

Prioritized Burst Segmentation and Composite Burst-Assembly Techniques for QoS Support in Optical Burst-Switched Networks

Vinod M. Vokkarane, *Student Member, IEEE*, and Jason P. Jue, *Member, IEEE*

Abstract—In this paper, we address the issue of providing quality-of-service (QoS) in an optical burst-switched network. QoS is provided by introducing prioritized contention resolution policies in the network core and a composite burst-assembly technique at the network edge. In the core, contention is resolved through prioritized burst segmentation and prioritized deflection. The burst segmentation scheme allows high-priority bursts to preempt low-priority bursts and enables full class isolation between bursts of different priorities. At the edge of the network, a composite burst-assembly technique combines packets of different classes into the same burst, placing lower class packets toward the tail of the burst. By implementing burst segmentation in the core, packets that are placed at the tail of the burst are more likely to be dropped than packets that are placed at the head of the burst. The proposed schemes are evaluated through analysis and simulation, and it is shown that significant differentiation with regard to packet loss and delay can be achieved.

Index Terms—Burst assembly, burst segmentation, Internet protocol (IP), optical burst switching (OBS), quality-of-service (QoS), wavelength-division multiplexing (WDM).

I. INTRODUCTION

THE EXPLOSIVE growth of the Internet is resulting in an increased demand for higher transmission rate and faster switching technologies. Internet protocol (IP) over wavelength-division multiplexing (WDM) is a promising framework that can support the bandwidth and flexibility requirement of the next generation networks. In order to efficiently utilize the amount of raw bandwidth in WDM networks, an all-optical transport method, which avoids optical buffering while handling bursty traffic, must be developed. IP over optical burst switching (OBS) is one such method for transporting traffic directly over a bufferless WDM network [1].

In an OBS network, a data burst consisting of multiple IP packets is switched through the network all-optically. A burst-header packet (BHP) is transmitted ahead of the burst in order to configure the switches along the burst's route. In the just-enough-time (JET) signaling scheme, the burst transmission follows an out-of-band BHP after a fixed offset time. The offset

time allows for the BHP to be processed before the burst arrives at the intermediate node; thus, no fiber delay lines are necessary at the intermediate nodes. The BHP also specifies the duration of the burst in order to let a node know when it may reconfigure its switch for the next arriving burst [1].

An important issue in OBS networks is contention resolution. Contention occurs when multiple bursts contend for the same network resources. Existing contention resolution schemes include deflection [2]–[4], wavelength conversion [4], [5], and buffering [5], [6]. An approach for reducing packet loss due to contention is *burst segmentation* [7]. Burst segmentation is the process of dropping only those parts of a burst which overlap with another burst. A variation of segmentation in which overlapping segments of the head of the latter-arriving burst are dropped is described in [8].

Another issue in OBS is burst assembly. Burst assembly is the process of aggregating and assembling IP packets into a burst at the edge of the network. The most common burst-assembly approaches are *timer-based* and *threshold-based*. In a timer-based burst assembly approach, a burst is created and sent into the optical network at periodic time intervals [9]; hence, the network may have variable length input bursts. In a threshold-based approach, a limit is placed on the number of packets contained in each burst; hence, the network will have fixed-size input bursts [10]. Timer-based and threshold-based approaches may also be combined into a single burst-assembly scheme.

In an IP over OBS network, it is desirable to provide QoS support for applications with diverse QoS demands. Several solutions have been proposed to support QoS in the OBS core network. In [11], an additional-offset-based scheme was proposed. In this offset-based reservation scheme, higher priority bursts are given a larger offset time than the lower priority bursts. By providing a higher offset time, the probability of reserving the resources for the higher priority burst is increased, and therefore, the loss of higher priority bursts is decreased. The limitations of the additional offset-based scheme are unfavorable end-to-end delay and unfairness [12], [13].

In this paper, we focus on the issue of providing QoS support in OBS through prioritized contention resolution and composite burst assembly. In the prioritized contention resolution scheme, priorities are included as a field in the BHP. This priority field is used to preferentially segment and deflect bursts when resolving contentions in the core. The composite burst-assembly technique is implemented at the OBS network edge and attempts to meet the delay and loss constraints of each IP packet class

Manuscript received August 10, 2002; revised April 11, 2003. This work was supported in part by the National Science Foundation (NSF) under Grant ANI-01-33899. This paper was presented in part at the IEEE/OSA OFC 2002, Anaheim, CA, March 2002 and in part at the IEEE GLOBECOM 2002, Taipei, Taiwan, November 2002.

The authors are with the Department of Computer Science at the University of Texas at Dallas, Richardson, TX 75083 USA (e-mail: vinod@utdallas.edu; jjue@utdallas.edu).

Digital Object Identifier 10.1109/JSAC.2003.815681

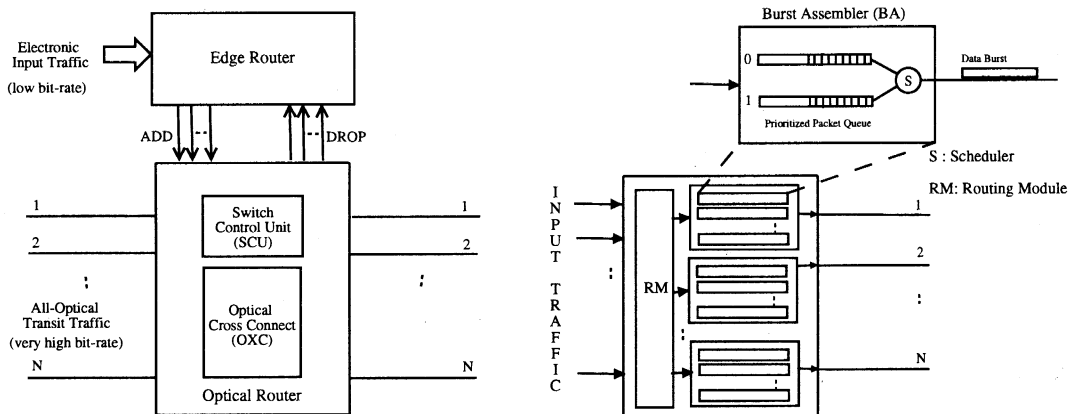


Fig. 1. (a) Node architecture. (b) Architecture of edge router.


without introducing any extra offset time. We develop a generalized framework for describing a wide range of burst assembly schemes and provide specific examples of composite burst-assembly schemes. Analytical and simulation models are developed to evaluate the packet loss probability of the various QoS schemes.

In this paper, we assume that JET signaling is used and that there are no fiber delay lines or wavelength converters in the network. The QoS requirements of an IP packet are defined by the packet's *class*, whereas bursts are differentiated in the core based on assigned *priorities*.

The remainder of the paper is organized as follows. Section II discusses the architecture of core and edge routers. Section III discusses the prioritized contention resolution policies employing burst segmentation and deflection. Section IV describes the generalized burst assembly framework. Section V describes the proposed burst assembly techniques. In Section VI, we develop an analytical model to calculate the packet loss probability for the proposed prioritized burst segmentation technique. Section VII provides numerical results from simulation and analysis, and compares the results of the different burst-assembly schemes. Section VIII concludes the paper.

