# Dynamic Traffic Grooming in Optical Burst-Switched Networks

Farid Farahmand, Qiong Zhang, *Member, IEEE*, and Jason P. Jue, *Senior Member, IEEE*

***Abstract*—In this paper, we address the problem of data-burst grooming in optical burst-switched (OBS) networks. In OBS networks, IP packets with the same edge-node destination are assembled into larger packets called data bursts. Depending on the core node's switching technology, data bursts are required to have a minimum length. On the other hand, each IP packet in a burst has a time delay constraint, called maximum end-to-end delay, which determines the upper time limit before which the packet must reach its destination. Thus, a data burst cannot wait indefinitely until a sufficient number of IP packets are assembled and the minimum burst length requirement is met. In order to satisfy the packet maximum end-to-end delay requirement, many bursts will be timed out and released before they reach the minimum-length requirement. Under such circumstances, padding overhead must be added to these short bursts, called sub-bursts. Excessive padding results in high overhead and high data-burst blocking probability. One approach to minimize the amount of padding overhead, while maintaining the end-to-end delay requirement of IP packets, is to groom multiple sub-bursts together. That is, sub-bursts with different destinations are aggregated together at the edge node and transmitted as a single burst until they are separated at some downstream node. In this paper, we present an edge-node architecture enabling burst-grooming capability. We also develop two basic grooming approaches, namely No routing overhead (NoRO) and minimum total overhead (MinTO). Through a comprehensive simulation study, we show that, in general, our proposed grooming algorithms can significantly improve the performance compared to the case of no grooming. However, careful considerations must be given to network-loading conditions and the number of sub-bursts allowed to be groomed together. We show that although simple greedy algorithms can reduce network overhead, they may alter the traffic characteristics and increase its burstiness, resulting in high packet blocking probability.**

***Index Terms*—Burst assembly, dynamic traffic, edge-node architecture, grooming, optical burst switching, padding overhead, routing overhead.**

## I. Introduction

THE AMOUNT of raw bandwidth available on fiber-optic links has increased dramatically with advances in dense wavelength division multiplexing (DWDM) technology; however, existing optical-network architectures are unable to fully utilize this bandwidth to support highly dynamic and bursty traffic. Optical burst switching [1], [2] has been proposed as a new paradigm to provide the flexible and dynamic bandwidth allocation required to support such traffic. In optical burst-switched (OBS) networks, incoming data are assembled into basic units, referred to as data bursts, which are then transported over the optical core network. Control signaling is performed out of band by control packets that carry information such as the length, the edge-node-destination address, and the QoS requirements of the optical data burst. The control packet is separated from the data burst by an offset time, allowing the control packet to be processed at each intermediate node before the data burst arrives. Aggregating IP packets into large-sized bursts can compensate for slow switching time at core nodes. This is motivated by the fact that the relatively mature MEMS-based optical crossconnects can provide a connection switching time of about 10 ms [5]; on the other hand, the typical switch reconfiguration time requirement for optical packets can be in the order of microseconds (or even nanoseconds). Consequently, core nodes with slower switching times require larger minimum burst lengths in order to minimize the switching overhead.

An important issue in OBS networks is data burst assembly. Burst assembly is the process of aggregating IP packets with the same characteristics, such as edge-node destination, class of service, etc., into a burst at the edge node. 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 when the time-out event is triggered. In a threshold-based approach, a limit is placed on the number of packets contained in each burst. A more efficient assembly scheme can be achieved by combining the timer-based and threshold-based approaches [6]–[9].

IP packets assembled in a data burst have a time delay constraint, called maximum end-to-end delay tolerance, determining the deadline by which the packet must reach its OBS destination. Thus, the main motivation for implementing the timer-based burst-assembly approach is to ensure an IP packet does not wait at the edge node's assembly unit indefinitely before its maximum end-to-end delay tolerance is violated. If the arrival rate of incoming IP packets with the same characteristics is low, bursts are timed out and released before they reach their minimum burst length requirement determined by the core-node switching time. Under such conditions, the timed-out burst is smaller than the minimum-length requirement. We refer to these short bursts as sub-bursts. Padding overhead must be added to sub-bursts in order to satisfy the minimum-length

Manuscript received December 16, 2004; revised June 17, 2005. This work was supported in part by the National Science Foundation (NSF) under grant ANI-01-33899.

F. Farahmand is with the Department of Computer Electronics and Graphics Technology, Central Connecticut State University, New Britain, CT 06050 USA (e-mail: farahmandfar@ccsu.edu).

Q. Zhang is with the Department of Mathematical Sciences and Applied Computing, Arizona State University—West Campus, AZ 85069 USA (e-mail: Qiong.Zhang@asu.edu).

J. P. Jue is with the Department of Computer Science, The University of Texas at Dallas, Richardson, TX 75083 USA (e-mail: jjue@utdallas.edu).

Digital Object Identifier 10.1109/JLT.2005.856236

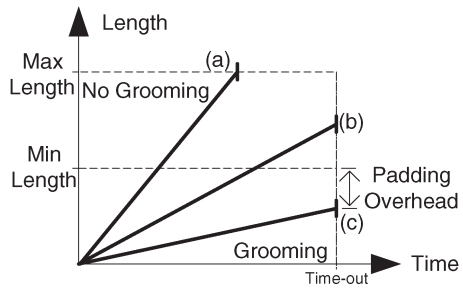


Fig. 1. Illustrating the timer-based and threshold-based burst-assembly approaches.

requirement. However, excessive padding results in high link utilization and data-burst blocking probability. Furthermore, when data bursts are timed out, their aggregated IP packets will experience higher average delay. These concepts are illustrated in Fig. 1. In case (a), the data burst reaches its maximum size before it is timed out. Case (b) represents a situation in which the burst is timed out before it reaches its maximum size. In case (c), the data burst is timed out before it reaches the minimum required length and padding overhead must be added. Note that in this paper, we mainly focus on case (c), representing instances when the incoming IP-packet arrival rate of sub-bursts is low. Consequently, in such cases, the burst-assembly approach will be timer based, and bursts will be released prior to reaching their minimum-length requirement. The padding overhead will increase the network load and can lead to increased blocking in the network.

One approach to minimize the amount of padding overhead, as well as the average end-to-end IP-packet delay due to low IP-packet arrival rate, is to groom bursts. Burst grooming is defined as aggregating multiple sub-bursts with different characteristics (i.e., edge-node destinations) together at the edge node and transmitting them as a single burst. In situations where the overall load is high, if there are still several sub-bursts with low arrival rate, the padding overhead for these sub-bursts can still have a significant impact on the network performance, particularly on bottleneck links. Thus, even under higher overall network loads, burst grooming may potentially improve network performance.

The problem of aggregating and routing sub-bursts together, as well as determining their wavelength assignment, is referred to as the data-burst-grooming problem. Heuristic algorithms that attempt to solve the data-burst-grooming problem are referred to as burst-grooming algorithms. These algorithms differ depending on their aggregation and routing criteria. For example, issues such as which sub-bursts and how many sub-bursts can be groomed together, or how long the accumulated length of the groomed burst should be, can have a significant impact on the efficiency of the grooming algorithm under different network-loading conditions. We note that the general burst-grooming concept can be implemented in conjunction with any given scheduling and routing algorithms.

The concept of traffic grooming has been extensively studied for various circuit-switched WDM network topologies (ring, mesh, etc.) under different traffic scenarios (static or dynamic) [10]–[14]. The basic idea in all these problems is to share lightpaths, defined as wavelength channels dedicated to established connections. Hence, we refer to these problems as lightpath-based grooming problems. The objective of data-burst grooming in OBS over WDM networks, however, is to aggregate multiple sub-bursts to share the data burst created to satisfy a request. Data-burst grooming in OBS has not received much attention in the literature. In [15], the authors consider data-burst grooming at core nodes where several sub-bursts sharing a common path can be aggregated together in order to reduce switching overhead. The aggregated sub-bursts can be separated at a downstream node prior to reaching their final destinations.

In this paper, we address the problem of data-burst grooming in OBS networks. In our study, we concentrate on grooming data bursts at the edge nodes. This study is motivated by the following network constraints: 1) the data traffic through the network is bursty in nature and connections are short lived; 2) at low IP-packet arrival-rate instances, the core-node switching time is much larger than the average IP-packet size; and 3) incoming IP packets passing through the network have a maximum end-to-end delay tolerance. We emphasize that under such conditions, traffic-aware assembly and routing schemes may not be efficient due to the bursty nature of the traffic. Similarly, lightpath-based grooming with static connections will not be suitable because it does not support on-demand network reconfigurability. On the other hand, dynamically reconfigurable lightpath-based grooming may not efficiently utilize the available bandwidth because data connections have short duration relative to the setup time of the lightpaths. Note that without the delay-tolerance constraint, packets can stay in the assembly unit indefinitely and there will be no need for burst grooming.

