A Comparison of the JIT, JET, and Horizon Wavelength Reservation Schemes on A Single OBS Node *

Jing Teng, George N. Rouskas
Department of Computer Science, North Carolina State University, Raleigh, NC 27695-7534

ABSTRACT

We present a detailed analysis of the JIT, JET, and Horizon wavelength reservation schemes for optical burst switched (OBS) networks. Our analysis accounts for several important parameters, including the burst offset length, and the optical switching and hardware processing overheads associated with bursts as they travel across the network. The contributions of our work include: **(i)** analytical models of JET and Horizon (on a single OBS node) that are more accurate than previously published ones, and which are valid for general burst length and offset length distributions; and **(ii)** the determination of the regions of parameter values in which a more complex reservation scheme reduces to a simpler one. We compare the performance of the three wavelength reservation schemes on a single OBS node under various sets of parameter values. Our major finding is that, under reasonable assumptions regarding the current and future state-of-the-art in optical switch and electronic hardware technologies, the simplicity of JIT seems to outweigh any performance benefits of Horizon and JET.

1. INTRODUCTION

Optical burst switching (OBS) is a technology positioned between wavelength routing (i.e., circuit switching) and optical packet switching. All-optical circuits tend to be inefficient for traffic that has not been groomed or statistically multiplexed, and optical packet switching requires practical, cost-effective, and scalable implementations of optical buffering and optical header processing, which are several years away. OBS is a technical compromise that does not require optical buffering or packet-level parsing, and it is more efficient than circuit switching when the sustained traffic volume does not consume a full wavelength. The transmission of each burst is preceded by the transmission of a `setup` (also referred to as burst header control) message, whose purpose is to inform each intermediate node of the upcoming data burst so that it can configure its switch fabric in order to switch the burst to the appropriate output port. An OBS source node does not wait for confirmation that an end-to-end connection has been set-up; instead it starts transmitting a data burst after a delay (referred to as *offset*), following the transmission of the `setup` message. We assume that OBS nodes have no buffers, therefore, in case of congestion or output port conflict, they may drop bursts.

OBS networks have received considerable attention recently, mainly through theoretical investigations. A number of wavelength reservation schemes have been proposed for OBS, including just-enough-time (JET) [1], Horizon [2], just-in-time (JIT) [3,4], and wavelength-routed OBS [5] which uses two-way reservations. The burst loss performance of OBS networks has been studied extensively using either simulation or simple analytical models [2,6–11]. Typically, an output port of an OBS node has been analyzed assuming Poisson arrivals and no buffering [7–10]. Under these assumptions, an output port can be modeled by a finite number of servers, each representing a wavelength, with no queue. Then, the probability that a burst destined to this output port is lost can be obtained from the Erlang-B formula. An output port can also be modeled as an $M/M/m/K$ queue by assuming Poisson arrivals and buffering [2, 11], where m is the number of wavelengths and $K - m$ is the capacity of the buffer. A similar model that accounts for multiple classes of bursts, each class characterized by a different offset length, was developed in [6].

The JIT protocol is significantly simpler than either JET or Horizon, since it does not involve complex scheduling or void filling algorithms; therefore, it is amenable to hardware implementation [12]. On the other hand, previous studies have shown that JIT performs worse than either JET or Horizon in terms of burst loss probability. Indeed, given the sophisticated scheduling and void filling algorithms that JET and Horizon require,

* This work was supported by the Intelligence Technology Innovation Center under contract MDA904-00-C-2133.

the fact that these schemes should outperform JIT might seem a reasonable one at first thought. However, most of the existing studies ignore many important parameters such as the offset length, the processing time of `setup` messages, and the optical switch configuration time. Therefore, there is a need for more detailed studies in order to explore in depth the differences among the various wavelength reservation schemes, and to establish the regions of network operation where one scheme may outperform the others.

In this paper, we present models for an OBS node operating under the JET, JIT, and Horizon wavelength reservation schemes. The analytical models assume Poisson arrivals, but are valid for arbitrary burst length distributions and arbitrary offset length distributions. The models also account for the processing time of `setup` messages and the optical switch configuration times, and thus, are very general. One contribution made possible by our analysis is the characterization of the regions of network operation in which a more complex reservation scheme reduces to a simpler one (i.e., when JET reduces to Horizon, and Horizon to JIT). In Section 2 we describe the OBS node and introduce important system parameters used in our analysis. Section 3 provides a detailed description of the JIT, JET, and Horizon wavelength reservation schemes, and in Section 4 we present analytical models of a single OBS node under the three reservation schemes. In Section 5 we present numerical results, and in Section 6 we conclude the paper.

2. THE OBS NODE UNDER STUDY

We consider a node in an OBS network whose links support $W + 1$ wavelengths. One wavelength is used for signaling (i.e., it carries `setup` messages) and the other W wavelengths carry data bursts. The OBS node consists of two main components:

1. A *signaling engine*, which implements the OBS signaling protocol and related forwarding and control functions. To avoid bottlenecks in the control plane and to achieve operation at wire speeds, we assume that the signaling engine is implemented in hardware. (For example, the JITPAC hardware [12], which was developed by MCNC-RDI, implements the JIT signaling engine in FPGA.)
2. An *optical cross-connect* (OXC), which performs the switching of bursts from input to output. We assume that the OXC consists of a non-blocking space-division switch fabric, with no optical buffers. We also assume that the OXC has full conversion capability, so that an optical signal on any wavelength at any input port can be converted to any wavelength at any output port.

Whereas burst wavelengths are optically switched at the OBS node, the signaling wavelength is terminated at the node, the information it carries is converted to electronic form, and the resulting signal is passed to the signaling engine. The signaling engine decodes the electronic signal and processes each incoming message using the appropriate rules. Processing a signaling message may involve one or more actions, including: (1) the determination of a next hop switch for a burst; (2) the forwarding of signaling messages to other nodes; (3) the configuration of the OXC switching elements; and (4) the handling of exception conditions. The following parameters play an important role in the performance of the OBS node, and will be used in our analysis.