II. OBS NETWORK ARCHITECTURE

An OBS network consists of a collection of edge and core routers. The edge routers assemble the electronic input packets into an optical burst which is sent over the OBS core. The ingress node presorts and schedules the incoming packets into electronic input buffers according to each packet's class and destination address. The packets are then aggregated into bursts that are stored in the output buffer. Since a separate buffer is required for each packet class and each destination, the limit on the maximum number of supported packet class is determined by the maximum electronic packet buffer size at each ingress node. The assembled bursts are transmitted all-optically over OBS core routers without any storage at intermediate core nodes. The egress node, upon receiving the burst, disassembles the burst into packets and provides the packets to the upper layer. Basic architectures for core and edge routers in an OBS network have been studied elsewhere [14].



In our network architecture, each node supports both new input traffic as well as all-optical transit traffic. Hence, each node consists of both an edge router and a core router, as shown in Fig. 1(a). The detailed architecture of the edge router is shown in Fig. 1(b).

The core routers primarily consist of an optical cross-connect (OXC) and a switch control unit (SCU). The SCU creates and maintains a forwarding table and is responsible for configuring the OXC. When the SCU receives a BHP, it identifies the intended destination and consults the forwarding table to find the intended output port. If the output port is available when the data burst arrives, the SCU configures the OXC to let the data burst pass through. If the port is not available, then the OXC is configured depending on the contention resolution policy implemented in the network.

The edge router performs the functions of presorting packets, buffering packets, and assembling packets into bursts. The architecture of the edge router consists of a routing module (RM), a burst assembler (BA), and a scheduler. The RM selects the appropriate output port for each packet and sends each packet to the corresponding BA module. Each BA module assembles bursts consisting of packets which are headed for a specific egress router. In the BA module, there is a separate packet queue for each class of traffic. The scheduler creates a burst based on the burst-assembly technique and transmits the burst through the intended output port. At the egress router, a burst-disassembly module disassembles the bursts into packets and send the packets to the upper network layers.

III. PRIORITIZED CONTENTION RESOLUTION

To overcome some of the limitations of OBS, burst segmentation can be used to minimize packet loss during contention. In burst segmentation, a burst is divided into multiple segments, and when contention occurs, only those segments of a given burst which overlap with segments of another burst will be dropped. If switching time is nonnegligible, then additional segments may be lost when the output port is switched from one burst to another. Segmentation can be used to minimize loss of packets during a contention, and can also allow high-priority bursts to preempt low-priority bursts. In these discussions,

the burst which arrives at a node first will be referred to as the *original* burst, and the burst which arrives later will be referred to as the *contending* burst. There are two approaches for segmenting a burst when contention occurs. The first approach is to segment the tail of the original burst, and the second approach is to segment the head of the contending burst. A significant advantage of segmenting the tail of bursts rather than segmenting the head is that there is a better chance of in-sequence delivery of packets at the destination, assuming that dropped packets are retransmitted at a later time. In this paper, we will assume that the remaining tail of the original burst will be dropped when segmentation takes place. Also, when a burst is segmented, its control message is updated accordingly.

Burst segmentation can also be implemented with deflection. Rather than dropping the tail segment of the original burst, we can either deflect the entire contending burst, or we can deflect the tail segment of the original burst. Implementing segmentation with deflection increases the probability that a burst's packets will reach the destination and, hence, improves performance. At each node, one or more alternate deflection ports can be specified for each destination. The order in which the alternate deflection ports are attempted is determined by a shortest-path policy.

The foundation for providing QoS in IP over OBS networks is service differentiation in the OBS core. We introduce and evaluate a new approach for such differentiation based on the concepts of burst segmentation and burst deflection. Burst segmentation enables the contending burst to preempt the original burst; hence, we have a choice of dropping either the contending burst or segmenting the original burst during a contention. Bursts are assigned priorities which are stored in the BHP, and contention between bursts is resolved through selective segmentation, deflection, and burst dropping based on these priorities.

We approach the general problem by first defining the possible segmentation and deflection policies which can be applied when a contention occurs. We then define the possible contention scenarios, which can take place between bursts of different priorities and lengths. Finally, we specify which policy to apply for each specific contention scenario.

When two bursts contend with one another, one of the following policies may be applied to resolve the contention.

- *Drop policy (DP)*: The original burst wins the contention. The entire contending burst is dropped.
- *Segment and drop policy (SDP)*: The contending burst wins the contention. The original burst is segmented and the tail segments of the original burst are dropped.
- *Deflect drop policy (DDP)*: The contending burst is deflected to an alternate port if an alternate port is available. If no alternate port is available, then the contending burst is dropped.
- *Segment first and deflect policy (SFDP)*: The original burst is segmented, and the tail segments of the original burst may be deflected if an alternate port is available, otherwise the tail segments of the original burst are dropped.
- *Deflect first, segment, and drop policy (DFSDP)*: The contending burst is deflected to an alternate port if an alternate




Fig. 2. (a) Contention of a low-priority burst with a high-priority burst. (b) Contention of a high-priority burst with a low-priority burst. (c) Contention of two equal-priority bursts with longer contending burst. (d) Contention of two equal-priority bursts with shorter contending burst.

port is available. If no alternate port is available, then the original burst is segmented and the tail segments of the original burst are dropped, while the contending burst is routed to the original output port.

We consider a total of four different contention scenarios which are based on the priorities and lengths of the original and contending bursts. When two bursts contend, the original burst may be of higher priority than the contending burst, the original burst may be of lower priority than the contending burst, or the two bursts may be of equal priority. For the situation in which bursts are of equal priority, we can break the tie by considering whether the length of the contending burst is longer or shorter than the remaining tail of the original burst. For each of these four contention scenarios, we specify one of the contention resolution policies described above.

Fig. 2 illustrates the possible contention scenarios. For the situation in which the priority of the contending burst is lower than the original burst, the contending burst should be deflected or dropped; thus, DDP will be applied. On the other hand, if the contending burst is of higher priority, then it should preempt the original burst. In this situation, SFDP will be applied. For the case in which both bursts are of equal-priority, we

TABLE I
QOS SCHEMES

Contention Scenario	Priority Relationship	Length Relationship	Scheme 1 Deflect/Segment	Scheme 2 Segment/Deflect	Scheme 3 Segment	Scheme 4 Deflect	Scheme 5 Drop
1	$P_o > P_c$	any	DDP	DDP	DP	DDP	DP
2	$P_o < P_c$	any	SFDP	SFDP	SDP	DDP	DP
3	$P_o = P_c$	$L_o > L_c$	DDP	DDP	DP	DDP	DP
4	$P_o = P_c$	$L_o < L_c$	DFSDP	SFDP	SDP	DDP	DP

should attempt to minimize the total number of packets which are dropped or deflected; thus, we compare the length of the contending burst with the remaining length (tail) of the original burst. If the contending burst is shorter than the tail of the original burst, then the contending burst should be deflected or dropped; thus, the DDP policy is applied. If the contending burst is longer than the tail of the original burst, then we have the option of either attempting to segment and deflect the tail of the original burst, or attempting to deflect the entire contending burst; thus, either DFSDP or SFDP may be applied. We consider both options, referring to the scheme in which DFSDP is applied as Scheme 1, and the scheme in which SFDP is applied as Scheme 2. For comparison, we further define schemes which do not take advantage of either segmentation or deflection. In Scheme 3, segmentation is supported but deflection is not, while in Scheme 4, deflection is supported but segmentation is not. In Scheme 5, neither deflection nor segmentation are supported. These schemes are summarized in Table I. The terms P_o and P_c refer to the priorities of the original burst and contending burst, respectively, and the terms L_o and L_c refer to the remaining length of the original burst and the length of the contending burst, respectively.