The main contribution of this paper is an edge-node architecture for enabling burst grooming, as well as several data-burst-grooming heuristic algorithms. Using simulation, we examine the performance of our proposed grooming algorithms under specific network conditions. We compare our results with those obtained without burst grooming in terms of blocking probability and average end-to-end IP-packet delay. We show that our proposed burst-grooming techniques lead to performance improvement when the IP-traffic arrival rate is low.

The remainder of this paper is organized as follows. In Section II, we describe the proposed edge-node architecture in OBS networks capable of supporting data-burst grooming. Section III formulates the data-burst-grooming problem and provides descriptions of two proposed grooming algorithms. The performance results for each algorithm are presented in Section IV. Finally, Section V concludes the paper.

## II. Node Architecture

The general core-node architecture is described in detail in [3] and [4]. We assume that the switching time for core nodes is given as $\tau$, and that the minimum required data-burst duration is defined as a function of $\tau$; i.e., $L^{\mathrm{MIN}} = f(\tau)$. Throughout this paper, we refer to sub-bursts as the aggregated IP packets with the same edge-node destination, whose total length is less than $L^{\mathrm{MIN}}$. Hence, a transmitted burst can contain multiple sub-bursts.

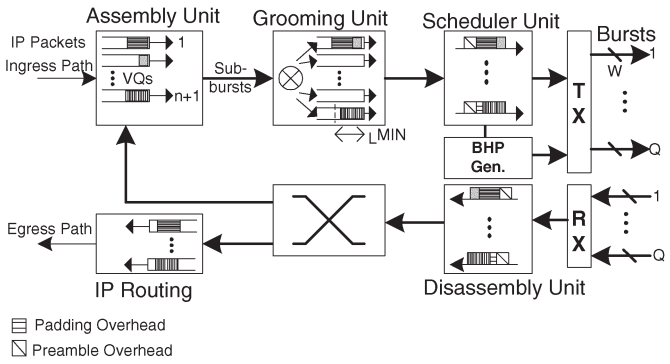


Fig. 2. Edge-node architecture supporting burst grooming with $Q$ ports and $W$ data channels and one control channel on each port.

TABLE I
SUMMARY OF PARAMETER DEFINITIONS

| Parameter | Description |
|---|---|
| $\mathbf{G} = \{b_0, b_1, ...\}$ | Groomed data burst set, $b_0$ is the timed-out sub-burst |
| $\lvert\mathbf{G}\rvert$ | Combined length of sub-bursts groomed together |
| $L_G$ | The number of sub-bursts groomed in a single burst |
| $b_i, L_{b_i}$ | Sub-burst $b_i$ with length $L_{b_i}$ |
| $G^{\mathrm{MAX}}$ | Max. number of sub-bursts allowed to be groomed |
| $T_e$ | Max. tolerable end-to-end delay for IP packet class |
| $L^{\mathrm{MIN}}$ | Minimum required burst size |
| $\delta_{b_0}$ | Stack time of sub-burst $b_0$ (Sec. III-A) |
| $\Delta(b_i, b_0)$ | Route deflection distance (Sec. III-C) |
| $\Psi(b_i, \mathbf{G})$ | Relative routing and padding overhead (Sec. III-C) |
| $\Upsilon(b_i, b_0)$ | Relative routing overhead of $b_i$ (Sec. III-C) |
| $S_{b_i}, D_{b_i}$ | Edge node source and destination of burst $b_i$ |
| $H_p(S_{b_i}, D_{b_i})$ | Shortest hop distance between $(S_{b_i}, D_{b_i})$ |
| $\rho_{net}, \rho_{act}$ | Average IP and data burst traffic load, respectively |

Fig. 2 shows the basic architecture of an edge node supporting data-burst grooming. An ingress edge node, which generates and transmits data bursts to core nodes, performs the following operations: 1) burst assembly—aggregating incoming IP packets with the same edge-node destination (or other similar characteristics) in a virtual queue (VQ); 2) sub-burst grooming—combining multiple sub-bursts from different VQs into a single burst; 3) burst scheduling—attaching padding and preamble (framing) overhead to the bursts and scheduling them for transmission on an appropriate channel; and 4) burst header packet (BHP) generation—constructing the header packets and transmitting them prior to their corresponding data bursts.

In the egress path, as shown in Fig. 2, an egress edge node performs two basic functions: burst disassembly and IP routing. Upon receiving a data burst, the edge node initially disassembles the burst. The extracted sub-bursts, which need to be retransmitted to the downstream nodes are sent to the assembly unit, while the remaining sub-burst will be directed to the IP-routing unit. The IP-routing unit is a line card responsible for disassembling each sub-burst and sending its embedded packets to appropriate IP routers in the access layer of the network.

## III. BURST GROOMING

In this section, we first introduce some basic definitions and formulate the edge-node grooming problem in the OBS network, and then describe our proposed grooming algorithms. A summary of notations used in the following sections is provided in Table I.

### A. Data-Burst Grooming

We denote a sub-burst $i$ as $b_i$. Each sub-burst $b_i$ consists of multiple IP packets with the same edge-node destination and can be characterized by its edge-node source, destination, and length: $S_{b_i}$, $D_{b_i}$, and $L_{b_i}$. As soon as an IP packet with destination $D_{b_i}$ arrives to a node, a timer is set for sub-burst $b_i$. The sub-burst will be released when it is timed out. The time-out value is based on the maximum end-to-end delay tolerance that IP packets can tolerate within the OBS network, denoted by $T_{\mathrm{e}}$.[1] Therefore, the time-out value for data bursts in each VQ is bounded by the difference between $T_{\mathrm{e}}$ and the sum of OBS source–destination propagation delay and node processing delays, which includes the burst disassembly time at the destination node. In addition to the aforementioned parameters, each sub-burst $b_i$ has a remaining slack time, denoted as $\delta_{b_i}$. The remaining slack time is defined as the remaining tolerable end-to-end delay the sub-burst can tolerate before it reaches its destination.

We represent a groomed data burst by $\mathbf{G} = \{b_0, b_1, b_2, \ldots\}$, which is constructed by aggregating a number of sub-bursts with different destinations. We consider the first element (sub-burst) in the grooming set $(b_0)$ as the timed-out sub-burst, which must be routed on a single-hop shortest path. Hence, the first hop for all sub-bursts in $\mathbf{G}$ will be the node corresponding to the destination $D_{b_0}$. In our notation, $|\mathbf{G}|$ and $L_G$ indicate the number of sub-bursts groomed together and their combined length, respectively. Clearly, if $|\mathbf{G}| = 1$, no grooming has been performed and $L_G = L_{b_0}$. Furthermore, we refer to $G^{\mathrm{MAX}}$ as the maximum number of sub-bursts that are allowed to be groomed together prior to transmission, hence, $|\mathbf{G}| \leq G^{\mathrm{MAX}}$.

We define the hop delay as the delay time imposed on an incoming sub-burst due to electronic processing. In our study, we only consider the maximum hop delay, expressed as $T_{\mathrm{h}}$, and assume it is the same for all nodes. It is clear that the timed-out sub-burst can only be groomed with any other sub-burst $b_i$, whose remaining slack time satisfies the following expression:

$$T_{\mathrm{p}}(S_{b_0}, D_{b_0}) + T_{\mathrm{p}}(D_{b_0}, D_{b_i}) + T_{\mathrm{h}} \leq \delta_{b_i} \leq T_{\mathrm{e}}. \quad (1)$$

In the above expression, $T_{\mathrm{p}}(s, d)$ is the propagation delay from node $s$ to node $d$. Note that $\delta_b$ for any given sub-burst is bounded by $T_{\mathrm{e}}$.

When $\mathbf{G}$ reaches its first destination node $D_{b_0}$, sub-burst $b_0$ is dropped. Then, each remaining sub-burst $b_i$, in the grooming set $\mathbf{G}$, is directed to its proper VQ and its slack time is reduced by $T_{\mathrm{h}} + T_{\mathrm{p}}(S_{b_0}, D_{b_0})$. Incoming sub-bursts may be aggregated with the existing IP packets waiting in the corresponding VQ. In this case, the remaining slack time of the combined sub-burst is set to the remaining slack time of the earliest packet in the queue.