- T_{OXC} is the amount of time it takes the OXC to configure its switch fabric to set up a connection from an input port to an output port. In other words, T_{OXC} is the delay incurred between the instant the OXC receives a command from the signaling engine to set up a connection from an input port to an output port, until the instant the appropriate path within the optical switch is complete and can be used to switch a burst. In this study, we assume that this configuration delay is largely independent of the pair of input/output ports that must be connected, as well as of the state of the optical switch at the time the connection must be performed; this assumption is valid for optical switch technologies under development, including MEMS mirror arrays [13]. Therefore, we take T_{OXC} as a constant in our study.
- $T_{setup}(X)$ is the amount of time it takes an OBS node to process the `setup` message under reservation scheme X, where X can be any of JIT, JET, or Horizon. This amount of time is a function of the reservation scheme employed, however, for a given scheme X, we assume that $T_{setup}(X)$ is constant across all bursts. This is a reasonable assumption since processing of signaling messages will most likely be performed in hardware, and thus, the processing time can be bounded.

- $T_{offset}(X)$ is the offset value of a burst under reservation scheme X. The offset value depends on (1) the wavelength reservation scheme, (2) the number of nodes the burst has already traversed, and (3) other factors, such as whether the offset is used for service differentiation [11]. The primary consideration in the calculation of the offset value is to ensure that the first bit of the burst arrives at the destination node shortly after this node is ready to receive it (i.e., just after the destination has processed the `setup` message announcing the burst). The delay between the `setup` message and the first bit of the burst shrinks as the two propagate along the path to the destination. This is because the `setup` message encounters processing delays at each OBS node in the path, whereas the burst travels transparently in the optical domain. In addition, one must account for the switch setup delay T_{OXC} of the last OXC in the path.

  Let k be the number of OBS nodes in the path of a burst from source to destination. Based on the above observations, it is easy to see that the minimum offset value to guarantee that the burst will arrive at the destination immediately after the `setup` message has been processed is equal to:

$$T_{offset}^{(min)}(X) \quad = \quad kT_{setup}(X) \; + \; T_{OXC} \tag{1}$$

  We note that the actual offset length can take any value larger than the minimum one shown in the above expression; in fact, the models we develop later can account for offset lengths of arbitrary distributions.

3. WAVELENGTH RESERVATION SCHEMES FOR OBS NODES

The manner in which output wavelengths are reserved for bursts is one of the principal differentiating factors among OBS variants. We distinguish between two types of reservations: *immediate* and *delayed*.

3.1. Immediate Reservation (JIT)

Immediate reservation, exemplified by the Just-In-Time (JIT) family of OBS protocols [3, 4], works as follows:

> *an output wavelength is reserved for a burst immediately after the arrival of the corresponding* `setup` *message; if a wavelength cannot be reserved at that time, then the* `setup` *message is rejected and the corresponding burst is dropped.*

We illustrate the operation of JIT in Figure 1. Let t be the time a `setup` message arrives at some OBS node along the path to the destination user. As the figure shows, once the processing of the `setup` message is complete at time $t+T_{setup}$, a wavelength is immediately reserved for the upcoming burst, and the operation to configure the OXC fabric to switch the burst is initiated. When this operation completes at time $t+T_{setup}+T_{OXC}$, the OXC is ready to carry the burst. Note that the burst will not arrive at the OBS node under consideration until time $t+T_{offset}$. As a result, the wavelength remains idle for a period of time equal to $(T_{offset}-T_{setup}-T_{OXC})$. Also, since the offset value decreases along the path to the destination, the deeper inside the network an OBS node is located, the shorter the idle time between the instant the OXC has been configured and the arrival of the burst.

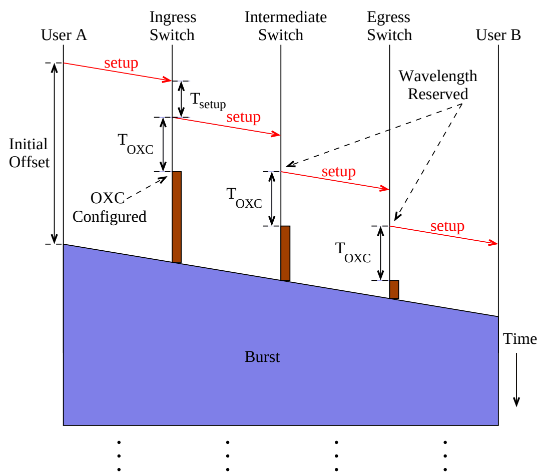

Figure 1. Immediate wavelength reservation

Figure 2 offers another perspective on how immediate reservation works, by considering the operation of a single output wavelength of an OBS node. Each such wavelength can be in one of two states: `reserved` or `free`. Figure 2 shows two successive bursts, i and $i+1$, *successfully transmitted* on the same output wavelength. As we can see, the `setup` message corresponding to the i-th burst arrives at the switch at time t_1, when we assume that the wavelength is `free`. This message is accepted, the status of the wavelength becomes `reserved` and, after an amount of time equal to the offset, the first bit of the optical burst arrives at the switch at time t_2. The last bit of the burst arrives at the switch at time t_3, at which instant the status of the wavelength is updated to `free`. Note that, any new `setup` message that arrives between t_1 and t_3 when the status of the wavelength is `reserved` is rejected, since the wavelength cannot be immediately reserved. The length of the interval, $t_3 - t_1$, during which new `setup` messages are rejected, is equal to the sum of the offset value and the length of burst i.

Suppose now that the next `setup` message for this wavelength arrives at time $t_4 > t_3$, while the wavelength is still `free`. Consequently, the burst corresponding to this message becomes the $(i+1)$-th burst to *successfully depart* on this wavelength; note that this burst may not be the $(i+1)$-th *arriving* burst, since some `setup` message(s) may have been rejected by the switch before time t_3. After an amount of time equal to the offset, the burst arrives at time t_5, and its transmission ends at time t_6, at which instant the wavelength becomes `free` again.

