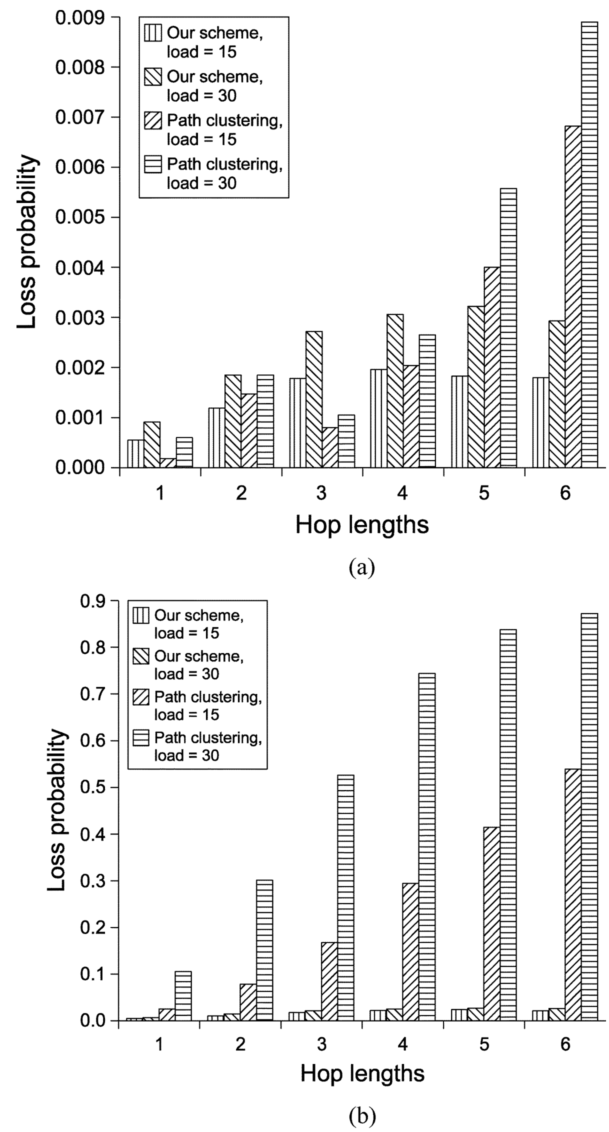

Fig. 8. Average e2e loss probability of LSPs with different hop lengths. (a) Group 0. (b) Group 1.

link with the lowest threshold $T_l = 0.0005$ and the ratio between two adjacent thresholds $\gamma = 2$.

New LSPs are generated at each node according to a Poisson process with rate λ_{LSP} and have exponentially distributed durations. For simplicity, we assume that LSPs do not change their bandwidth requirements. Two groups of LSPs are considered: group 0 and group 1 with required e2e loss probabilities of 0.01 and 0.05, respectively. A new LSP is destined to a random node and falls into one of the two groups with equal probability.

In the first experiment, we investigate the temporal loss behavior of the framework. To do this, we run the simulation for 11 s and monitor the e2e loss rate of traffic between node pair (1,24). The path between this node pair is 6 hops long, which is the longest path in the network, and it runs through the bottleneck link (9,10). The data in the first second is discarded. During the first 6 s, the total network load is 15 Erlang, which is equally distributed among all node pairs. After that, the offered load between node pair (1,24) is increased 10 folds. The loss rates of the two groups are plotted against time in Fig. 7. It is observed that the loss probabilities increase in response to the increase in the offered load at $t = 6$ s. Nevertheless, they are always kept below the respective thresholds. It shows that the reservation process is able to guarantee the loss bounds to admitted LSPs in real time regardless of the traffic load.

Another observation from Fig. 7 is that the maximum loss probabilities of the two traffic groups are 0.004 and 0.03, which are well below the required e2e loss probabilities. This is due to the fact that almost all of the burst loss on the path occurs at a single bottleneck link. Hence, the e2e loss probabilities are limited by the maximum thresholds that can be allocated to the bottleneck links. In this case, they are 0.004 and 0.032, respectively. If more per-hop classes are available, the gaps between adjacent thresholds will be reduced and the e2e loss probabilities can be pushed closer to the targets.