IV. GENERALIZED BURST ASSEMBLY FRAMEWORK FOR QOS SUPPORT

In this section, we formulate a generalized framework for burst assembly. The primary issues are which class of packets and how many of packets of each class to put into a burst. To provide QoS support, the burst assembly policies should take into account the number of packet classes as well as the number of burst priorities supported in the core. A burst can contain packets of a particular class [Fig. 3(a)], or a combination of packets of different classes [Fig. 3(b)]. Existing burst assembly techniques assemble packets of the same class into a burst. We introduce a new approach of assembling packets of different classes into a single burst, namely, *composite burst assembly*. This approach is motivated by the observation that, with burst segmentation, packets toward the tail of a burst are more likely to be dropped than packets at the head of a burst; thus, packet classes which have low loss tolerance may be placed toward the head of a burst while packet classes which have higher loss tolerance may be placed toward the tail of a burst. Furthermore, by implementing composite burst assembly, the network can support differentiation even if the number of IP packet classes exceeds the number of burst priorities supported in the core.

Another issue in burst assembly is when to create a burst. Typically, threshold and timer based approaches are utilized. In a timer-based approach, a timer is started when a packet arrives. When the timer expires, a burst is created from all packets received. In a threshold-based approach, an upper bound is placed on the number of packets in the burst. When the threshold is reached, a burst is created. Below, we provide a generalized framework for classifying various burst assembly approaches.

Let N be the number of input packet classes at the edge and let M be the number of burst priorities supported in the core network. Given N packet classes and M burst priorities, the objective is to meet the QoS requirements by defining a set of *burst types*, which specify how packets are aggregated, and by assigning an appropriate burst priority to each burst type. In this model, we define the length of the burst by the number of packets in the burst. Let K be the number of burst types, where $M \leq K \leq (2^N - 1)$. A burst type of type k is characterized by the following parameters:

- L_k^{MIN} : minimum length of burst of type k ;
- L_k^{MAX} : maximum length of burst of type k ;
- R_{jk}^{MIN} : minimum number of packets of Class j in a burst of type k ;
- R_{jk}^{MAX} : maximum number of packets of Class j in a burst of type k ;
- $S_k = \{j \mid R_{jk}^{\text{MAX}} > 0\}$: the set of packet classes which may be included in a burst of type k ;
- P_k : priority of burst of type k ;
- τ_k : timeout value for creating bursts of type k ;
- T_k : threshold value for creating bursts of type k ;
- C_k : $C_k \subseteq S_k$, subset of packet classes over which the threshold is evaluated. If x_j is the number of packets of Class j at the ingress node, then a burst is created if $\sum_{j \in C_k} x_j \geq T_k$.

The burst creation criterion for a burst of type k is satisfied either when the threshold value T_k for packets in C_k is satisfied, or when the timeout value, τ_k is reached. When the criterion is satisfied, a burst of type k is created, and the classes of packets to be included in the burst are specified by S_k . Packets are added to the burst until L_k^{MAX} is reached.

For example, in a threshold-based approach ($T_k \geq L_k^{\text{MIN}}$), if $S_k = \{1, 2\}$, then C_k can be $\{1, 2\}$, $\{1\}$, or $\{2\}$. If $C_k = \{1, 2\}$, then a burst of type k is created when the sum of packets of Class 1 and Class 2 is $\geq T_k$. If $C_k = \{1\}$, then a burst of type k is created when the number of packets of Class 1 is $\geq T_k$. If $C_k = \{2\}$, then a burst of type k is created when the number

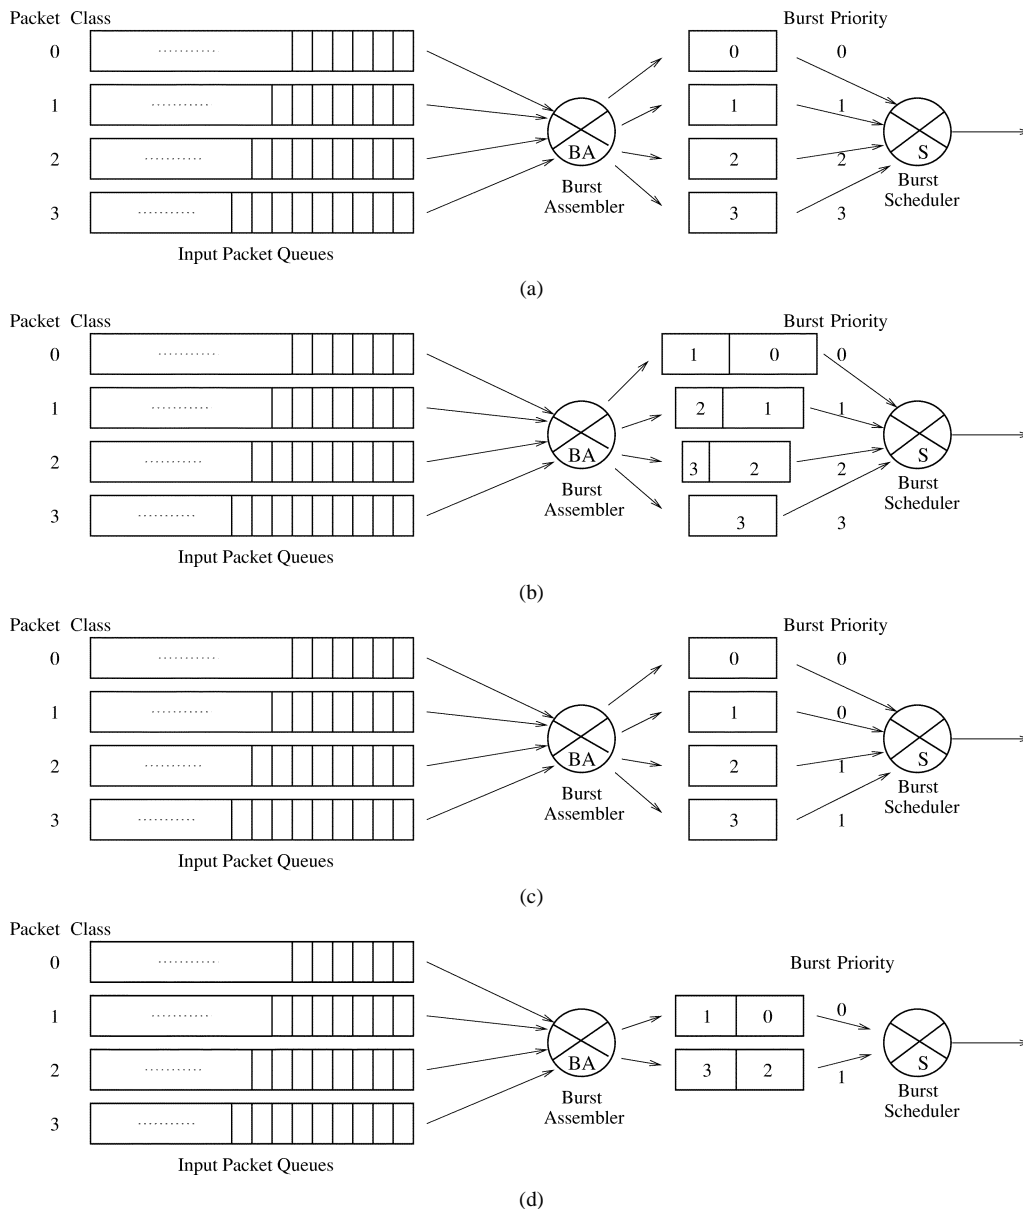


Fig. 3 (a) Creation of single-class burst with $N = 4$ and $M = 4$. (b) Creation of composite-class burst with $N = 4$ and $M = 4$. (c) Creation of single-class burst with $N = 4$ and $M = 2$. (d) Creation of composite-class burst with $N = 4$ and $M = 2$.

of packets of Class 2 is $\geq T_k$. In each of these cases, packets of both Class 1 and Class 2 may be included in the burst until L_k^{MAX} is reached.