We illustrate the above concepts using the example shown in Fig. 3. In this example, the sub-burst at node 1 going to node 3 is timed out and it is groomed with another sub-burst

[1] In general, depending on their class of service, IP packets can have different $T_{\mathrm{e}}$ values. In this paper, we assume all IP packets belong to the same class of service and hence, the same $T_{\mathrm{e}}$ value can be applied to all IP packets.

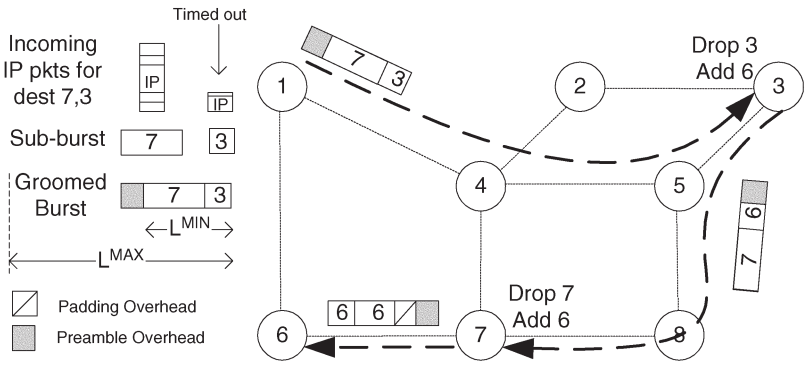


Fig. 3. Simple network carrying groomed data bursts.

with destination node 7, in order to meet the minimum-length requirement. At node 3, the sub-burst with destination node 3 is dropped. The remaining sub-burst going to node 7 will be groomed with another sub-burst with destination node 6. At node 7, the sub-burst going to node 6 is sent to the proper VQ and combined will all existing IP packets in the queue. When the timer is expired, the combined sub-burst going to node 6 must be transmitted. In this case, since the minimum length is not met, padding overhead is added.

When a sub-burst $b_0$ is timed out, the burst-grooming algorithm finds the appropriate $\mathbf{G}(b_0 \in \mathbf{G})$ among all possible grooming combinations. Selection of the grooming set is based on the optimization objective of the grooming algorithm. Aggregating multiple sub-bursts reduces the padding overhead, which in turn can improve the blocking probability. However, this can potentially result in routing the groomed sub-bursts over longer physical paths. This phenomena, referred to as the routing overhead, can impact the network throughput.

For example, consider Fig. 3, where at node 1, the timed-out sub-burst going to node 3 is groomed with the sub-burst going to node 7. We denote the shortest physical-hop distance between node pair $(s, d)$ by $H_{\mathrm{p}}(s, d)$. In this case, the sub-burst going to node 7, will be traveling over $H_{\mathrm{p}}(1, 3) + H_{\mathrm{p}}(3, 7) = 3 + 3 = 6$ physical hops, whereas the shortest physical-hop distance between node 1 and node 7 is 2: $H_{\mathrm{p}}(1, 7) = 2$. This example demonstrates that a simple greedy aggregation of sub-bursts can have adverse effects. Consequently, an effective grooming policy must minimize both the padding and the routing overhead, while minimizing additional hop delay.

It is evident in the above example that a potential drawback of burst grooming is the increase in number of electrical-to-optical converters/transmitters as incoming groomed sub-bursts must be retransmitted from intermediate nodes to their final edge-node destinations. Furthermore, burst grooming can result in higher buffering requirements at intermediate nodes. We defer these issues until later studies.

### B. Problem Formulation

In an OBS mesh network, data-burst grooming can be performed at the edge node. In this case, each individual edge node must decide how to aggregate individual sub-bursts with durations smaller than the minimum-length requirement, in order to optimize the throughput and reduce the probability of burst dropping. Hence, we can formulate the data-burst-grooming problem at the edge node as follows. Given the entire network information (including the physical network topology and full routing knowledge between all node pairs), the minimum required data-burst duration, the maximum end-to-end delay that each class of IP packet can tolerate, and a timed-out sub-burst with a length smaller than the minimum required length that has timed out, find any available sub-burst $b_i$ that can be aggregated with the timed-out sub-burst $b_0$ in order to minimize blocking probability.

We consider the following assumptions: All edge nodes have full grooming capability and equipped with full wavelength converters; all incoming IP packets have arbitrary lengths and a single destination; data bursts with durations shorter than the minimum burst length requirement will be subject to padding overhead; and all IP packets in a VQ must be transmitted together. In addition, in this study, we focus on networks with low IP-traffic arrival rate; thus, only a timer-based triggering scheme is assumed. We assume source routing, where the source edge node knows the entire path for all sub-bursts.

### C. Description of Grooming Algorithms

An intuitive approach to reduce IP-packet blocking probability is to develop effective grooming algorithms in order to reduce overall network overhead. The efficiency of the grooming algorithm can be affected by several parameters, including the number of sub-bursts that can be groomed together, the accumulated length of the groomed sub-bursts, and the way groomed sub-bursts with different destinations are routed. These parameters can have conflicting impacts under different network conditions. For example, at moderate loads, having fewer constraints on the above parameters may considerably reduce the network overhead, resulting in higher network throughput. On the contrary, at higher loads, asserting no constraints on the above parameters may notably alter the traffic characteristics and increase traffic burstiness, resulting in higher packet blocking.

We distinguish grooming algorithms by the way the source node calculates the padding and routing overheads due to burst grooming. Since the source node has no knowledge about the traffic between other node pairs, its padding-overhead calculations are based on worst case local estimations. In our study,

```
Initialization: G = {b_0}, S = {b_1, ..., b_i, ...}
while L_G < L^MIN, |G| < G^MAX and S ≠ ∅
    - Select b_i ∈S with the largest length such that
      δ_{b_i} satisfies eqn. (1) and ϒ(b_i, b_0) = 1
    - if b_i exists, move b_i from S to G
      update L_G and |G|
    - else S = ∅
end while
```

Fig. 4. NoRO algorithm.

```
Initialization: G = {b_0}, S = {b_1, ..., b_i, ...}
while L_G < L^MIN, |G| < G^MAX and S ≠ ∅
    - Select b_i ∈S with smallest Ψ(b_i, G) < 1
      and largest length such that δ_{b_i} satisfies eqn. (1)
      and Δ(b_i, b_0) <max. allowable route deflection
    - if b_i exists, move b_i from S to G
      update L_G and |G|
    - else S = ∅
end while
```

Fig. 5. MinTO algorithm.

we consider two grooming algorithms: No routing overhead (NoRO) and Minimum total overhead (MinTO).

*NoRO Algorithm:* The main objective in this algorithm is to ensure no routing overhead is added as sub-bursts are groomed together. In practice, this leads to routing all sub-bursts through their shortest cost path. Depending on the cost metric, the shortest cost path can be, for example, based on physical-hop distance or total link distance. Without loss of generality, in this paper, we consider the physical-hop distance between a node pair as the cost metric and refer to it as the shortest path. Note that NoRO does not distinguish between alternative shortest cost paths or interdependencies between them, as long as the cost remains the same.

The routing overhead for a sub-burst $b_i$, when groomed with $b_0$, can be quantified using the relative routing overhead $\Upsilon(b_i, b_0)$, which is defined as

$$\Upsilon(b_i, b_0) = \frac{H_{\rm p}(S_{b_0}, D_{b_0}) + H_{\rm p}(D_{b_0}, D_{b_i})}{H_{\rm p}(S_{b_0}, D_{b_i})}. \tag{2}$$

Sub-bursts $b_i$ and $b_0$ can only be groomed if $\Upsilon(b_i, b_0) = 1$, indicating that the destination of the timed-out sub-burst $D_{b_0}$ is on the shortest path to the destination of the groomed sub-burst $D_{b_i}$.

The details of the NoRO grooming algorithm as sub-burst $b_0$ with length $L_{b_0}$ is timed out are shown in Fig. 4. We denote all available sub-bursts (excluding $b_0$) in VQs as set $\mathbf{S} = \{b_1, \ldots, b_i, \ldots\}$. Note that the NoRO algorithm continues to groom $b_0$ with other sub-bursts until the combined length of the groomed burst $L_G$ is larger than $L^{\rm MIN}$, or the number of groomed sub-bursts $|\mathbf{G}|$ has exceeded $G^{\rm MAX}$.