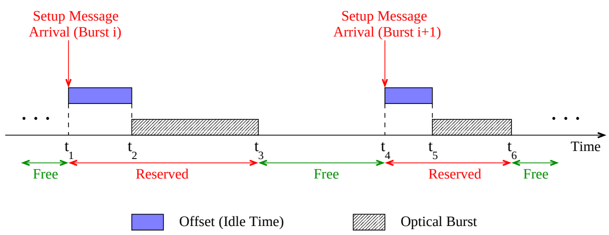

Figure 2. Operation and departure process of a wavelength with immediate reservation (JIT)

As Figure 2 illustrates, immediate reservation is simple. Time is divided into periods during which the wavelength is `reserved`, followed by periods during which it is `free`. The length of a `reserved` period is equal to the burst length plus the corresponding offset, while the length of a `free` period is equal to the time until the arrival of the next `setup` message. Also, service on each wavelength is first-come, first-served (FCFS), in the sense that bursts are served in the order in which their corresponding `setup` messages arrive at the switch.

3.2. Delayed Reservation

The Horizon [2] and JET [1, 14] protocols employ a delayed reservation scheme which operates as follows:

> *an output wavelength is reserved for a burst just before the arrival of the first bit of the burst; if, upon arrival of the* **`setup`** *message, it is determined that no wavelength can be reserved at the appropriate time, then the* **`setup`** *message is rejected and the corresponding burst is dropped.*

Figure 3 illustrates the operation of delayed reservation. Let us again assume that a `setup` message arrives at an OBS node at time t, in which case the first bit of the corresponding burst is expected to arrive at time $t + T_{offset}$. Assuming that the burst can be accepted, the `setup` message reserves a wavelength for the burst starting at time $t' = t + T_{offset} - T_{OXC}$. As shown in the figure, at time t', the OBS node instructs its OXC fabric to configure its switch elements to carry the burst, and this operation completes just before the arrival of the first bit of the burst. Thus, whereas immediate reservation protocols only permit a single outstanding reservation for each output wavelength, delayed reservation schemes allow multiple `setup` messages to make future reservations on a given wavelength (provided of course, that these reservations, i.e., the corresponding bursts, do not overlap in time). We also note that, when a burst is accepted, the output wavelength is reserved for an amount of time equal to the length of the burst *plus* T_{OXC}, in order to account for the OXC configuration time.

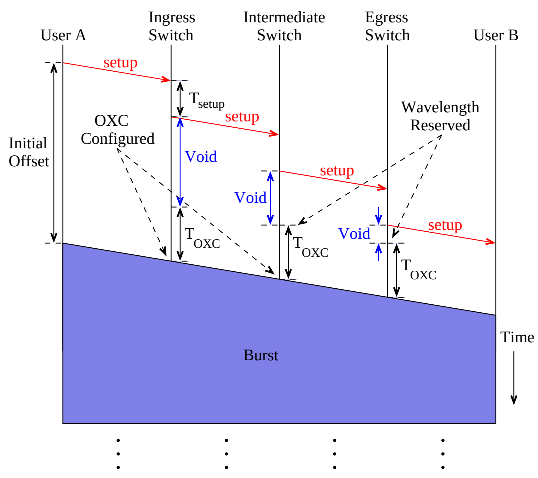

Figure 3. Delayed reservation

As we can see in Figure 3, a *void* is created on the output wavelength between time $t + T_{setup}$, when the reservation operation for the upcoming burst is completed, and time $t' = t + T_{offset} - T_{OXC}$, when the output wavelength is actually reserved for the burst. If the offset value T_{offset} is equal to the minimum value in expression (1), then the length of this void at some OBS node

x is equal to rT_{setup}, where r is the number of OBS nodes in the path from x to the destination of the burst. Consequently, the void created by a given burst decreases in size as the burst travels along its path.

Delayed reservation schemes can be further classified according to whether or not they employ specialized burst scheduling algorithms in an attempt to make use of the voids created by earlier `setup` messages, by transmitting bursts whose `setup` messages arrive later. Usually, such scheduling techniques are referred to as *void filling* algorithms.

3.2.1. Delayed Reservation *Without* Void Filling

Delayed reservation schemes, such as Horizon [2], that do not perform any void filling, are typically less complex than schemes with void filling, such as JET. The Horizon scheme takes its name from the fact that each wavelength is associated with a *time horizon* for burst reservation purposes. This time horizon is defined as "the earliest time after which there is no planned use of the channel (wavelength)". Under this scheme,

> *an output wavelength is reserved for a burst only if the arrival time of the burst is later than the time horizon of the wavelength; if, upon arrival of the* `setup` *message, it is determined that the arrival time of the burst is earlier than the smallest time horizon of any wavelength, then the* `setup` *message is rejected and the corresponding burst dropped.*

When a burst is scheduled on a given wavelength, then the time horizon of the wavelength is updated to the departure instant of the burst *plus* the OXC configuration time T_{OXC}. Consequently, under Horizon, a new burst can be scheduled on a wavelength only if the first bit of the burst arrives *after* all currently scheduled bursts on this wavelength have departed.

Figure 4 shows two bursts transmitted successively on a given wavelength out of an OBS node using the Horizon reservation scheme. The `setup` message of burst i arrives at the OBS node at time t_1, and the last bit of this burst leaves the node at time t_4. Since the OXC needs an amount of time equal to T_{OXC} to reconfigure its switching elements to perform a connection from another input port to this output wavelength, no new bursts can be scheduled on this wavelength until time $t_5 = t_4 + T_{OXC}$. Therefore, at time t_1, i.e., when burst i is accepted, t_5 becomes the time horizon of this channel.