In Fig. 8, we plot the e2e loss probabilities of LSPs with different hop lengths and at two different loads of 15 and 30 Erlang. The same loss probabilities of the path clustering scheme proposed in [17] are also plotted for comparison. Since no admission control implementation is provided in the paper, we do not include any admission control mechanism for the path clustering scheme. The cluster combination $\{1,2\}\{3,4,5,6\}$ is used as it is the best performing one. It groups LSPs with one

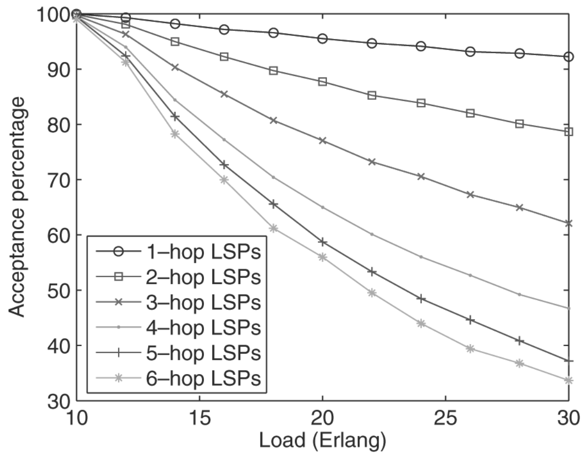

Fig. 9. Overall acceptance percentage versus load.

or two hop lengths into one cluster and all the remaining LSPs into the other cluster.

A number of observations can be made from Fig. 8. Firstly, we observe that the e2e loss probabilities of all LSP groups in our scheme are below their required e2e levels. This is true under both medium and heavy loading conditions. Secondly, the loss probabilities in our scheme increase from one-hop group to three-hop group but level off after that. The loss probability increase is due to the fact that burst traversing more hops will experience more loss. However, at a certain level, the effect of admission control dominates and the loss probabilities stay nearly constant. For the path clustering scheme, it is observed that it can keep the e2e loss probabilities for group 0 LSPs below the required level. However, it is achieved at great cost to group 1 LSPs, which experience very large loss probabilities. This happens because there is no admission control present. So core nodes must drop lower priority bursts excessively in order to keep the loss guarantees for high priority bursts. Another observation is that the loss probabilities of group 0 LSPs in the path clustering scheme vary significantly with hop lengths. This is because the scheme allocates the same per-hop threshold to LSPs within a cluster. Therefore, LSPs in a cluster with many different hop lengths such as $\{3, 4, 5, 6\}$ will experience significantly different e2e loss probabilities, some of which are far below the required level.

In the final experiment, the acceptance percentage of LSP groups with different hop lengths are plotted against the network loads in Fig. 9. It shows that the acceptance percentage of all LSP groups decrease with increasing load, which is expected. Among the groups, the longer the hop length, the worse the performance. There are two reasons for that. Firstly, the network must allocate more stringent per-hop thresholds to longer LSPs compared to shorter LSPs that have the same required e2e loss probability. Secondly, longer LSPs are more likely to encounter congestion on one of their intermediate links. This situation can be remedied by a fairness scheme that gives more favorable treatment to longer LSPs in the reservation process. However, that is beyond the scope of this paper.

VII. Conclusion

In this paper, we have proposed a novel absolute QoS framework that can offer quantitative e2e loss guarantees for individual LSPs in OBS networks. The framework consists of two parts. The first part includes a preemptive differentiation mechanism and a hop-based admission control mechanism. They enable core nodes to offer per-hop loss guarantees to individual LSPs. The second part includes edge-to-edge signalling and reservation mechanisms. For each new LSP, the framework divides its required e2e loss probability into small loss probabilities that are allocated to the intermediate links through the signalling and reservation process. When each core node fulfills its loss guarantee, the e2e loss guarantee is achieved. We analyze the differentiation scheme and verify it through simulation. The whole framework is also evaluated through network level simulation. It is shown to be able to reliably guarantee e2e loss probabilities under different network scenarios.

References

[1] C. Qiao and M. Yoo, "Optical Burst Switching—A new paradigm for an optical Internet," *J. High Speed Netw.*, vol. 8, no. 1, pp. 69–84, Jan. 1999.

[2] J. S. Turner, "Terabit burst switching," *J. High Speed Netw.*, vol. 8, no. 1, pp. 3–16, Jan. 1999.