*MinTO Algorithm:* The main objective of this algorithm is to reduce the combined routing and padding overheads by grooming multiple sub-bursts together. In practice, this leads to relaxing the NoRO constraint and allowing sub-bursts to travel through additional physical hops, when compared to their shortest path, before reaching their edge-node destinations. The combined routing and padding overheads for a sub-burst $b_i$, if groomed with a set of sub-bursts $\mathbf{G}$, can be quantified by the relative routing and padding overhead $\Psi(b_i, \mathbf{G})$, which is defined in (3) shown at the bottom of the page. In (3) $\hbar(x) = \max(L^{\rm MIN}, x)$ and $\mathbf{G} \cup b_i = \{b_0, b_i\}$, if $\mathbf{G} = \{b_0\}$. Note that the necessary condition for $b_i$ to be groomed with set $\mathbf{G}$, where $b_0 \in \mathbf{G}$, is $\Psi(b_i, \mathbf{G}) < 1$.

The additional physical hops a sub-burst $b_i$, when groomed with $b_0 \in \mathbf{G}$, must traverse before it reaches its edge-node destination is referred to as route-deflection distance, and we define it by

$$\Delta(b_i, b_0) = \left(H_{\rm p}\left(S_{b_0}, D_{b_0}\right) + H_{\rm p}\left(D_{b_0}, D_{b_i}\right)\right) - H_{\rm p}\left(S_{b_0}, D_{b_i}\right). \tag{4}$$

For example, referring to Fig. 3, the sub-burst going to node 7 from node 1 will have to tolerate a route-deflection distance of $\Delta = 6 - 2 = 4$.

Clearly, if $\Delta$ is limited to 0, no route deflection will be allowed, and all sub-bursts must traverse along their shortest paths. Note that having $\Delta = 0$ also implies no routing overhead: $\Upsilon = 1$.

Details of the MinTO grooming algorithm as sub-burst $b_0$ with length $L_{b_0}$ is timed out are shown in Fig. 5.

### D. Algorithm Analysis

In this section, we take a closer look at the MinTO algorithm and examine its performance under three different loading conditions. For simplicity, we assume that the maximum number of sub-bursts that can be groomed in a single burst is 2, $G^{\rm MAX} = 2$.

*1) Light Loads ($L_G$, $L_{b_0}$, $L_{b_i} < L^{\rm MIN}$):* In this case, (3) will be reduced to

$$\Psi(b_i, \mathbf{G}) = \frac{H_{\rm p}\left(S_{b_0}, D_{b_0}\right) + H_{\rm p}\left(D_{b_0}, D_{b_i}\right)}{H_{\rm p}\left(S_{b_0}, D_{b_0}\right) + H_{\rm p}\left(S_{b_0}, D_{b_i}\right)}. \tag{5}$$

Using (4), the necessary condition for $\Psi(b_i, \mathbf{G}) < 1$, indicating $b_i$ can be groomed with $b_0 \in \mathbf{G}$, is

$$\Delta(b_0, b_i) \le H_{\rm p}\left(S_{b_0}, D_{b_0}\right). \tag{6}$$

If the route-deflection distance is 0, $\Delta = 0$, under the low-loading assumption, (5) is reduced to

$$\Psi(b_i, \mathbf{G}) = \frac{H_{\rm p}\left(S_{b_0}, D_{b_i}\right)}{H_{\rm p}\left(S_{b_0}, D_{b_0}\right) + H_{\rm p}\left(S_{b_0}, D_{b_i}\right)} < 1. \tag{7}$$

$$\Psi(b_i, \mathbf{G}) = \frac{\hbar\left(L_G + L_{b_i}\right) \cdot H_{\rm p}\left(S_{b_0}, D_{b_0}\right) + \sum_{b_j \in \mathbf{G} \cup b_i}^{b_j \neq b_0} \hbar\left(L_{b_j}\right) \cdot H_{\rm p}\left(D_{b_0}, D_{b_j}\right)}{\sum_{b_j \in \mathbf{G} \cup b_i} \hbar\left(L_{b_j}\right) \cdot H_{\rm p}\left(S_{b_0}, D_{b_j}\right)} \tag{3}$$

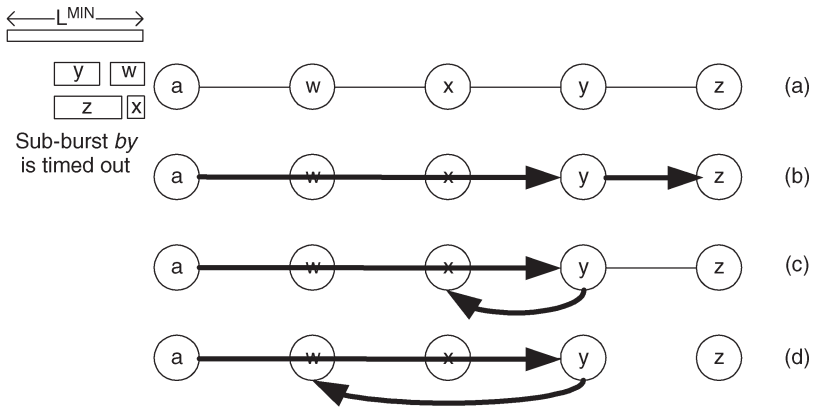


Fig. 6. Example of a five-node network where sub-burst $b_y$ going to node $y$ is timed out and it can be groomed with any one of the available sub-bursts: $b_w$, $b_x$, or $b_z$. Note that we assume the size of the grooming set is limited to $G^{\mathrm{MAX}} = 2$.

In this case, $\Psi(b_i, \mathbf{G})$ will be smaller in value for sub-bursts $b_i$ with shorter hop distance from $S_{b_0}$ to $D_{b_i}$: $H_{\mathrm{p}}(S_{b_0}, D_{b_i})$.

*2) Moderate Loads ($L_G \geq L^{\mathrm{MIN}},\ L_{b_0},\ L_{b_i} < L^{\mathrm{MIN}}$):* In this case, (3) will be reduced to

$$\Psi(b_i, \mathbf{G}) = \frac{H_{\mathrm{p}}\left(S_{b_0}, D_{b_0}\right) \cdot \left(\frac{L_G}{L^{\mathrm{MIN}}}\right) + H_{\mathrm{p}}\left(D_{b_0}, D_{b_i}\right)}{H_{\mathrm{p}}\left(S_{b_0}, D_{b_0}\right) + H_{\mathrm{p}}\left(S_{b_0}, D_{b_i}\right)}. \quad (8)$$

Rewriting the above expression in terms of $\Delta$, we obtain

$$\Delta(b_0, b_i) \leq H_{\mathrm{p}}\left(S_{b_0}, D_{b_0}\right)\left(\frac{L_G}{L^{\mathrm{MIN}}}\right). \quad (9)$$

Comparing (5) and (8) suggests that as long as $L_G < L^{\mathrm{MIN}}$ and $H_{\mathrm{p}}(D_{b_0}, D_{b_i}) < H_{\mathrm{p}}(S_{b_0}, D_{b_i})$, the timed-out sub-burst can be groomed with $b_i$. However, as the load increases and $L_G > L^{\mathrm{MIN}}$, fewer burst grooming can be expected.

*3) Higher Loads ($L_{b_i} \approx L_G \geq L^{\mathrm{MIN}},\ L_{b_0} < L^{\mathrm{MIN}}$):* In this case, (3) can be expressed as

$$\Psi(b_i, \mathbf{G}) = \frac{H_{\mathrm{p}}\left(S_{b_0}, D_{b_0}\right) + H_{\mathrm{p}}\left(D_{b_0}, D_{b_i}\right)}{H_{\mathrm{p}}\left(S_{b_0}, D_{b_0}\right) \cdot \left(\frac{L^{\mathrm{MIN}}}{L^G}\right) + H_{\mathrm{p}}\left(S_{b_0}, D_{b_i}\right)}. \quad (10)$$

Using the definition for $\Delta$, the above expression can be rewritten as

$$\Delta(b_0, b_i) \leq H_{\mathrm{p}}(S_{b_0}, D_{b_0})\left(\frac{L^{\mathrm{MIN}}}{L_G}\right) \quad (11)$$

where $L^{\mathrm{MIN}}/L_{\mathrm{G}} \leq 1$.

In the above discussion, we can clearly see that, in order to minimize routing and padding overhead, MinTO continuously attempts to groom multiple small sub-bursts whose destinations are closest to $D_{b_0}$. On the contrary, the NoRO algorithm mainly attempts to find the largest available sub-burst traveling along the timed-out sub-burst's path. An interesting observation in comparing (6), (9), and (11) is that as the network load increases, smaller route-deflection distance will be allowed, and hence, less grooming opportunities will be provided by MinTO.