V. BURST-ASSEMBLY TECHNIQUES

We now provide general guidelines for defining various burst types. The important design considerations when defining the burst types are packet loss probability, delay constraints, and bandwidth requirements. By appropriately mapping packet classes to burst types and by assigning appropriate priorities, P_k , to burst types, differentiated levels of packet loss may be achieved. End-to-end delay constraints can be met by setting appropriate timeout values τ_k for each burst type. Bandwidth requirements can be met by choosing an appropriate R_{jk}^{MIN} and R_{jk}^{MAX} for each packet class. In this paper, we focus primarily on achieving differentiated loss and delay. A fixed value of T_k

is assigned for all burst types and a timeout value τ_k is assigned only to the highest priority burst. We investigate the following approaches for selecting mappings S_k and priorities P_k to achieve differentiated QoS.

A. Approach 1: Single-Class Burst (SCB) With $N = M$

For the case in which $N = M$, we can create $K = M$ burst types such that each burst only contains a single class of packets ($S_k = \{k\}$). The priority of a burst will be equal to the class of packets contained in the burst ($P_k = k$). If a threshold-based approach is adopted, then the threshold, T_k for a priority k burst will be evaluated over Class k packets ($C_k = \{k\}$).

For example, if $N = 4$ and $M = 4$, as shown in Fig. 3(a), we set the number of burst types K equal to four. We set $S_0 = C_0 = \{0\}$, $S_1 = C_1 = \{1\}$, $S_2 = C_2 = \{2\}$, and $S_3 = C_3 = \{3\}$. If we consider the Class 2 packets that are collected in an input queue, once the number of Class 2 packets exceeds T_2 , a burst

consisting of Class 2 packets is created and sent into the network with a burst Priority 2. This process is followed for each class; thus, the priority of a burst will directly correspond to a specific class of packets contained in the burst.

B. Approach 2: Composite-Class Burst (CCB) With $N = M$

In composite bursts, each burst can consist of packets of different classes. One approach is to have $K = M$ burst types with a burst of type k containing packets of both Class k and Class $k + 1$, i.e., $S_k = \{k, k + 1\}$. In this approach, packets are placed in the burst in decreasing order of class, such that the higher class packets are at the head of the burst. A burst of type k is generated if the number of packets of Class k is equal to the threshold T_k ($C_k = \{k\}$) or if the timeout τ_k has expired. The priority of the burst is given by the burst type ($P_k = k$).

For example, if $N = M = 4$, as shown in Fig. 3(b), the number of burst types, K is equal to four. We set $S_0 = \{0, 1\}$, $S_1 = \{1, 2\}$, $S_2 = \{2, 3\}$, and $S_3 = \{3\}$. Here, $C_0 = \{0\}$, $C_1 = \{1\}$, $C_2 = \{2\}$, and $C_3 = \{3\}$. If the threshold of packet Class 1 is met, then a burst of type 1 is created with packets of class $S_1 = \{1, 2\}$, where Class 1 packets are placed at the head of the burst and Class 2 packets are placed at the tail of the burst. It is important to notice that there is no additional overhead incurred when ordering packets during the creation of the burst, since it is possible to access a particular input packet queue, place its contents in a burst, then go to the next lower class queue. This process can be repeated for all packet classes in S_k .

C. Approach 3: Single-Class Burst (SCB) With $N > M$

We now consider single-class bursts for the case $N > M$. In this approach, we have $K = N$ types of bursts, where each burst consists of packets of a single class ($S_k = \{k\}$). However, several burst types will have the same burst priority given by $P_k = \lfloor kM/N \rfloor$.

For example, if $N = 4$ and $M = 2$, as shown in Fig. 3(c), we set the number of burst types K equal to four. We have four unique types of bursts, each containing a single class of packets, i.e., $S_0 = C_0 = \{0\}$, $S_1 = C_1 = \{1\}$, $S_2 = C_2 = \{2\}$, and $S_3 = C_3 = \{3\}$. Each burst is assigned one of the two burst priorities. Bursts containing either Class 0 or Class 1 packets have Priority 0, while bursts containing either Class 2 or Class 3 packets have Priority 1.

D. Approach 4: Composite-Class Burst (CCB) With $N > M$

We now consider composite-class bursts for the case $N > M$. In this case, we have $K = M$ types of burst, where each burst consists of packets of class given by, $S_k = \{kN/M, \dots, (k + 1)N/M - 1\}$. A burst of type k is generated if the sum of packets of classes in C_k is equal to the threshold T_k ($C_k = S_k$). Once the threshold or timer criterion is met, a burst of type k containing packets defined by S_k is generated by appending all constituent class packets into the burst in decreasing order of class, such that the highest class packet in that burst type is at the head of the burst. The priority of the burst is same as type of burst ($P_k = k$).

For example, if $N = 4$ and $M = 2$, as shown in Fig. 3(d), we set the number of burst types K is equal to two. We select

$S_0 = C_0 = \{0, 1\}$ and $S_1 = C_1 = \{2, 3\}$. If the sum of Class 0 and Class 1 packets meet the threshold T_0 , then a burst of type 0 is created with packets of class $S_0 = \{0, 1\}$. The two types of composite bursts $\{0, 1\}$ and $\{2, 3\}$ are assigned burst Priority 0 and Priority 1, respectively.

VI. ANALYTICAL MODEL

In this section, we develop an analytical model for evaluating the packet loss probabilities with prioritized burst segmentation and composite burst assembly. We evaluate a modified version of Scheme 3 in which no length comparison is done. If two bursts are of equal-priority, we give priority to the contending burst. We assume that high- and low-priority bursts arrive to the network according to a Poisson process with rate λ^{sd} and γ^{sd} bursts per second for source-destination pair sd , respectively. Fixed routing is assumed and no buffering is supported at core nodes. We also assume that all bursts have the same offset time. This implies that the BHP of the original burst always arrives before the BHP of the contending burst. Traffic on each link is assumed to be independent. We consider a two-priority OBS network such that, Priority 0 bursts have higher priority than Priority 1 bursts. First, we analyze the packet loss probability for the high-priority bursts. We begin by defining the following notation:

- λ_l^{sd} : arrival rate of high-priority bursts to link l on the path between source s and destination d ;
- γ_l^{sd} : arrival rate of low-priority bursts to link l on the path between source s and destination d ;
- $\lambda_l = \sum_{sd} \lambda_l^{sd}$: arrival rate of high-priority bursts to link l , due to all source-destination pairs sd ;
- $\gamma_l = \sum_{sd} \gamma_l^{sd}$: arrival rate of low-priority bursts to link l , due to all source-destination pairs sd ;
- r_{sd} : route from source s to destination d .