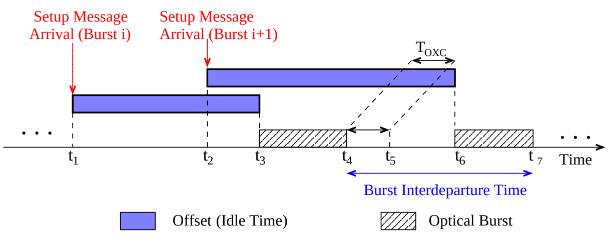

Figure 4. Departure process of a wavelength with delayed reservation and no void filling (Horizon)

Let us now suppose that, as Figure 4 illustrates, the `setup` message of burst $i+1$ arrives at the OBS node at time $t_2 > t_1$. The node uses the offset length information carried in the `setup` message to calculate that the first bit of this burst will arrive at time t_6. Since $t_6 > t_5$, burst $i+1$ is scheduled for transmission on this wavelength, and the time horizon is updated accordingly to $t_7 + T_{OXC}$, where t_7 is the instant the transmission of burst $i+1$ ends. This example shows that the offset of a burst (in this case, burst $i+1$) may overlap with the offset and/or transmission of another burst (i.e., burst i). However, bursts are scheduled in a strict FCFS manner determined by the order of arrival of their respective `setup` messages.

3.2.2. Delayed Reservation *With* Void Filling

JET [14] is the best known delayed wavelength reservation scheme that uses void filling. Under JET,

> *an output wavelength is reserved for a burst if the arrival time of the burst (1) is later than the time horizon of the wavelength,* ***or*** *(2) coincides with a void on the wavelength, and the end of the burst (plus the OXC configuration time T_{OXC}) occurs before the end of the void; if, upon arrival of the* `setup` *message, it is determined that none of these conditions are satisfied for any wavelength, then the* `setup` *message is rejected and the corresponding burst dropped.*

Note that, bursts which are accepted because their arrival and departure instants satisfy condition (2) above would have been rejected by an OBS node using Horizon. Consequently, JET is expected to perform better than Horizon in terms of burst drop probability. On the other hand, the void filling algorithm must keep track of, and search, the starting and ending times of all voids on the various wavelengths, resulting in a more complex implementation than either Horizon or JIT.

Figure 5 illustrates the void-filling operation of JET. The figure shows two bursts, A and B, which are both transmitted on the same output wavelength. The `setup` message for burst A arrives first, followed by the `setup` message for burst B. As we show in the figure, burst A has a long offset. Upon receipt of its `setup` message, the switch notes the later arrival of burst A, but does not initiate any connection within its cross-connect fabric. Once burst A has been accepted, a void is created, which is the interval of time until the arrival of the first bit of the burst at time t_6. Let us assume that at time t_2 when the `setup` message for burst B arrives, no other burst transmissions have been scheduled within this void.

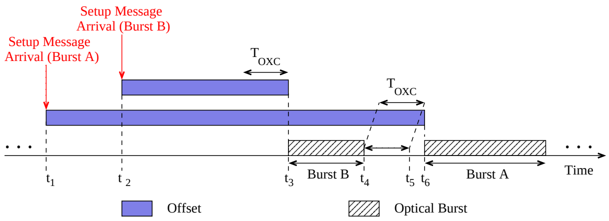

Figure 5. Non-FCFS service of a wavelength in an OBS node with delayed reservation and void filling (JET)

Upon the arrival of the `setup` message for burst B at time t_2, the switch notes that burst B will arrive *before* the arrival of burst A, and runs a void filling algorithm [15, 16] to determine whether it can accept the new burst. In order to accept the new burst, there must be sufficient time between the end of the transmission of burst B and the arrival of burst A for the switch to reconfigure its cross-connect fabric to accommodate burst A. For the scenario depicted in Figure 5, burst B is accepted, and it completes service before the arrival of the first bit of burst A. Since the `setup` message for burst B arrived after the `setup` message for burst A, this operation results in a non-FCFS service of bursts.

4. MODELS OF AN OBS NODE

In this section, we develop three analytical models for an output port p of an OBS node, one for each of the three reservation schemes JIT, JET, and Horizon. In our analysis, we make the following assumptions:

- `Setup` messages corresponding to bursts destined to output port p arrive at the OBS node according to a Poisson process with rate λ; this arrival rate is the total rate over all input ports. This assumption is made mainly for mathematical tractability, and is common in the OBS literature [2, 6–11].
- Burst lengths follow a general distribution with CDF $B(l)$ and Laplace transform $B^\star(s)$. We let $1/\mu$ denote the mean of the burst length distribution.
- Offset lengths follow a general distribution with CDF $G(z)$ and Laplace transform $G^\star(s)$. We also let $\overline{T}_{offset}(X)$ denote the mean offset length under reservation scheme X.
- An output wavelength is reserved for a given burst for a period of time that is *larger than* the length of the burst; at a minimum, the wavelength must be reserved for the duration of the burst length *plus* the OXC configuration time T_{OXC}, to allow for setting up the optical switch fabric to establish a connection from the input to the output port. Therefore, we define the *effective service time* of a burst as the amount of time that an output wavelength is reserved for the burst. As we shall see, the effective service time of the burst depends on the wavelength reservation scheme used.

We note that, while the burst arrival rate λ and the burst length distribution are *not* affected by the reservation scheme (JIT, JET, or Horizon), the offset length distribution is affected by the choice of reservation scheme.

Note that we have assumed that `setup` messages arrive as a Poisson process with rate λ. Let us now concentrate on the arrival process of the corresponding bursts, rather than that of the `setup` messages. The arrival time of a burst is the arrival time t of its `setup` message plus an offset, which is distributed according to a general distribution $G(z)$. One way of thinking about this burst arrival process is to assume that bursts arrive *at the same time* as their corresponding `setup` messages (i.e., as a Poisson process with rate λ), but they have to be served by a *fictitious* infinite server (i.e., an $M/G/\infty$ queue) before they enter the OBS node. The service time at this infinite server is distributed according to the CDF of the offset length, $G(z)$. As a result, the actual arrival of a burst to the OBS node is indeed the arrival time of its `setup` message plus an offset time distributed according to CDF $G(z)$. It is well-known that the departure process of an $M/G/\infty$ queue is a Poisson process with rate λ, the same as the arrival process. Therefore, burst arrivals to the OBS node are also Poisson with rate λ.