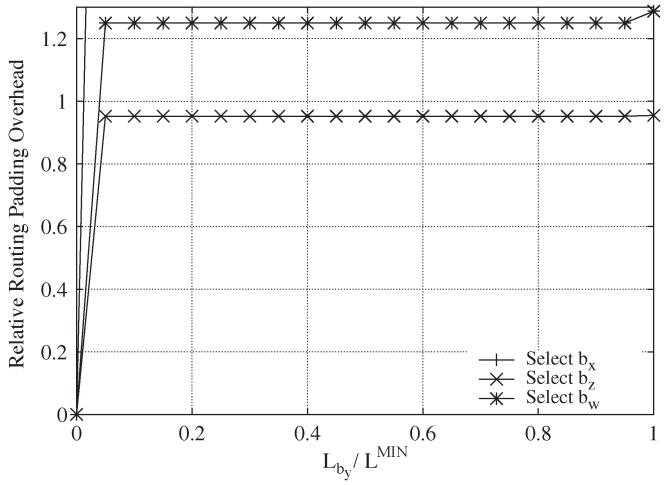


Fig. 7. Calculating the minimum routing and padding overhead for $\mathbf{G} = \{b_y, b_x\}$, $\{b_y, b_z\}$, and $\{b_y, b_w\}$ as a function of $L_{b_y}$.

Furthermore, the above relationships show that under certain network conditions, MinTO reduces the overall overhead in the network by introducing minimum routing overhead, $\Delta \neq 0$. This is different from NoRO, which aggressively attempts to search for the largest available sub-bursts to be groomed, regardless of the network load.

We illustrate the behavior of the NoRO and MinTO using the example shown in Fig. 6, where a five-node network with a single optical channel between each node pair is considered. We assume, at node $a$, that sub-burst $b_y$ is timed out and can be groomed with one of the available sub-bursts: $b_w$, $b_x$, or $b_z$. Using the NoRO algorithm, if we groom sub-burst $b_y$ with $b_z$, the lowest $\Upsilon$ value can be obtained. On the other hand, using the MinTO algorithm, the grooming choice changes depending on the length ratio of the available sub-bursts, namely, $b_w$, $b_x$, and $b_z$, over $L^{\mathrm{MIN}}$. For example, assuming the length of $b_z$ is much larger than $L_{b_x}$ and $L_{b_w}$, the value of $\Psi$ for $b_x$, $b_z$, and $b_w$ varies depending on the length of the timed-out sub-burst $b_y$, as shown in Fig. 7. It can be seen that for high values of $L_{b_y}/L^{\mathrm{MIN}}$, $\Psi(b_x, b_y)$ will be the smallest, and hence, $b_x$ will be selected to be groomed with $b_y$. This shows that under special circumstances, the MinTO algorithm prefers to groom with

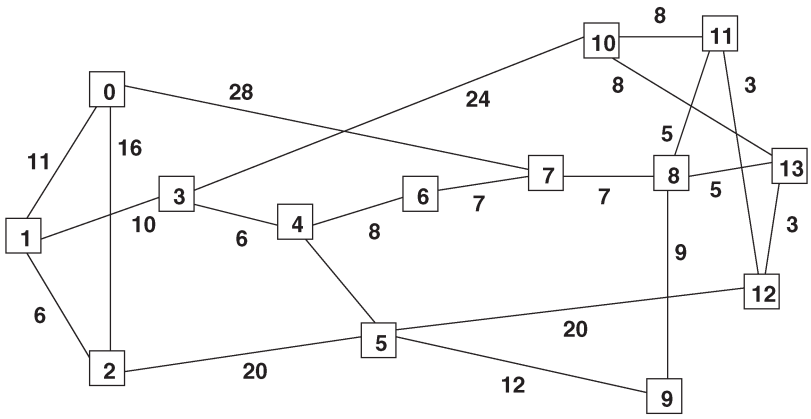


Fig. 8. NSF network with 14 nodes.

an available sub-burst, which results in larger route-deflection distance.

## IV. Performance Results

In this section, we present the simulation results obtained by implementing the NoRO and MinTO algorithms. We have chosen the National Science Foundation Network (NSFNet) backbone, shown in Fig. 8, as our test network. In this network, we assume each link is bidirectional with a fiber in each direction. Our simulation model was developed based on the following assumptions: IP-packet arrivals into the OBS network are Poisson, with $\lambda$ denoting their arrival rate, and they are uniformly distributed over all sender–receiver pairs; IP-packet length is fixed with 1250 B; the maximum end-to-end IP-packet delay tolerance is 50 ms; the switching time at the core node is 25 $\mu$s, requiring a minimum burst length of $L^{\mathrm{MIN}} = 250$ packets; each data burst can carry a maximum of 2500 IP packets; and each switch port has eight wavelengths, each of which has a transmission rate of 10 Gbit/s. We also assume all nodes support data-burst-grooming capacity and are equipped with full wavelength converters. We adopt the latest available unscheduled channel (LAUC) algorithm to schedule data bursts at the core nodes. Furthermore, we only consider a timer-based assembly and assume all sub-bursts can be groomed as long as their accumulated length is less than the minimum required length.

In our simulation study, we focus on traffic-load scenarios where sub-bursts typically time out before they reach their minimum required length, and hence, the mean burst length is less than $L^{\mathrm{MIN}}$. Throughout this section, we refer to the offered IP-packet load into the network as load, denoted by $\rho$. In our results, we focus on two basic performance metrics: IP-packet blocking probability and average end-to-end IP-packet delay. We define the former as the ratio of the number of IP packets that did not reach their destination over the total number of incoming IP packets.

In our C-based simulation model, we used confidence-interval accuracy as the controlling factor. For each case of interest, the simulation was run until a confidence-interval level of 90% was observed and an acceptably tight confidence interval were achieved. Calculations of the confidence interval were based on the variance within the collected observations [16].

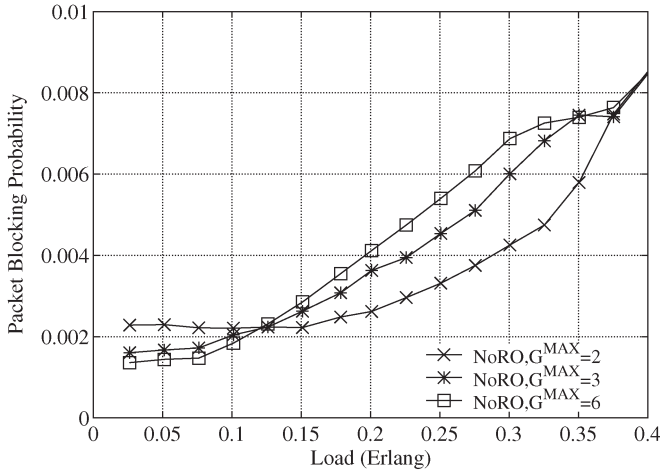


Fig. 9. Packet blocking probability using NoRO for $G^{\mathrm{MAX}} = 2$, 3, and 6.

### A. Characterizing the NoRO Algorithm

Fig. 9 shows the performance of the NoRO grooming algorithm for $G^{\mathrm{MAX}} = 2$, 3, and 6. As this figure suggests, under light loads ($\rho < 0.1$), allowing more sub-bursts to be groomed together results in lower packet blocking probability. Note that in our simulation, further increase in $G^{\mathrm{MAX}} > 6$ does not result in further performance improvement. This is because there is no more sub-burst available in the VQs.

At moderate loads ($0.1 \leq \rho \leq 0.36$), the IP-packet blocking probability increases for higher $G^{\mathrm{MAX}}$ values. As the load continues to increase, only a small percentage of sub-bursts will be shorter than $L^{\mathrm{MIN}}$, and hence, less grooming will take place. Eventually, at higher loads ($\rho > 0.36$), no grooming will be performed and, as Fig. 9 suggests, the blocking probability for different $G^{\mathrm{MAX}}$ values will become the same.

A surprising observation in Fig. 9 is that at moderate loads, as more sub-bursts are allowed to be groomed, the packet blocking probability increases. To understand this behavior, we examine two basic traffic characteristics, namely the padding-overhead ratio and traffic burstiness. The former is defined as the ratio of total padding overhead with and without grooming.