[3] Y. Xiong, M. Vandenhoute, and H. C. Cankaya, "Control architecture in Optical Burst-Switched WDM networks," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 1838–1851, Oct. 2000.

[4] D. K. Hunter, M. C. Chia, and I. Andonovic, "Buffering in optical packet switches," *J. Lightw. Technol.*, vol. 16, no. 12, pp. 2081–2094, Dec. 1998.

[5] M. Yoo, C. Qiao, and S. Dixit, "QoS performance of Optical Burst Switching in IP-over-WDM networks," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 2062–2071, Oct. 2000.

[6] J. Wan, X. Liu, X. Sun, and M. Zhang, "QoS in Optical Burst Switching networks," *Proc. SPIE*, vol. 4604, pp. 123–129, Nov. 2001.

[7] V. M. Vokkarane and J. Jue, "Prioritized burst segmentation and composite burst-assembly techniques for QoS support in Optical Burst-Switched networks," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 7, pp. 1198–1209, Sep. 2003.

[8] Z. Zhang, F. Cheng, J. Wang, J. Luo, Q. Zeng, and X. Xuan, "A new burst assembly and dropping scheme for service differentiation in Optical Burst Switched networks," *Proc. SPIE*, vol. 5281, pp. 236–245, Nov. 2003.

[9] Y. Chen, M. Hamdi, and D. Tsang, "Proportional QoS over OBS networks," in *Proc. IEEE GLOBECOM*, 2001, pp. 1510–1514.

[10] M. Yang, S. Zheng, and D. Verchere, "A QoS supporting scheduling algorithm for Optical Burst Switching DWDM networks," in *Proc. IEEE GLOBECOM*, 2001, pp. 86–91.

[11] D. Q. Liu and M. T. Liu, "Burst scheduling for differentiated services in Optical Burst Switching WDM networks," *Int. J. Commun. Syst.*, vol. 17, no. 2, pp. 127–140, 2004.

[12] H. C. Cankaya, S. Charcranoon, and T. S. El-Bawab, "A preemptive scheduling technique for OBS networks with service differentiation," in *Proc. IEEE GLOBECOM*, 2003, p. 2704.

[13] W. Liao and C.-H. Loi, "Providing service differentiation for Optical-Burst-Switched networks," *J. Lightw. Technol.*, vol. 22, no. 7, pp. 1651–1660, Jul. 2004.

[14] F. Farahmand and J. P. Jue, "Supporting QoS with look-ahead window contention resolution in Optical Burst Switched network," in *Proc. IEEE GLOBECOM*, 2003, pp. 2699–2703.

[15] L. Yang, Y. Jiang, and S. Jiang, "A probabilistic preemptive scheme for providing service differentiation in OBS networks," in *Proc. IEEE GLOBECOM*, 2003, pp. 2689–2693.

[16] H. Overby and N. Stol, "Quality of service in asynchronous bufferless Optical Packet Switched networks," *Telecommun. Syst.*, vol. 27, no. 2–4, pp. 151–179, 2004.

[17] Q. Zhang, V. M. Vokkarane, J. Jue, and B. Chen, "Absolute QoS differentiation in Optical Burst-Switched networks," *IEEE J. Sel. Areas Commun.*, vol. 22, no. 9, pp. 1781–1795, Nov. 2004.

[18] M. H. Phùng, K. C. Chua, G. Mohan, M. Motani, and T. C. Wong, "A preemptive differentiation scheme for absolute loss guarantees in OBS networks," presented at the IASTED Int. Conf. Opt. Commun. Syst. Netw., Banff, AB, Canada, Jul. 2004.

[19] M. H. Phùng, K. C. Chua, G. Mohan, M. Motani, and T. C. Wong, "Absolute QoS signalling and reservation in Optical Burst-Switched networks," in *Proc. IEEE GLOBECOM*, Nov. 2004, pp. 2009–2013.

[20] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol label switching architecture," RFC 3031, Jan. 2001.

[21] T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido, "A nonstationary Poisson view of Internet traffic," in *Proc. IEEE INFOCOM*, 2004, vol. 3, pp. 1558–1569.

[22] H. Overby, "Performance modelling of optical packet switched networks with the Engset traffic model," *OSA Optics Express*, vol. 13, no. 5, pp. 1685–1695, Mar. 2005.