We note that the above $M/G/\infty$ model assumes *optimal* scheduling and void filling algorithms, in the sense that no burst is dropped if it can be carried by the switch; in practice, fast suboptimal algorithms may be used, in which case some bursts may be dropped even if they would be scheduled under an optimal algorithm. Furthermore, the $M/G/\infty$ model is an approximation since the underlying assumption is that the decision to accept or drop the burst is taken *at the moment the first bit of the burst arrives.* In other words, this model is exact only under the assumption that processing of `setup` messages and the OXC configuration takes zero time. In reality, the decision to accept or drop a burst is taken at the instant its `setup` message arrives, and if a `setup` message is rejected then the corresponding burst never arrives at the OBS node, resulting in a non-Poisson arrival process for bursts. However, the $M/G/\infty$ model is both conceptually simple and reasonably accurate, and we will make use of it in the analysis of some of the reservation schemes.

We model the output port of an OBS node as a multiple server loss system, and we use the Erlang-B formula to obtain the burst drop probability. The Erlang-B formula for an m-server system with traffic intensity ρ is:

$$Erl(\rho, m) \quad = \quad \frac{\rho^m/m!}{\sum_{i=0}^{m} \rho^i/i!} \tag{2}$$

In the following subsections, we determine accurate values for the intensity ρ under each reservation scheme. Since the loss probability in an m-server loss system is insensitive to the service time distribution, we use the Erlang-B formula above for any distribution of the effective service time of bursts.

4.1. A Model of JIT

In order to determine the effective service time of a burst under the JIT reservation scheme, let us refer again to Figure 2. We observe that, for a given burst, a wavelength is reserved for a length of time that is equal to the sum of two time periods. The duration of the first period is equal to the burst offset, and is distributed according to CDF $G(z)$ with a mean $\overline{T}_{offset}(JIT)$. The duration of the second period is equal to the burst length, and is distributed according to CDF $B(l)$ with a mean $1/\mu$. Consequently, the Laplace transform of the distribution of the effective service time of bursts is given by $G^\star(s)B^\star(s)$, with mean $1/\mu + \overline{T}_{offset}(JIT)$.

Based on these observations, an output port of an OBS node using JIT behaves as an $M/G/W/W$ loss system, where W is the number of wavelengths of the port. The traffic intensity $\rho(JIT)$ of the queue is:

$$\rho(JIT) \quad = \quad \lambda\left(\frac{1}{\mu} \; + \; \overline{T}_{offset}(JIT)\right) \tag{3}$$

and the burst drop probability is given by $Erl(\rho(JIT), W)$. We also note that, under the assumption that `setup` messages arrive as a Poisson process, the $M/G/W/W$ queue is an exact model for JIT. This model has been used in earlier studies, e.g., in [8], where, however, the assumption was made that burst (rather than `setup` message) arrivals are Poisson; in that case, the model is only approximate.

4.2. A Model of JET

The operation of an OBS node under the delayed reservation scheme is more complicated than under immediate reservation (i.e., JIT). Let us first consider the case in which void filling is employed [15, 16] when allocating a wavelength to a burst, as in the JET [14] reservation scheme. The difficulty in this case arises from two observations regarding burst transmissions on a given output wavelength. First, the offset of a given burst may overlap with the offset and/or transmission of one or more other bursts. Second, bursts are not necessarily served in an FCFS fashion. This overlap feature and resulting non-FCFS service were illustrated in Figure 5.

To overcome the difficulty introduced by the offset overlap and the non-FCFS service, let us concentrate on the departure process of a given output wavelength. In Figure 6, we show two bursts transmitted successively out of the switch on a given wavelength. We number the bursts in the order in which they *depart* the switch, so that burst $i+1$ is the first burst to be transmitted out on this wavelength after burst i; note that, due to the possibility for void filling, this may not be the order in which the `setup` messages of the two bursts arrived.

As Figure 6 illustrates, the first bit of burst i arrives at the OBS node at time t_1, and the last bit of the same burst leaves the switch at time t_2. Recall that the OXC needs an amount of time equal to T_{OXC} to reconfigure its switching elements to perform a connection from another input port to this output wavelength. Therefore, the switch cannot accommodate a new burst on this wavelength until time t_3, which is such that $t_3 = t_2 + T_{OXC}$. In fact, any `setup` message for a burst scheduled to arrive at the switch in the time interval between t_2 and t_3 would have been rejected by the switch scheduling algorithm. Therefore, we can think of a burst as occupying the channel not only during its transmission time (equal to its length), but also for an additional amount of time equal to T_{OXC}. Consequently, the effective service time of a burst follows a general distribution with Laplace transform $B^\star(s)e^{-sT_{OXC}}$ and mean $1/\mu + T_{OXC}$.

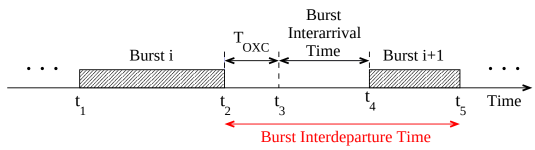

Figure 6. Departure process of a wavelength in an OBS node with delayed reservation and void filling (JET)

Based on the above observations, an output port p with W burst wavelengths can be modeled using the $M/G/W/W$ loss system. The traffic intensity $\rho(JET)$ for this system is given by

$$\rho(JET) \quad = \quad \lambda\left(\frac{1}{\mu} \; + \; \overline{T}_{OXC}\right) \tag{4}$$

and the probability of burst loss at the output port is given by the Erlang-B formula $Erl(\rho(JET), W)$. Note that, as we discussed above, the $M/G/W/W$ model for JET is approximate since it assumes a Poisson arrival process for bursts (or equivalently, that scheduling decisions are made at the instant a burst arrives, rather than at the time the `setup` message arrives). It also implies optimal scheduling decisions, when in practice a fast suboptimal algorithm may be used. Nevertheless, numerical results to be presented shortly indicate that this model is quite accurate.