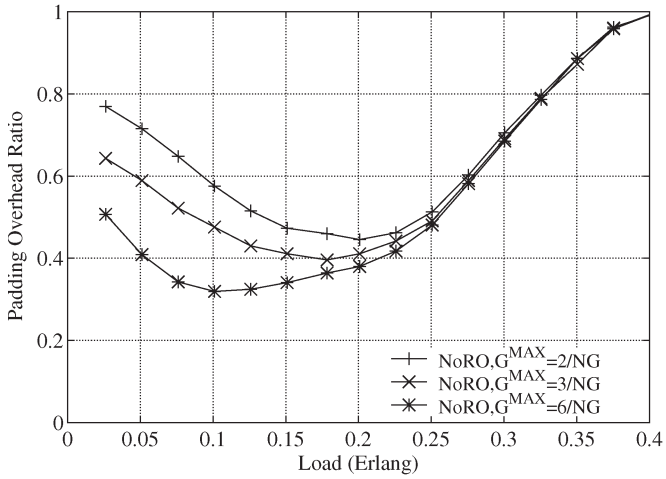


Fig. 10. Padding-overhead ratio over the OBS network using NoRO for $G^{\mathrm{MAX}} = 2$, 3, and 6.

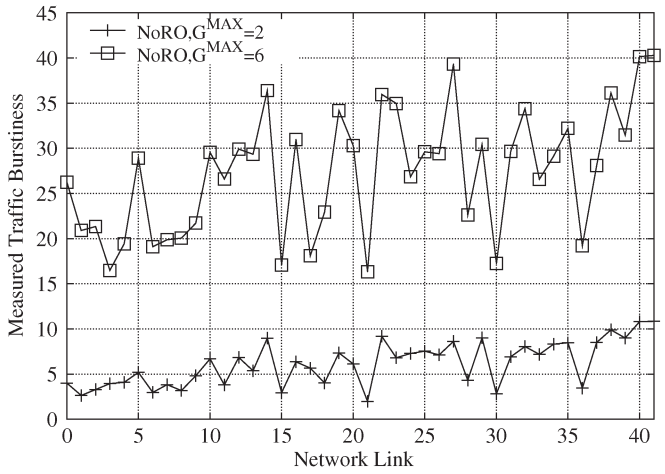


Fig. 11. Traffic burstiness measured on each switch egress port at $\rho = 0.25$ for $G^{\mathrm{MAX}} = 2$ and 6 using NoRO.

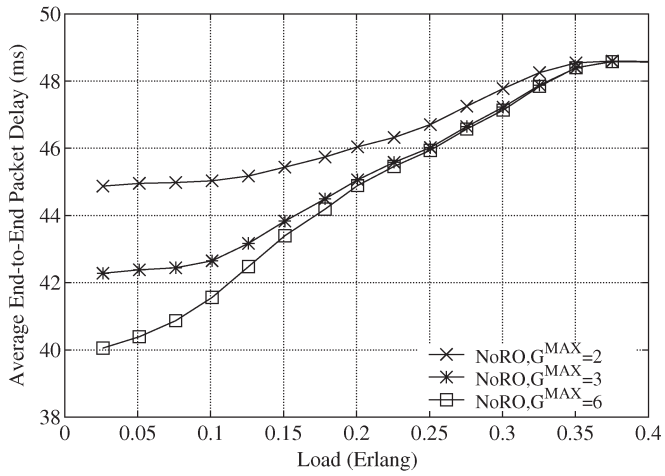


Fig. 12. Average end-to-end packet delay (milliseconds) using NoRO for $G^{\mathrm{MAX}} = 2$, 3, and 6.

Clearly, having smaller padding-overhead ratio implies higher link utilization compared to the case with no grooming.

We define traffic burstiness over each switch egress port $i$ (or unidirectional link between a node pair) as variation of burst load in time interval $s$ [17] and express it as

$$\beta_i(s) = \sqrt{\frac{\sigma_i^2(s)}{\mu_i^2(s)}} \tag{12}$$

where $\sigma_i^2(s)$ and $\mu_i(s)$ are the variance and mean of the burst load measured on link $i$ over some time period $s$, respectively. Burst load is defined as the product of burst arrival rate and mean burst length. Assuming the entire simulation period is $T_{\mathrm{sim}} = m \cdot s$, with $m$ discrete intervals of $s$, we will have $\mu_i(T_{\mathrm{sim}}) = \sum_{k=0}^{m-1} v_i^k(s)/m = E\{v_i(s)\}$, where $v_i(s)$ is the burst load measured on link $i$ over a time interval $s$. Similarly, $\sigma_i(T_{\mathrm{sim}}) = \sqrt{E\{v_i(s)^2\} - E^2\{v_i(s)\}}$.

Fig. 10 indicates that at light loads, having larger $G^{\mathrm{MAX}}$ value can considerably reduce the padding-overhead ratio. However, as the load increases, the padding-overhead ratios for different values of $G^{\mathrm{MAX}}$ tend to become the same and approach 1. In Fig. 10, note that, at moderate loads, the grooming algorithm results in minimum padding-overhead ratio. This is attributed to the fact that, at light loads, fewer and smaller sub-bursts are available to be groomed. On the other hand, at higher loads, fewer sub-bursts require padding and hence, grooming impact is minimized.

Fig. 11 compares the traffic burstiness, defined in (12), on unidirectional links 1 through 42 for $G^{\mathrm{MAX}} = 2$ and 6 when $\rho = 0.25$. This figure shows that as $G^{\mathrm{MAX}}$ increases from 2 to 6, the traffic burstiness increases as well. Increasing traffic burstiness results in higher peak burst load, leading to higher link congestion and blocking probability. Similar results can be obtained until $\rho \approx 0.36$, where grooming impact starts diminishing.

Based on the above traffic characteristics, we observe that at light loads, the padding-overhead ratio of $G^{\mathrm{MAX}} = 6$ is significantly lower than that of $G^{\mathrm{MAX}} = 2$. Consequently, the link utilization when $G^{\mathrm{MAX}} = 6$ will be higher, leading to lower blocking probability. As the load increases, the difference between the padding-overhead ratios for $G^{\mathrm{MAX}} = 2$ and 6 is reduced, and the traffic burstiness becomes the dominant factor. Hence, the blocking probability for $G^{\mathrm{MAX}} = 6$ will be higher than $G^{\mathrm{MAX}} = 2$.

The average end-to-end packet delay obtained from NoRO is shown in Fig. 12. As $G^{\mathrm{MAX}}$ increases, lower average delay can be achieved. This is due to the fact that by allowing higher number of sub-bursts to be groomed in a single burst, fewer sub-bursts will have to wait until they are timed out.

### B. Characterizing the MinTO Algorithm

The overall performance of MinTO, in terms of packet blocking probability and average end-to-end packet delay, follows similar trends we described for NoRO. A major issue with MinTO, however, is that it can potentially send some groomed sub-bursts through longer paths before reaching their destinations. Consequently, such sub-bursts will be more vulnerable to blocking. In fact, further simulation shows that at light loads, sub-bursts experience an average route-deflection distance of $\Delta \approx 1.9$. As the load increases, $\Delta$ tends to become

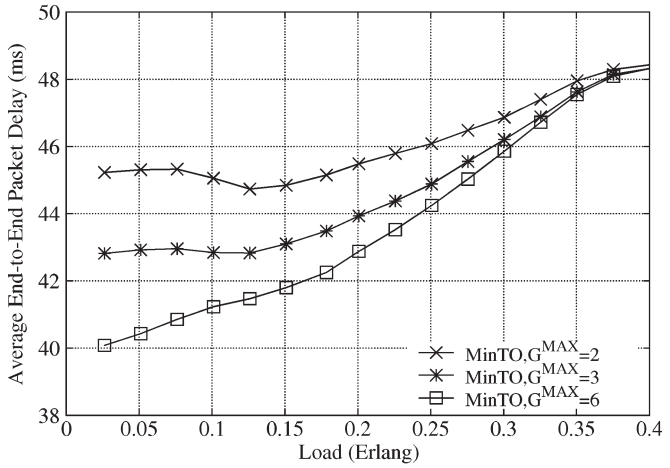


Fig. 13. Average end-to-end IP-packet delay (milliseconds) using MinTO for $G^{\mathrm{MAX}} = 2$, 3, and 6.