The load placed on a link l by traffic going from source s to destination d depends on whether link l is on the path to destination d . If link l is on the path to d , then the load applied to link l by sd traffic is simply λ^{sd} . Thus

$$\lambda_l^{sd} = \begin{cases} \lambda^{sd}, & \text{if } l \in r_{sd} \\ 0, & \text{if } l \notin r_{sd} \end{cases}. \quad (1)$$

Also, the total high-priority (new) burst arrival into the network, λ , is given by

$$\lambda = \sum_s \sum_d \lambda^{sd}. \quad (2)$$

We calculate the packet loss probability by finding the distribution of the burst length at the destination and comparing the mean burst length at the destination to the mean burst length at the source. Let the initial cumulative distribution function of the burst length be $G_{l_0^{sd}}^0(t)$ for high-priority bursts transmitted from source s to destination d , where l_0^{sd} is the zeroth hop link between source s to destination d . The cumulative distribution function of the burst after k hops is $G_{l_k^{sd}}^0(t)$. Let $F_{l_k^{sd}}^0(t)$ be the cumulative distribution function for the arrival

time of the next high-priority burst on the k th hop link l between source-destination pair sd

$$F_{l_k}^0(t) = 1 - e^{-\lambda_{l_k}^{sd} t} \quad (3)$$

where $\lambda_{l_k}^{sd}$ is the arrival rate of all high-priority bursts on the k th hop link of the path between source s and destination d .

We note that a high-priority burst is segmented only if the next arriving burst is also of high-priority, but is not affected by the arrival of low-priority bursts. The burst length will be reduced if another high-priority burst arrives while the original burst is being transmitted; thus, the probability that the burst length is less than or equal to t after the first hop is equal to the probability that the initial burst length is less than or equal to t or that the next high-priority burst arrives in time less than or equal to t . Therefore

$$\begin{aligned} G_{l_1}^0(t) &= 1 - \left(1 - G_{l_0}^0(t)\right) \left(1 - F_{l_1}^0(t)\right) \\ &= 1 - \left(1 - G_{l_0}^0(t)\right) e^{-\lambda_{l_1}^{sd} t}. \end{aligned} \quad (4)$$

Similarly, let $G_2(t)$ be the cumulative distribution function of the burst after the second hop

$$\begin{aligned} G_{l_2}^0(t) &= 1 - \left(1 - G_{l_1}^0(t)\right) \left(1 - F_{l_2}^0(t)\right) \\ &= 1 - \left(1 - G_{l_0}^0(t)\right) e^{-(\lambda_{l_1}^{sd} + \lambda_{l_2}^{sd})t}. \end{aligned} \quad (5)$$

In general

$$\begin{aligned} G_{l_k}^0(t) &= 1 - \left(1 - G_{l_{k-1}}^0(t)\right) \left(e^{-\lambda_{l_k}^{sd} t}\right) \\ &= 1 - \left(1 - G_{l_0}^0(t)\right) e^{-\left(\sum_{i=1}^k \lambda_{l_i}^{sd}\right)t}. \end{aligned} \quad (6)$$

We now find the expected length after k hops and compare this length with the expected length at the source node in order to obtain the expected loss that a particular burst will experience. Let $L_{l_k}^0$ be the expected length of the high-priority burst at the k th hop.

Case (1): If we have fixed-sized bursts of length $1/\mu = T^0$, then the initial distribution of the burst length is given by

$$G_{l_0}^0(t) = \Pr(T \leq t) = \begin{cases} 1, & \text{if } t \geq T^0 \\ 0, & \text{if } t < T^0 \end{cases}. \quad (7)$$

Substituting (7) into (6) and taking the expected value, we obtain

$$L_{l_k}^0 = \frac{1 - e^{-\sum_{i=1}^k \lambda_{l_i}^{sd} T^0}}{\sum_{i=1}^k \lambda_{l_i}^{sd}}. \quad (8)$$

Case (2): If the initial burst length is exponentially distributed, we have

$$G_{l_0}^0(t) = 1 - e^{-\mu t}. \quad (9)$$

Substituting (9) into (6) and taking the expected value, we obtain

$$L_{l_k}^0 = \frac{1}{\sum_{i=1}^k \lambda_{l_i}^{sd} + \mu}. \quad (10)$$

We now find the expected length after K hops, where K is the total number of hops between s and d . Let Loss_{sd}^0 be the expected length of the burst lost per high-priority burst for a burst traveling from s to d

$$\text{Loss}_{sd}^0 = \frac{1}{\mu} - L_{l_K}^0. \quad (11)$$

Hence, the packet loss is proportional to the length of the route and the length of the burst. The packet loss probability of high-priority bursts, $P_{\text{loss}0}^{sd}$, is then given by

$$\begin{aligned} P_{\text{loss}0}^{sd} &= \frac{E[\text{Length Lost}]}{E[\text{Initial Length}]} \\ &= \text{Loss}_{sd}^0 \cdot \mu. \end{aligned} \quad (12)$$

We can then find the average packet loss probability of high-priority bursts for the system by finding the individual loss probability for each source-destination pair and taking the weighted average of the loss probabilities

$$P_{\text{loss}}^0 = \sum_s \sum_d \frac{\lambda^{sd}}{\lambda} P_{\text{loss}}^{sd}. \quad (13)$$

We also calculate the average service time on a link l , where l is the k th link from source s to destination d

$$\frac{1}{\bar{\mu}_l} = \sum_{s,d: \lambda_{l_i}^{sd} > 0} \frac{\lambda_{l_i}^{sd}}{\lambda_l} \cdot \frac{1}{\mu_{l_k}^{sd}} \quad (14)$$

where $\mu_{l_k}^{sd} = 1/L_{l_k}^0$.

Using $\bar{\mu}_l$, we can calculate the utilization for high-priority bursts on link l

$$\rho_l = 1 - e^{-\frac{\lambda_l}{\bar{\mu}_l}}. \quad (15)$$

Now, we calculate the probability of low-priority packet loss. The entire low-priority burst is dropped if a high-priority burst is occupying the channel. Thus, the arrival rate of low-priority bursts depends upon the link utilization of high-priority bursts. The offered load on the first hop is the total offered load from source to destination. On subsequent hops, the offered load is the load from the previous hop that was not blocked by high-priority traffic, thus

$$\gamma_i^{sd} = \begin{cases} \gamma^{sd}, & \text{if } l \in r_{sd}, l = l_0^{sd} \\ \gamma_h^{sd}(1 - \rho_h), & \text{if } l, h \in r_{sd}, l = l_i^{sd}, h = l_{i-1}^{sd}, i \geq 1 \\ 0, & \text{if } l \notin r_{sd}. \end{cases} \quad (16)$$

The calculation of low-priority packet loss probability is similar to that of high-priority packet loss. Let the initial cumulative distribution function of the burst length be $G_{l_0}^1(t)$, and the cumulative distribution function of the burst after k hops be $G_{l_k}^1(t)$ for low-priority bursts transmitted from source s to destination d . Let $F_{l_k}^1(t)$ be the cumulative distribution function for the arrival time of the next burst on the k th hop link. Here we consider the total arrival rate of bursts of both high and low priorities

$$F_{l_k}^1(t) = 1 - e^{-(\lambda_{l_k}^{sd} + \gamma_{l_k}^{sd})t} \quad (17)$$

where $\gamma_{l_k^{sd}}$ and $\lambda_{h_k^{sd}}$ are the arrival rates of all low and high-priority bursts on the k th hop link of the path between source s and destination d .