4.3. A Model of Horizon

Similar to JET, the length of a wavelength reservation in Horizon is equal to the duration of a burst's transmission *plus* the OXC configuration time T_{OXC}. In order to account for the "no-void-filling" feature of Horizon compared to JET, we let the mean effective service time of bursts be equal to the mean wavelength reservation, $1/\mu + T_{OXC}$, *plus* a quantity $\Delta \geq 0$. That is, we use the following value for the traffic intensity of Horizon:

$$\rho(Horizon) \quad = \quad \lambda\left(\frac{1}{\mu} \; + \; \overline{T}_{OXC} \; + \; \Delta\right) \tag{5}$$

We first note that, when the values of the system parameters T_{OXC}, T_{setup}, and $1/\mu$ are such that no void filling is possible in the OBS network (refer to our discussion in Section 4.4), then obviously, $\Delta = 0$ and Horizon has the same burst drop probability as JET. However, if void filling is possible, then $\Delta > 0$, and the traffic intensity of Horizon is greater than that of JET (refer to expression (4)), resulting in higher burst drop probability. Using $\Delta > 0$ in (5) implies that the effective service time of bursts is larger than under JET. This increase in the effective service time of bursts has two consequences: first, voids become smaller, and second, the "larger" bursts will not fit within the "smaller" voids. Therefore, the essence of our approximation is to account for the lack of void filling by appropriately increasing the effective service time of bursts, and in turn, the traffic intensity. In Appendix A, we show how to estimate the value of Δ in expression (5).

4.4. Discussion

If we ignore the differences in the `setup` message processing time $T_{setup}(X)$ among the three reservation schemes X, then, in general, JET will result in the lowest burst drop probability, followed by Horizon and JIT. In practice, however, the relative performance of the three schemes depends on the actual values of certain system parameters. Let $X \equiv Y$ denote that reservation scheme X is *equivalent* to scheme Y (in the sense that both result in *the same* burst drop probability), and $X \approx Y$ denote that schemes X and Y result in *approximately the same* burst drop probability. Then, we can make the following observations.

- $T_{OXC} > kT_{setup} \Rightarrow$ JET $\equiv$ Horizon
  Referring to (1), if T_{OXC} is larger than the sum of `setup` message processing times, then no void filling may take place. This is because two OXC configuration operations are needed for a burst with a later `setup` message to fill a void created by a burst with an earlier `setup` message: one operation to switch the former burst, and one to switch the latter. The total time required for these operations is $2T_{OXC}$, while the void is at most equal to $T_{offset} = T_{OXC} + kT_{setup} \leq 2T_{OXC}$. Therefore, JET reduces to Horizon in this case.

- Minimum burst length $+T_{OXC} > kT_{setup} \Rightarrow$ JET $\equiv$ Horizon
  For similar reasons, if the minimum burst length plus the OXC configuration time T_{OXC} is larger than the sum of processing times, then no void filling is possible, hence JET reduces to Horizon.

- $T_{offset} =$ constant $\Rightarrow$ JET $\equiv$ Horizon
  If the offset value is constant (rather than equal to the minimum value in (1)), then no void filling is possible therefore JET reduces to Horizon. Note that a constant offset value may be of practical importance. For example, rather than estimating the number of hops to the destination in order to compute the minimum offset value according to (1), it may be desirable to set the offset to a large value that can accommodate any source-destination pair; this is similar to setting the TTL of an IP packet to a high value rather than one based on a given source-destination pair. Furthermore, if alternate routing algorithms are used to reduce the burst loss probability, as has been suggested in the literature, then the number of hops in the actual path may not be easy to estimate; a large constant offset value might then be appropriate.

- ($1/\mu \gg T_{OXC}$ and $1/\mu \gg T_{setup}$) $\Rightarrow$ JET $\approx$ Horizon $\approx$ JIT
  If the mean burst size $1/\mu$ is large relative to the values of T_{OXC} and T_{setup}, then from (1), it is also large with respect to T_{offset}. As a result, there are few opportunities for void filling or delayed reservations, and the performance of all three schemes will be very similar. We can reach the same conclusion by observing that, in this case, the traffic intensity value of JIT, JET, and Horizon (see (3), (4), and (5)) is dominated by $1/\mu$, resulting in similar burst drop probabilities for the three schemes. Note that T_{OXC} and T_{setup} represent the overheads associated with switching bursts in the network. Therefore, it is reasonable to assume that, whatever the actual values of these parameters, the mean burst length must be significantly larger, otherwise the network will waste a large fraction of its resources on overhead operations rather than on transmitting bursts, resulting in low throughput or high burst drop probability regardless of the reservation scheme used.

- As a burst travels along its path, its offset value decreases by an amount equal to T_{setup} for each OBS node visited. As a result, inside the network, the offset value becomes dominated by T_{OXC} (refer to (1)), and all three reservation schemes will have similar performance. Consequently, the JET or Horizon schemes may offer the highest benefit at edge nodes, rather than inside the network.

5. NUMERICAL RESULTS

In this section we compare the JIT, JET, and Horizon schemes on a single OBS node in terms of burst loss probability. We use the Erlang-B formula (2) with the appropriate traffic intensity to obtain the burst loss probability. Since this formula is exact only for JIT, we also use simulation for the other two reservation schemes to estimate the burst loss probability. In obtaining the simulation results, we have estimated 95%

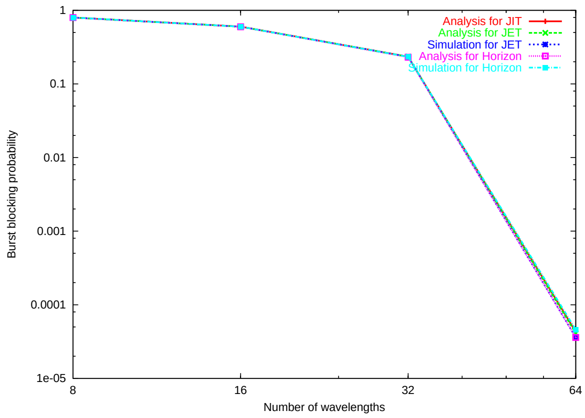

Figure 7. Current scenario: $T_{OXC} = 10ms$, $T_{setup}(JIT) = 12.5\mu s$, $T_{setup}(JIT) = 50\mu s$, $T_{setup}(Horizon) = 25\mu s$, $1/\mu = 50ms$, $\lambda/\mu = 32$