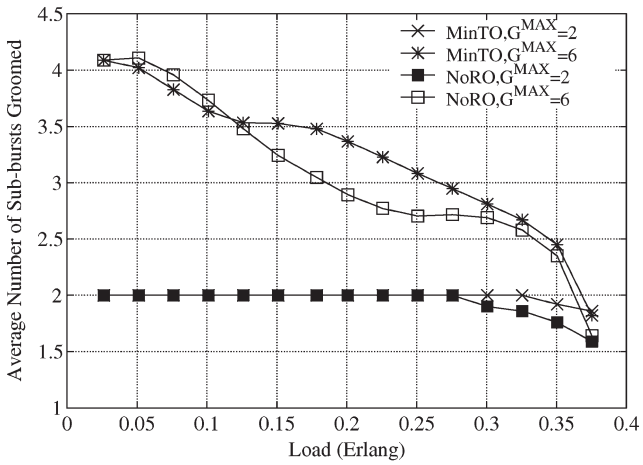


Fig. 14. Comparing the average number of sub-bursts groomed in a single burst using NoRO and MinTO for $G^{\mathrm{MAX}} = 2$ and 6.

smaller, which can be verified by comparing (5) and (8). This effect can also be observed in Fig. 13. When $G^{\mathrm{MAX}} = 2$, the average number of sub-bursts groomed together remains the same for all loads and the average end-to-end packet delay at $\rho = 0.05$ is slightly larger than when $\rho = 0.13$. However, when $G^{\mathrm{MAX}} = 6$, this effect is not evident because the average number of sub-bursts groomed together changes.

One way to avoid excessive route deflection is to impose an upper bound on the maximum route-deflection distance, for example, $\Delta \leq 1$. Our simulation results showed that under such constraint, at higher loads, slightly lower packet blocking can be achieved, which are consistent with our analysis in Section III-D. The tradeoff for such constraint is the higher average end-to-end packet delay due to limited grooming opportunities. In the remainder of this section, we consider MinTO, where $\Delta \geq 0$.

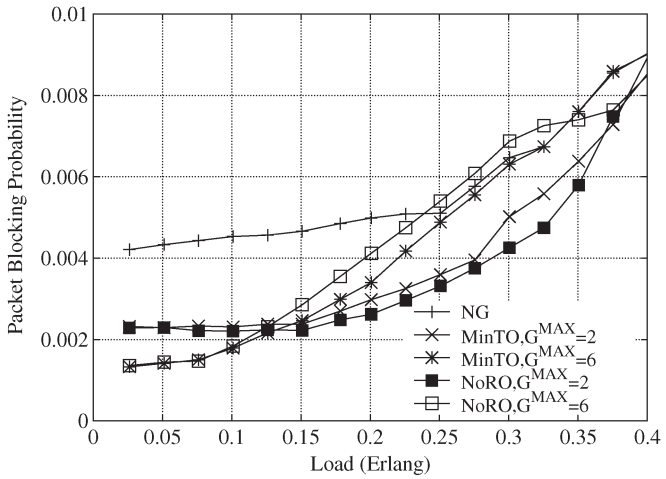


Fig. 15. Comparing the packet blocking probability using NoRO and MinTO for $G^{\mathrm{MAX}} = 2$ and 6.

### C. Grooming Algorithm Comparison

In this section, we compare the performance of NoRO and MinTO with the case when no grooming is applied. We start by examining the average number of sub-bursts groomed obtained by each algorithm as the load changes. Then, we demonstrate how the average number of sub-bursts groomed impacts the performance metrics.

Fig. 14 shows that, when $G^{\mathrm{MAX}} = 6$, at light loads, MinTO is less aggressive and provides fewer grooming opportunities compared to NoRO. This is due to the fact that in MinTO, burst grooming depends on the latest value of $\Psi(b_i, \mathbf{G})$ and the sub-bursts that have already been included in $\mathbf{G}$. As the load increases, there are more sub-bursts available. Hence, MinTO provides more grooming opportunities by relaxing the routing overhead constraint and allowing route-deflection distance. When $G^{\mathrm{MAX}} = 2$, the average number of sub-bursts groomed will be similar for both algorithms. However, NoRO is more aggressive because it tends to select the largest available sub-bursts for grooming, as described in Fig. 4.

Fig. 15 shows the packet blocking probability obtained by implementing the NoRO and MinTO for $G^{\mathrm{MAX}} = 1$, 2, and 6. At light loads, the more aggressive grooming approach, i.e., NoRO, with higher $G^{\mathrm{MAX}}$, provides slightly lower blocking probability. At higher loads, however, a more aggressive grooming approach results in higher traffic burstiness on links, as shown in Fig. 11. Consequently, NoRO with $G^{\mathrm{MAX}} = 2$ outperforms MinTO with $G^{\mathrm{MAX}} = 6$ in terms of packet blocking probability. Recall that even limited grooming, $G^{\mathrm{MAX}} = 2$, can considerably reduce the padding-overhead ratio, and hence, improve the overall blocking probability.

Fig. 16 shows that, in general, the average end-to-end packet delay due to grooming is much less than the case in which no grooming is implemented. Furthermore, the relative performance of MinTO and NoRO consistently follows the algorithm's grooming aggressiveness, as shown in Fig. 14. That is, more grooming results in lower end-to-end packet delay.

The aforementioned results demonstrate that MinTO and NoRO perform differently depending on the load. In general, grooming higher number of sub-bursts together can considerably improve the average end-to-end packet delay. Burst grooming can also improve packet blocking probability throughout the network at moderate loads. However, at higher loads, depending on network constrains, such as $L^{\mathrm{MIN}}$ and $T_{\mathrm{e}}$,

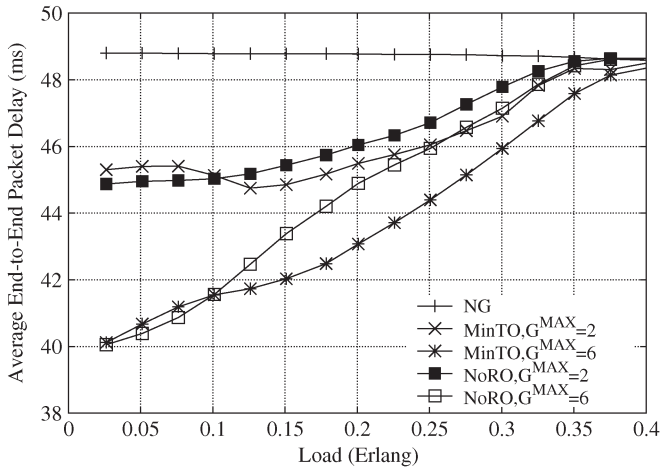


Fig. 16. Comparing the average end-to-end packet delay (milliseconds) using NoRO and MinTO for $G^{\mathrm{MAX}} = 2$ and 6.

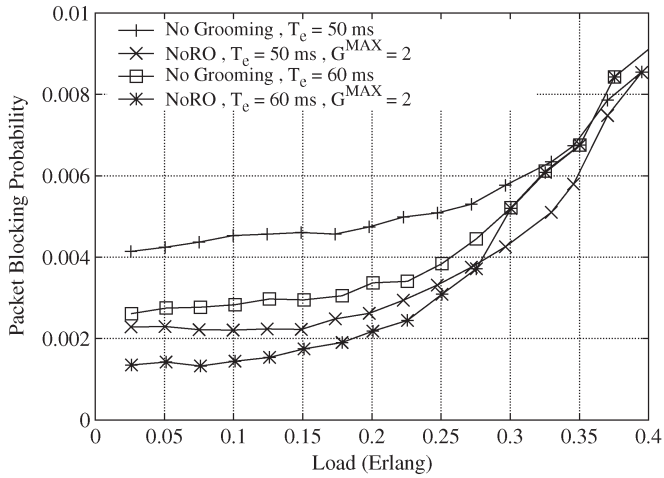


Fig. 17. Comparing the packet blocking probability using NoRO and no grooming for $G^{\mathrm{MAX}} = 2$, $T_{\mathrm{e}} = 50$ and 60 ms.

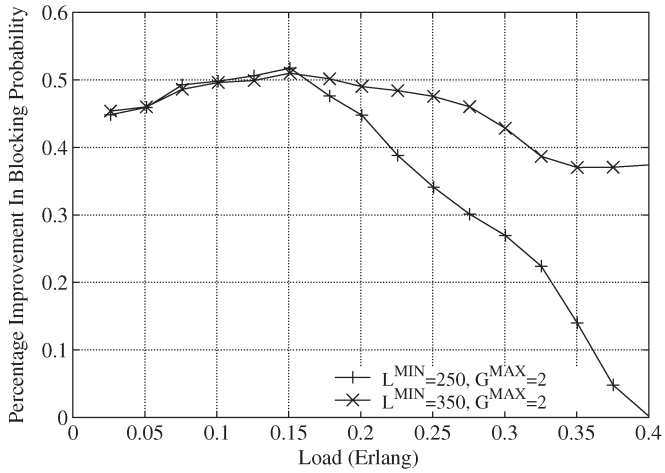


Fig. 18. Percentage improvement in packet blocking probability of NoRO compared to no grooming, assuming $G^{\mathrm{MAX}} = 2$ and $L^{\mathrm{MIN}}$ changes from 250 to 350.

limited grooming must be considered to reduce the padding overhead, and hence, to reduce the packet blocking probability. An adaptive approach, which aggressively grooms sub-bursts at low loads and gradually decreases $G^{\mathrm{MAX}}$ as the load increases, can improve the overall network performance in terms of both blocking probability and average end-to-end packet delay.