The burst length will be reduced if another burst of any priority arrives while the original burst is being transmitted; thus, the cumulative distribution function after the first hop is equal to the probability that the initial burst length is less than or equal to t or the next burst arrives in time less than or equal to t

$$\begin{aligned} G_{l_1^{sd}}^1(t) &= \rho_1 + (1 - \rho_1) \left[1 - \left(1 - G_{l_0^{sd}}^1(t) \right) \left(1 - F_{l_1^{sd}}^1(t) \right) \right] \\ &= 1 - (1 - \rho_1) \left(1 - G_{l_0^{sd}}^1(t) \right) e^{-(\lambda_{l_1^{sd}} + \gamma_{l_1^{sd}})t}. \end{aligned} \quad (18)$$

Similarly, $G_{l_2^{sd}}^1(t)$ is the cumulative distribution function of the burst length after the second hop

$$\begin{aligned} G_{l_2^{sd}}^1(t) &= \rho_2 + (1 - \rho_2) \left[1 - \left(1 - G_{l_1^{sd}}^1(t) \right) \left(1 - F_{l_2^{sd}}^1(t) \right) \right] \\ &= 1 - (1 - \rho_2) \left(1 - G_{l_0^{sd}}^1(t) \right) \\ &\quad \cdot e^{-(\lambda_{l_2^{sd}} + \lambda_{l_1^{sd}} + \gamma_{l_2^{sd}} + \gamma_{l_1^{sd}})t}. \end{aligned} \quad (19)$$

In general

$$\begin{aligned} G_{l_k^{sd}}^1(t) &= \rho_k + (1 - \rho_k) \left[1 - \left(1 - G_{l_{k-1}^{sd}}^1(t) \right) \left(1 - F_{l_k^{sd}}^1(t) \right) \right] \\ &= 1 - \prod_{i=1}^k (1 - \rho_i) \left(1 - G_{l_0^{sd}}^1(t) \right) e^{-\left(\sum_{j=1}^k \lambda_{l_j^{sd}} + \gamma_{l_j^{sd}} \right) t}. \end{aligned} \quad (20)$$

We now find the expected length after k hops and compare with the expected length at the source node to obtain the expected loss. Let $L_{l_k^{sd}}^1$ be the expected length of the low-priority burst at the k th hop.

Case (1): If we have fixed-sized bursts of length, $1/\mu = T^1$, the initial distribution of the burst length is given by

$$G_{l_0^{sd}}^1(t) = \Pr(T \leq t) = \begin{cases} 1, & \text{if } t \geq T^1 \\ 0, & \text{if } t < T^1. \end{cases} \quad (21)$$

Therefore, $L_{l_k^{sd}}^1$ is given by

$$L_{l_k^{sd}}^1 = \frac{\prod_{i=1}^k (1 - \rho_i) \left(1 - e^{-\sum_{i=1}^k (\lambda_{l_i^{sd}} + \gamma_{l_i^{sd}}) T^1} \right)}{\sum_{i=1}^k (\lambda_{l_i^{sd}} + \gamma_{l_i^{sd}})}. \quad (22)$$

Case (2): If the initial burst length is exponentially distributed, we have

$$G_{l_0^{sd}}^1(t) = 1 - e^{-\mu t}. \quad (23)$$

Therefore, $L_{l_k^{sd}}^1$ is given by

$$L_{l_k^{sd}}^1 = \frac{\prod_{i=1}^k (1 - \rho_i)}{\sum_{j=1}^k (\lambda_{l_j^{sd}} + \gamma_{l_j^{sd}}) + \mu}. \quad (24)$$

Let Loss_{sd}^1 be the expected length of the burst lost per low-priority burst for a burst traveling from s to d

$$\text{Loss}_{sd}^1 = \frac{1}{\mu} - L_{l_k^{sd}}^1. \quad (25)$$

The probability of packet loss for low-priority bursts is given by

$$P_{\text{loss}1}^{sd} = \text{Loss}_{sd}^1 \cdot \mu. \quad (26)$$

We can then find the average packet loss probability of low-priority bursts for the system by finding the individual loss probability for each source-destination pair, and taking the weighted average of the loss probabilities

$$P_{\text{loss}}^1 = \sum_s \sum_d \frac{\gamma^{sd}}{\gamma} P_{\text{loss}1}^{sd}. \quad (27)$$

Note that, if two contending bursts follow the same route, then the original burst will only be segmented at the first instance of contention; however, the model assumes that the arrivals of the two contending bursts are uncorrelated on the subsequent links in the route. Thus, the model overestimates the packet loss.

Also, if a burst is segmented in the middle of a packet, the model does not account for the entire packet loss, which leads to a slight under-estimation of packet loss. However, this under-estimation of packet loss is insignificant compared with the over-estimation of the packet loss due to the uncorrelated arrival assumption.

This analysis may be extended to any arbitrary number of priorities in a straightforward manner. Also, a more accurate model may be obtained by using a reduced load approximation for the arrival of the low-priority bursts and by taking into account the link correlation effect.

We now compute the packet loss probability for different packet classes in a CCB. We consider an OBS network with four packet classes and two burst priorities. Let Class 0, Class 1, Class 2, and Class 3 be the four packet classes with Class 0 being the highest packet class and Class 3 being the lowest packet class, in that order. The following are the assumptions:

- Initial burst length is fixed;
- T^0 : high-priority burst length;
- T^1 : low-priority burst length;
- α : ratio of Class 0 packets in the high-priority burst;
- β : ratio of Class 2 packets in the low-priority burst;
- The ratio of traffic of Class 1 and Class 4 will be $(1 - \alpha)$ and $(1 - \beta)$ in the high and low-priority bursts, respectively;
- Class 0 packets are placed toward the head and Class 1 packets are placed toward the tail of the high-priority burst;
- Class 2 packets are placed toward the head and Class 3 packets are placed toward the tail of the low-priority burst.

Based on the ratio of packets of each class, we can find the packet loss probabilities of each class. The packet loss for Class 0, P_{loss}^{00} , is the same as the loss probability of a high-priority burst of length $\alpha \cdot T^0$; therefore, we can obtain P_{loss}^{00} by replacing T^0 in (8) with $\alpha \cdot T^0$. The packet loss probability for Class 1 is found by considering the total packet loss probability




Fig. 4. NSFNET with 14 nodes.

in a burst and the packet loss probability of Class 0 packets; thus, P_{loss}^{01} is given by

$$P_{\text{loss}}^{01} = \frac{P_{\text{loss}}^0 - \alpha \cdot P_{\text{loss}}^{00}}{1 - \alpha}. \quad (28)$$

Similarly, the packet loss probability for Class 2, P_{loss}^{12} , is same as the packet loss probability of a low-priority burst of length $\beta \cdot T^1$ and can be found by replacing T^1 in (22) with $\beta \cdot T^1$. The packet loss probability for Class 3 is given by

$$P_{\text{loss}}^{13} = \frac{P_{\text{loss}}^1 - \beta \cdot P_{\text{loss}}^{12}}{1 - \beta}. \quad (29)$$

VII. NUMERICAL RESULTS

In order to evaluate the performance of the proposed schemes and to verify the analytical models, a simulation model is developed. Burst arrivals to the network are assumed to be Poisson. Burst lengths are exponentially distributed with average length of $100 \mu\text{s}$. The link transmission rate is 10 Gb/s . Packets are assumed to be 1250 bytes and each segment consists of a single packet. The configuration time of the switching is assumed to be $10 \mu\text{s}$. There is no buffering or wavelength conversion at the core nodes. Burst arrivals are uniformly distributed over all sender-receiver pairs, and shortest-path routing is assumed. Fig. 4 shows the 14-node NSFNET on which the simulation was implemented, with distances in km.