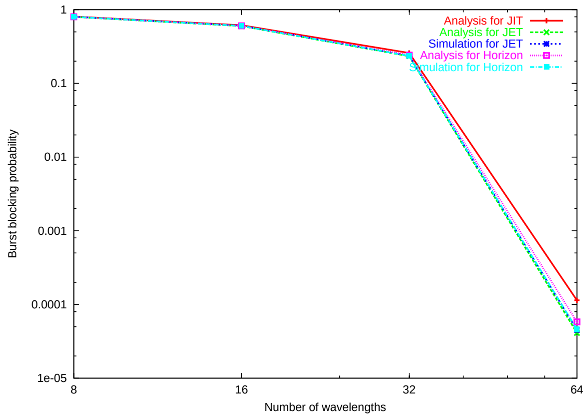

Figure 8. Future scenario: $T_{OXC} = 20\mu s$, $T_{setup}(JIT) = 1\mu s$, $T_{setup}(JIT) = 4\mu s$, $T_{setup}(Horizon) = 2\mu s$, $1/\mu = 100\mu s$, $\lambda/\mu = 32$

confidence intervals using the method of batch means. The number of batches is 30, with each batch run lasting until at least $120,000$ bursts are transmitted to the switch. However, we have found that the confidence intervals are very narrow. Therefore, to improve readability, we do not plot the confidence intervals in the figures presented in this section.

Because of space constraints, we show results for two sets of values for the system parameters: one set corresponds to the current state of technology, and one set corresponds to projections regarding the future state of technology. For the current scenario, we let $T_{OXC} = 10\ ms$, a value that represents the configuration time of existing MEMS switches [13], and $T_{setup}(JIT) = 12.5\ \mu s$, a value that corresponds to the processing time of JIT signaling messages in our JITPAC controllers [12]. For the future scenario, we let $T_{OXC} = 20\ \mu s$ (an improvement of three orders of magnitude over the current scenario) and $T_{setup}(JIT) = 1\ \mu s$ (an improvement of one order of magnitude). Since we do not have actual values for $T_{setup}(JIT)$ and $T_{setup}(Horizon)$, we estimate their values to be four and two times, respectively, the value of $T_{setup}(JIT)$. In other words, $T_{setup}(JIT) = 50\ \mu s, T_{setup}(Horizon) = 25\ \mu s$, for the current scenario, and $T_{setup}(JIT) = 4\ \mu s, T_{setup}(Horizon) = 2\ \mu s$, for the future scenario; while these values are best guess estimates, we have found that the relative performance of the three schemes is not significantly affected as long as these values are a small multiple of $T_{setup}(JIT)$.

In our study, the mean burst size $1/\mu$ was set to 50 ms for the current scenario, and 100 μs for the future scenario; note that these values are equal to five times the respective values of T_{OXC}. We also assume that the number of hops in the path of a burst is uniformly distributed between 1 and 10, and we calculate the offset using (1). The burst arrival rate λ is such that, for both scenarios, $\lambda/\mu = 32$. Finally, in the simulation, we used the latest available unused channel (LAUC) algorithm [15, 16] in JET and Horizon to select an available wavelength for an arriving burst; for JIT, on the other hand, any of the available wavelengths was selected with equal probability to transmit a new burst.

Figures 7 and 8 plot the burst drop probability of JIT, JET and Horizon as the number W of wavelengths varies from 8 to 64. Because of the high value of the arrival rate λ relative to the mean burst size ($\lambda/\mu = 32$), the burst drop probability is high for up to $W = 32$ wavelengths, but it drops dramatically for $W = 64$, and becomes zero for $W = 128$ (not shown in the figures). We observe the good match between analytical and simulation results for JET and Horizon, across both sets of values for the system parameters as well as across the various values of W. We also see that the burst drop probability curves in the two figures are very similar, despite the fact that the ratio $T_{OXC}/T_{setup}(JIT)$ drops from 800 in the current scenario to 20 in the future scenario (the relative values of the processing times of JIT, JET, and Horizon do not change in the two scenarios). More importantly, we observe that the burst probability of the three reservations schemes is identical, except for $W = 64$ in the future scenario, where JET and Horizon slightly outperform JIT. In other words, our results show that, with these sets of values for T_{setup} ad T_{OXC}, there is little opportunity for performing void filling or scheduling of multiple bursts on a wavelength, and thus JET and Horizon behave similarly to JIT. Similar results have been obtained for a wide range of values for T_{setup} and T_{OXC}. On the other hand, we have found that JET and Horizon perform much better than JIT when the mean burst size is significantly smaller than T_{setup} and/or T_{OXC}, in which case there is ample opportunity for void filling and/or scheduling of multiple bursts in the future. However, as we mentioned earlier, it is highly unlikely that OBS networks will be designed to operate under such a scenario, since the high switching and processing overhead would imply low throughput.

6. CONCLUDING REMARKS

We have presented a detailed analysis of the JIT, JET, and Horizon wavelength reservation schemes for OBS networks. We have also presented numerical results to compare the performance of the three schemes in terms of burst drop probability under a range of network scenarios. Our work accounts for the switching and processing overheads associated with bursts as they travel across the network, and it provides new insight into the relative capabilities of the various schemes. Our findings indicate that the simpler JIT reservation scheme appears to be a good choice for the foreseeable future.