Minh Hoang Phùng received the B.Eng degree in telecommunications (first-class honors) from the University of Sydney, Australia, in 2000, and the Ph.D. degree in electrical engineering from the National University of Singapore (NUS) in 2006.

He is currently a Research Fellow in the Department of Electrical and Computer Engineering, NUS. His research interests focus on networking issues in WDM optical networks.

Kee Chaing Chua received the Ph.D. degree in electrical engineering from the University of Auckland, New Zealand, in 1990.

He joined the National University of Singapore (NUS) as a Lecturer in 1990 and is now a Professor in the Department of Electrical and Computer Engineering. He served as the Faculty of Engineering's Vice Dean for Research from June 2003 to March 2006. From 1995 to 2000, he was seconded to the Center for Wireless Communications (now part of the Institute for Infocomm Research), a national telecommunication R&D center funded by the Singapore Agency for Science, Technology and Research as its Deputy Director. From 2001 to 2003, he was on leave of absence from NUS to work at Siemens Singapore, where he was the founding head of the Mobile Core R&D Department funded by Siemens' ICM Group. Since March 2006, he has been seconded to the National Research Foundation as a Director. He has carried out research in various areas of communication networks, and has published more than 180 papers in these areas in international refereed journals and conferences. His current research interests are in wireless networks (in particular, wireless sensor networks) and optical burst switched networks.

Dr. Chua has been an active member of the IEEE, and is a recipient of an IEEE 3rd Millennium Medal.

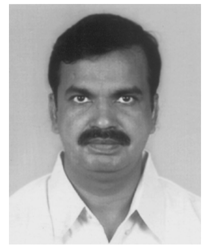

Gurusamy Mohan (M'00) received the Ph.D. degree in computer science and engineering from the Indian Institute of Technology (IIT), Madras, in 2000.

He joined the National University of Singapore in June 2000, where he is currently an Assistant Professor in the Department of Electrical and Computer Engineering. He has held a visiting position at Iowa State University, Ames, during January–June 1999. His current research interests are in high-speed multi-wavelength optical circuit and burst switching networks, IP-MPLS, and sensor networks. He is a coauthor of the textbook *WDM Optical Networks: Concepts, Design, and Algorithms* (Prentice Hall PTR, 2001).

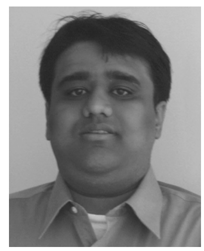

Mehul Motani (S'92–M'98) received the Ph.D. degree in electrical and computer engineering from Cornell University, Ithaca, NY, in 2000.

Since August 2000, he has been an Assistant Professor in the Department of Electrical and Computer Engineering at the National University of Singapore. From 2000 to 2003, he was also with the Institute for Infocomm Research, Singapore, as a Research Scientist. From 1992 to 1996, he was a Member of Technical Staff at Lockheed Martin, Ocean Radar and Sensor Systems, Syracuse, NY. His current research interests are in information theory, wireless networks including the design of wireless ad-hoc and sensor network systems, and optical networks.

Dr. Motani was awarded the Intel Foundation Fellowship in 2000 for his Ph.D. work on information theory and coding for CDMA systems. He is on the organizing committees for ISIT 2006 and 2007 and the technical program committees of MobiCom 2007, INFOCOM 2008, and several other conferences. He participates actively in IEEE and ACM and has served as the secretary of the IEEE Information Theory Society Board of Governors.

David Tung Chong Wong (S'90–M'95–SM'03) received the B.Eng. and M.Eng. degrees from National University of Singapore (NUS) in 1992 and 1994, respectively, and the Ph.D. degree from the University of Waterloo, Waterloo, ON, Canada, in 1999, all in electrical engineering.

He is with the Institute for Infocomm Research, Singapore (formerly Centre for Wireless Communications, NUS and Institute for Communications Research, NUS), first as a Research Engineer and currently as a Scientist, since 1994. His research interests are in communications networks, 3G/4G, WLAN, WPAN, and ultrawideband wireless mobile multimedia networks. His areas of research are in the medium access control, resource allocation with quality of service constraints, traffic policing with heterogeneous traffic and cross-layer design.

Dr. Wong was on the Technical Program Committee of the IEEE WCNC 2003, IEEE WCNC 2005, IEEE GLOBECOM 2005, ICCS 2006, and VTC Fall.