A. Analytical Results

Let us consider a network with two priorities. The fraction of high-priority (Priority 0) bursts is 20% and the fraction of low-priority (Priority 1) bursts is 80%. In the analytical model, we ignore the switching time and header processing time.

Fig. 5 plots the packet loss probability versus load for high-priority and low-priority packets for Scheme 1, with exponential burst length, and for fixed-sized bursts. In Scheme 1, the contending burst preempts the original burst if the contending burst is of equal or higher priority, otherwise, the contending burst is dropped. We observe that the analytical model slightly over-estimates the packet loss probabilities due to the independent link assumption. We also observe that the packet loss with fixed-sized bursts is lower than packet loss with exponentially distributed burst sizes, since the maximum number of packets lost per contention is potentially less with a fixed initial burst




Fig. 5. Packet loss probability versus load for both exponential initial burst size $1/\mu = 100 \text{ ms}$ and fixed initial burst size = 100 packets using Scheme 3 without length comparison.




Fig. 6. Packet loss probability versus alpha and beta values for composite bursts of fixed initial burst size = 100 packets length using Scheme 3 without length comparison.

size. This observation may be useful when determining the burst assembly policy.

We now consider an OBS network with composite bursts. The network supports four packet classes. Class 0 is the highest packet class and Class 3 is the lowest packet class. Fig. 6 plots the packet loss probability versus α and β for Scheme 3 without length comparison, where α is the ratio of Class 0 packets in the high-priority burst and β is the ratio of Class 2 packets in the low-priority burst. The graphs are plotted for a fixed load of 1 Erlang with fixed-sized bursts. We observe that the packet loss probability of the different classes obtained through the analytical models match with the simulation results. By choosing a specific value of α and β , we can ensure that a certain level of performance is guaranteed. For example, for the case shown in Fig. 6, if we choose $\alpha = 55\%$, then the packet loss probability of Class 0 will be less than 1%.




Fig. 7. Packet loss probability versus load for different schemes with fixed burst size = 100 packets with the traffic ratio being 20% Priority 0 and 80% Priority 1 bursts.

B. Prioritized Burst Segmentation Results

Fig. 7 plots the packet loss probability versus load for high-priority (Priority 0) and low-priority (Priority 1) packets for Scheme 1 through Scheme 5, with fixed-sized bursts. The graph shows packet losses for the case in which 20% of the traffic is high priority and 80% of the traffic is low priority. We observe that the loss of high-priority packets is lower than that for low priority packets in schemes which employ burst segmentation (Schemes 1, 2, and 3), while schemes without segmentation do not provide service differentiation (Schemes 4 and 5). We also observe that Scheme 1 performs the best under the observed load values, while Scheme 2, performs better at higher loads; thus, at low loads, it is better to attempt deflection before segmentation when two bursts are of equal priority. At higher loads, schemes with deflection as the primary contention resolution technique (Schemes 1 and 4) suffer from higher loss compared with schemes with no or controlled deflection (Schemes 2 and 3) due to the increased load due to deflection. Also, by varying the number of alternate deflection ports at each switch, we can achieve different levels of packet loss.

Fig. 8 plots the average end-to-end packet delay versus load for high-priority and low-priority packets for Scheme 1 through Scheme 5 with fixed-sized bursts. We observe that the delay of high-priority packets are lower than that for low-priority packets in schemes which employ burst segmentation (Schemes 1, 2, and 3). Schemes without segmentation do not provide service differentiation (Schemes 4 and 5) and, hence, have the same delays for both priorities. The delay for high-priority bursts remains in a consistent range, while the low-priority bursts have higher delay due to multiple deflections. At very high load, bursts which are further from their destination are less likely to reach their destination compared with those bursts which are close to their destination; thus, the average delay will eventually decrease at very high load. Schemes 1 and 4 suffer high delays compared with other schemes, since the contending burst (either lower or equal priority) is deflected first.

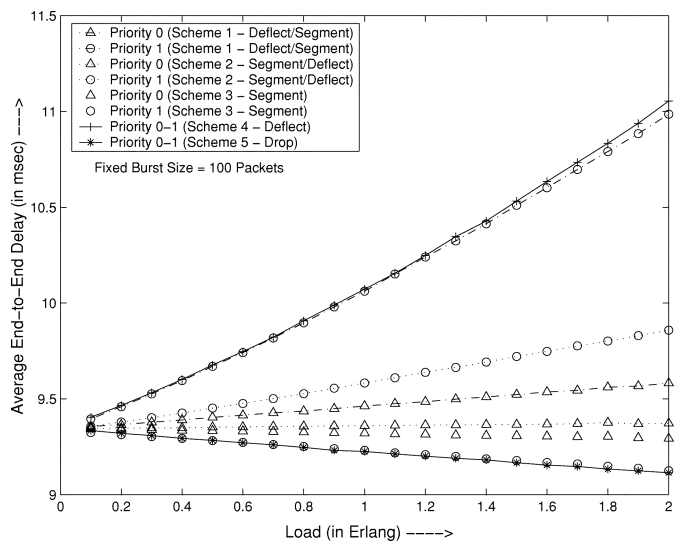


Fig. 8. Average end-to-end packet delay versus load for different schemes with fixed burst size = 100 packets with the traffic ratio being 20% Priority 0 and 80% Priority 1 bursts.




Fig. 9. Packet loss probability versus load for $N = 4$ and $M = 4$ for single- and composite-class bursts, with the traffic ratio of the packets classes 0, 1, 2, and 3 being 10%, 20%, 30%, and 40%, respectively.

C. Composite Burst-Assembly Results

The different burst assembly techniques are compared through simulation. We consider composite- and single-burst assembly while utilizing Scheme 3 without length comparisons for contention resolution in the core. The input traffic ratios of individual packet classes are 10%, 20%, 30%, and 40% for Class 0, Class 1, Class 2, and Class 3, respectively. We set a threshold value of 100 packets for each burst type, and a timeout value of 50 ms for the highest priority burst. We also avoid contentions between multiple bursts at the source by delaying the contending bursts until the desired output port is free. The remaining assumptions remain the same as the prioritized burst segmentation case.

Figs. 9 and 10 plot packet loss probability and average end-to-end delay versus load for both CCB and SCB with




Fig. 10. Average end-to-end packet delay versus load for $N = 4$ and $M = 4$ for single- and composite-class bursts, with the traffic ratio of the packets classes 0, 1, 2, and 3 being 10%, 20%, 30%, and 40%, respectively.




Fig. 12. Average end-to-end packet delay versus load for $N = 4$ and $M = 2$ for single- and composite-class bursts, with the traffic ratio of the packets classes 0, 1, 2, and 3 being 10%, 20%, 30%, and 40%, respectively.