### *D. Performance of NoRO Under Different Network Parameters*

In this section, we investigate the performance of the grooming algorithms as the maximum tolerable end-to-end packet delay $T_{\mathrm{e}}$ and the minimum burst length requirement $L^{\mathrm{MIN}}$ vary. Since both NoRO and MinTO behave similarly under such changes, we only focus on the performance of the NoRO grooming algorithm.

In general, for a given switching time and load, as $T_{\mathrm{e}}$ decreases, data bursts time out earlier, and hence, the average number of IP packets aggregated in each burst tends to become smaller. Consequently, more padding overhead will be generated and higher packet blocking probability is expected. Fig. 17 shows the packet blocking probability using NoRO with $G^{\mathrm{MAX}} = 2$ for $T_{\mathrm{e}} = 50$ and 60 ms. This figure suggests that, for a given load and switching time, NoRO becomes more effective in terms of lowering the packet blocking probability as $T_{\mathrm{e}}$ is reduced. This implies that burst grooming can particularly benefit IP packets with lower end-to-end delay tolerance.

Fig. 18 shows the percentage performance improvement of NoRO with $G^{\mathrm{MAX}} = 2$ compared to when no grooming is implemented, as $L^{\mathrm{MIN}}$ changes from 250 to 350. As $L^{\mathrm{MIN}}$ increases, NoRO becomes more effective in terms of lowering the blocking probability for higher loads. Similarly, our simulation results confirmed that, for a given load and $T_{\mathrm{e}}$, as $L^{\mathrm{MIN}}$ increases, burst grooming can become more effective in terms of lowering the average end-to-end packet delay.

We also examined the impact of burst grooming when no wavelength converters were used. The results demonstrate that burst grooming provides lower blocking probability compared to the case with no grooming, when no wavelength converters are utilized. This is because burst grooming allows the incoming groomed sub-bursts, which have been reassembled, to be retransmitted on any available channel. Such advantage is diminished when all nodes have wavelength converters. In terms of average end-to-end packet delay, having wavelength converters appears to have no impact.

Based on the above results, it can be concluded that burst grooming can particularly be advantageous for networks that are constrained by cost (e.g., having no wavelength converters) or technology (e.g., having core nodes with slow switching times) and which carry on-demand traffic with relatively low arrival rate sub-bursts.

## V. Conclusion

In this paper, we discussed the problem of data-burst grooming in OBS networks. The main motivation for this study is improving network performance when the sub-bursts have low arrival rate, and the core node's switching time is larger than the average size of sub-bursts. Under such assumptions, sub-bursts will time out before they reach their minimum required length,

and hence, padding overhead must be added. We developed two grooming algorithms, namely MinTO and NoRO, which aggregate multiple small sub-bursts together in order to reduce the padding overhead, while minimizing any added routing overhead.

Through a comprehensive simulation study, we investigated the performance of the MinTO and NoRO algorithms in terms of packet blocking probability and average end-to-end packet delay. Our results show that, in general, the proposed grooming algorithms can improve the performance when compared with the case with no grooming. However, careful considerations must be given to loading conditions and the number of sub-bursts allowed to be groomed together. This is due to the fact that they alter the network traffic characteristics negatively and make the traffic more sporadic.

One area of future work would be to extend the proposed burst-grooming framework such that it can support service differentiation and QoS. As we mentioned before, two potential drawbacks of burst grooming are increase in number of electrical-to-optical converters/transmitters and additional buffering requirements. Further studies are required to examine such cost increases. In addition, analyzing the cost–performance comparison between two networks, one with burst-grooming capability but no wavelength converters and the other with wavelength converters but no grooming capability, can also be interesting. Another open problem to study is the data-burst grooming under static-traffic scenario, where the average traffic between each node pair is known in advance.

## REFERENCES


[1] C. Qiao and M. Yoo, "Optical burst switching (OBS)—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. Vanderhoute, 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] F. Farahmand, V. M. Vokkarane, and J. P. Jue, "Practical priority contention resolution for slotted optical burst switching networks," in *Proc. IEEE/SPIE 1st Int. Workshop Optical Burst Switching (WOBS)*, Dallas, TX, Oct. 2003.

[5] S. Hardy. (2000, May). "All-optical-switching groundswell builds," *Lightwave* [Online]. 17(6), pp. 45–47. Available: http://lw.pennnet.com/

[6] 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 Optical Networking and Communications (OptiComm)*, Boston, MA, Jul. 2002, vol. 4874, pp. 125–136.

[7] M. Casoni, E. Luppi, and M. Merani, "Impact of assembly algorithms on end-to-end performance in optical burst switched networks with different QoS classes," in *Proc. IEEE/SPIE 3rd Workshop Optical Burst Switching*, San Jose, CA, Oct. 2004.

[8] X. Yu, Y. Chen, and C. Qiao, "A study of traffic statistics of assembled burst traffic in optical burst switched networks," in *Proc. SPIE Optical Networking and Communications (OptiComm)*, Boston, MA, Jul. 2002, vol. 4874, pp. 149–159.

[9] S. Oh and M. Kang, "A burst assembly algorithm in optical burst switching networks," in *Proc. Optical Fiber Communication Conf. and Exhibit (OFC)*, Anaheim, CA, Mar. 2002, pp. 771–773.

[10] R. Dutta and G. N. Rouskas, "Traffic grooming in WDM networks: Past and future," *IEEE Network*, vol. 16, no. 6, pp. 46–56, Nov.–Dec. 2002.

[11] P. J. Lin and E. Modiano, "Traffic grooming in WDM networks," *IEEE Commun. Mag.*, vol. 39, no. 7, pp. 124–129, Jul. 2001.

[12] O. Gerstel and R. Ramaswami, "Cost effective traffic grooming in WDM rings," *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 618–630, Oct. 2000.

[13] X. Zhang and C. Qiao, "An effective and comprehensive approach to traffic grooming and wavelength assignment in SONET/WDM rings," *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 608–617, Oct. 2000.

[14] K. Zhu and B. Mukerjee, "Traffic grooming in an optical WDM mesh network," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 1, pp. 122–133, Jan. 2002.

[15] S. Sheeshia and C. Qiao, "Burst grooming in optical-burst-switched networks," in *Proc. IEEE/SPIE 1st Workshop Traffic Grooming WDM Networks*, San Jose, CA, Oct. 2004.

[16] R. Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. New York: Wiley Interscience, Apr. 1991.

[17] Y. Arakawa, N. Yamanaka, and I. Sasase, "Performance of optical burst switched WDM ring network with TTFR system," in *Proc. 1st IFIP Optical Networks & Technologies Conf. (OpNeTec)*, Pisa, Italy, Oct. 2004, pp. 95–102.


**Farid Farahmand** received the Ph.D. degree in electrical engineering from the University of Texas at Dallas in 2005.

In 1993, he was a Hardware Design Engineer at Alcatel, USA. In January 2000, he moved to Alcatel Corporate Research and was involved with the development of terabit optical routers. He is currently an Assistant Professor at the Department of Computer Electronics and Graphics Technology, Central Connecticut State University, New Britain. His research interests include high-speed packet switching and all-optical networks focusing on architecture and protocol designs for optical burst-switched networks.

**Qiong Zhang** (S'02–M'05) received the B.S. degree in computer science from Hunan University, China, in 1999, and the M.S. and Ph.D. degrees in computer science from the University of Texas at Dallas in 2000 and 2005, respectively.

In 2005, she joined the Arizona State University West, Phoenix, where she is currently an Assistant Professor at the Department of Mathematical Sciences and Applied Computing. Her current research interests include optical networks, network security, and grid computing.

**Jason P. Jue** (S'95–A'99–M'00–SM'04) 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.

He is currently an Associate Professor at the Department of Computer Science, University of Texas, Dallas. His research interests include optical networks, network control and management, and network survivability.