Fig. 11. Packet loss probability versus load for $N = 4$ and $M = 2$ for single- and composite-class bursts, with the traffic ratio of the packets classes 0, 1, 2, and 3 being 10%, 20%, 30%, and 40%, respectively.

$N = M = 4$. We refer to this case as the 4:4 mapping. We observe that, by using CCB, the loss of packets is more proportional to the packet class than in SCB. We observe that the loss of lower class packets is better in CCB, since some of the lower class packets are placed into higher priority bursts, which, in turn, decreases the loss probability. Also, the highest class packets in CCB perform as well as in SCB, since at every contention between highest priority bursts, the lower-class packets are more likely to be dropped. We see that the average delay decreases with the increase in load. This decrease is due to the higher arrival rate of packets which causes the threshold to be satisfied more frequently. The delay of highest class packets is fairly constant, since we enforce an upper limit on the aggregation time by using a timeout.

Figs. 11 and 12 plot loss probability and average end-to-end delay versus load for both CCB and SCB with

$N = 4$ and $M = 2$. We refer to this case as the 4:2 mapping. We observe that the performance of CCB is much better than SCB for the highest class packets. This is due to the fact that in a 4:2 mapping, both packets of Class 0 and Class 1 are assigned Priority 0, and in an equal-priority contention, packets of Class 1 may preempt packets of Class 0. In SCB, the loss of Class 0 packets and Class 1 packets will be the same if the input ratio are the same, and if the same threshold and timeout values are used. In our example, a timeout value is assigned to bursts carrying Class 0 packets, but not to bursts carrying Class 1 packets. This difference results in lower loss and delay for Class 0 packets, even though the burst are of equal priority. Also, we see that the average end-to-end delay for Class 0 and Class 1 in the case of CCB are similar in both 4:4 and 4:2 mapping, since Class 1 packets are included in the same bursts as Class 0 packets when the timeout is reached. The difference in delay between Class 0 and Class 1 packets is due to their different arrival rates.

VIII. CONCLUSION

In this paper, we introduce the concept of prioritized contention resolution through differentiated burst segmentation and deflection to provide QoS in the optical burst-switched core. The prioritized contention resolution policies can provide QoS with 100% class isolation without requiring additional offset times. An analytical model for prioritized burst segmentation was developed to calculate the packet loss probabilities for a two-priority network and the model was verified through simulation. The high-priority bursts have significantly lower losses and delay than the low-priority bursts, and the schemes which incorporate deflection tend to perform better than the schemes with limited deflection or no deflection.

We also introduced the concept of composite burst assembly to handle the differentiated service requirements of the IP packets at edge nodes of the optical burst-switched network, and we described a generalized framework for burst assembly.

We considered four different burst assembly approaches and evaluated their performance in terms of delay and loss. We observe that approaches with composite bursts perform better than approaches with single-class bursts with respect to providing differentiated QoS for different classes of packets. This was verified by the analytical model results. The developed model can be useful for selecting the class ratios for composite bursts in a manner which can satisfy the packet loss requirement. In order to further reduce the packet loss, the proposed techniques can be employed in conjunction with all-optical wavelength conversion and buffering through fiber delay lines.

REFERENCES

- [1] C. Qiao and M. Yoo, "Optical Burst Switching (OBS)—A new paradigm for an optical Internet," *J. High Speed Networks*, vol. 8, no. 1, pp. 69–84, Jan. 1999.
- [2] A. S. Acampora and I. A. Shah, "Multihop lightwave networks: A comparison of store-and-forward and hot-potato routing," *IEEE Trans. Commun.*, vol. 40, pp. 1082–1090, June 1992.
- [3] F. Forghieri, A. Bononi, and P. R. Prucnal, "Analysis and comparison of hot-potato and single-buffer deflection routing in very high bit rate optical mesh networks," *IEEE Trans. Commun.*, vol. 43, pp. 88–98, Jan. 1995.
- [4] A. Bononi, G. A. Castanon, and O. K. Tonguz, "Analysis of hot-potato optical networks with wavelength conversion," *J. Lightwave Technology*, vol. 17, pp. 525–534, Apr. 1999.
- [5] S. Yao, B. Mukherjee, S. J. B. Yoo, and S. Dixit, "All-optical packet-switched networks: A study of contention resolution schemes in an irregular mesh network with variable-sized packets," in *Proc. SPIE OptiComm 2000*, Dallas, TX, Oct. 2000, pp. 235–246.
- [6] I. Chlamtac *et al.*, "CORD: Contention resolution by delay lines," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 1014–1029, June 1996.
- [7] V. M. Vokkarane, J. P. Jue, and S. Sitaraman, "Burst segmentation: An approach for reducing packet loss in optical burst switched networks," in *Proc. IEEE Int. Conf. Communications (ICC) 2002*, vol. 5, New York, NY, Apr. 2002, pp. 2673–2677.
- [8] A. Detti, V. Eramo, and M. Listanti, "Performance evaluation of a new technique for IP support in a WDM optical network: Optical Composite Burst Switching (OCBS)," *J. Lightwave Technology*, vol. 20, pp. 154–165, Feb. 2002.
- [9] A. Ge, F. Callegati, and L. S. Tamil, "On optical burst switching and self-similar traffic," *IEEE Commun. Lett.*, vol. 4, pp. 98–100, Mar. 2000.

- [10] V. M. Vokkarane, K. Haridoss, and J. P. Jue, "Threshold-based burst assembly policies for QoS support in optical burst-switched networks," in *Proc. SPIE OptiComm 2002*, vol. 4874, Boston, MA, July 2002, pp. 125–136.
- [11] M. Yoo, C. Qiao, and S. Dixit, "QoS performance of optical burst switching in IP-over-WDM networks," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 2062–2071, Oct. 2000.
- [12] Y. Chen, M. Hamdi, and D. H. K. Tsang, "Proportional QoS over OBS networks," in *Proc. IEEE GLOBECOM 2001*, vol. 3, San Antonio, TX, Nov. 2001, pp. 1510–1514.
- [13] F. Poppe, K. Laevens, H. Michiel, and S. Molenaar, "Quality-of-service differentiation and fairness in optical burst-switched networks," in *Proc. SPIE OptiComm 2002*, vol. 4874, Boston, MA, July 2002, pp. 118–124.
- [14] Y. Xiong, M. Vanderhoute, and H. C. Cankaya, "Control architecture in optical burst-switched WDM networks," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 1838–1851, Oct. 2000.



Vinod M. Vokkarane (S'02) received the B.Eng. degree in computer science and engineering from the University of Mysore, India, in 1999 and the M.S. degree in computer science from the University of Texas, Dallas, in 2001. He is working toward the Ph.D. degree in computer science at the University of Texas, where he is currently a Research Assistant at the Advanced Networks Research Laboratory.

His active area of research is in optical networks, focusing on the design and analysis of optical burst-switched network architectures and protocols.



Jason P. Jue (S'95–M'99) received the B.S. degree in electrical engineering from the University of California, Berkeley, in 1990, the M.S. degree in electrical engineering from the University of California, Los Angeles, in 1991, and the Ph.D. degree in computer engineering from the University of California, Davis, in 1999.

In 1999, he joined the faculty at the University of Texas at Dallas, where he is currently an Assistant Professor of computer science. He is actively conducting research in the area of wavelength-division multiplexing (WDM) optical networks, focusing on the design and analysis of optical network architectures and protocols.

Dr. Jue was a recipient of the NSF Career Award in 2002.