REFERENCES

1. C. Qiao and M. Yoo. Optical burst switching (OBS)-A new paradigm for an optical Internet. *Journal of High Speed Networks*, 8(1):69–84, January 1999.
2. J. S. Turner. Terabit burst switching. *Journal of High Speed Networks*, 8(1):3–16, January 1999.
3. J. Y. Wei and R. I. McFarland. Just-in-time signaling for WDM optical burst switching networks. *Journal of Lightwave Technology*, 18(12):2019–2037, December 2000.
4. I. Baldine, G. N. Rouskas, H. G. Perros, and D. Stevenson. `JumpStart`: A just-in-time signaling architecture for WDM burst-switched networks. *IEEE Communications*, 40(2):82–89, February 2002.
5. M. Duser and P. Bayvel. Analysis of a dynamically wavelength-routed, optical burst switched network architecture. *IEEE/OSA Journal of Lightwave Technology*, 20(4):574–585, April 2002.
6. J. Y. Wei, J. L. Pastor, R. S. Ramamurthy, and Y. Tsai. Just-in-time optical burst switching for multi-wavelength networks. In *IFIP TC6 Fifth Inter. Conf. Broadband Commun.*, pp. 339–352, Nov. 1999.
7. H. M. Chaskar, S. Verma, and R. Ravikanth. A framework to support IP over WDM using optical burst switching. In *IEEE/ACM/SPIE Optical Network Workshop*, January 2000.
8. K. Dolzer, C. Gauger, J. Späth, and S. Bodamer. Evaluation of reservation mechanisms for optical burst switching. *AEÜ International Journal of Electronics and Communications*, 55(1), January 2001.
9. K. Dolzer and C. Gauger. On burst assembly in optical burst switching networks - a performance evaluation of Just-Enough-Time. In *Proceedings of ITC 17*, pages 149–161, September 2001.
10. S. Verma, H. Chaskar, and R. Ravikanth. Optical burst switching: a viable solution for terabit IP backbone. *IEEE Network*, pages 48–53, November/December 2000.
11. M. Yoo, C. Qiao, and S. Dixit. QoS performance of optical burst switching in IP-over-WDM networks. *Journal on Selected Areas in Communications*, 18(10):2062–2071, October 2000.
12. I. Baldine, M. Cassada, A. Bragg, G. Karmous-Edwards, and D. Stevenson. Just-in-time optical burst switching implementation in the ATDnet all-optical networking testbed. In *Proceedings of Globecom 2003*.
13. P. B. Chu, S-S. Lee, and S. Park. MEMS: The path to large optical crossconnects. *IEEE Communications*, 40(3):80–87, March 2002.
14. M. Yoo and C. Qiao. Just-enough-time (JET): A high speed protocol for bursty traffic in optical networks. In *IEEE/LEOS Technol. Global Information Infrastructure*, pages 26–27, August 1997.
15. J. Xu, C. Qiao, J. Li, and G. Xu. Efficient channel scheduling algorithms in optical burst switched networks. In *Proceedings of IEEE INFOCOM*, 2003.
16. Y. Xiong, M. Vandenhoute, and H.C. Cankaya. Control architecture in optical burst-switched WDM networks. *IEEE Journal on Selected Areas in Communications*, 18(10):1838–1851, October 2000.
17. L. Kleinrock. *Queueing Systems, Volume 1: Theory.* John Wiley & Sons, New York, 1975.

APPENDIX A. ESTIMATION OF THE PARAMETER Δ FOR HORIZON

We now consider the problem of estimating the value of parameter Δ in the expression (5) for the traffic intensity of Horizon. Recall that Δ represents the increase in the effective service time of bursts under Horizon over that under JET, to prevent any void filling from taking place. In the following analysis, we consider a single wavelength $w, w = 1, \cdots, W$, of the output port in isolation. Assuming that the burst scheduling algorithm is not biased to favor some wavelengths over the others, then, in the long run, we can assume that the arrival rate of bursts to each wavelength is equal to λ/W. Reasoning about the departure process of Horizon becomes much easier when there is a single output wavelength, and, comparing to simulation results, we have found that the results of considering each wavelength in isolation are reasonably accurate.

Let us refer to Figure 4 which shows the burst departure process on a single wavelength. We note that, because of the additional burst dropping (compared to JET) due to the lack of void filling, the mean length of the interval $t_6 - t_5$ is greater than the mean burst interarrival time W/λ. The essence of our approximation is to increase the effective service time of bursts by an amount equal to the difference between the mean length of this interval and the mean burst interarrival time.

We now show how to find the distribution of the length u of the interval of time between t_5 and t_6 in Figure 4. This interval corresponds to the time until the next burst arrival, since any burst arriving after time t_5 is accepted. We let $Prob^{noburst}(u)$ denote the probability that no burst arrives in an interval of length u; note that we assume that this probability depends only on the length of the interval, not its start time.

Let us define the *holding time* of a burst as the sum of three quantities: (1) the burst offset, (2) the burst length, and (3) the OXC configuration time T_{OXC}. From Figure 4, we observe that burst $i+1$ is the first burst whose `setup` message arrives after the arrival of burst i's `setup` message and whose first bit arrives after the end of the holding time of burst i (i.e. t_5). In other words, all the bursts with `setup` messages arriving between t_1 and t_2 must have completed their offset before t_5. Therefore, to analyze the interval between the end of the holding time of burst i and the arrival of burst $i+1$, we only need to consider those bursts whose `setup` messages arrive between t_1 and t_2. Thus we can initiate a new busy period at time t_1, so t_1 is time 0 in this new busy period.

Let s denote the holding time of a burst, which is distributed according to CDF $H(s)$; the Laplace transform of this CDF can be easily obtained from the definition above. Let also $t = t_2 - t_1$ denote the interval between the arrival times of the `setup` messages of bursts i and burst $i+1$.

From [17], we know that for a Poisson arrival process, with a certain number of customers arriving within a given period, the arrival times of these customers are uniformly distributed in that period. Thus, the probability that a customer arriving in $(0, u)$ is still in the system at time u' is $\frac{1}{u}\int_0^u [1 - G(u' - x)]\, dx$, where $G(z)$ is the CDF of the offset length. Then, the probability that the k bursts whose `setup` message arrives in the period $(0, t)$ would have their first bit arrive before time s is $\left[\frac{1}{t}\int_0^t G(s-x)\, dx\right]^k$.

The sum of $k+1$ exponentially distributed intervals follows a $(k+1)$-stage Erlang distribution, so the PDF of t is: $\frac{\lambda/W(\lambda t/W)^k e^{-\lambda t/W}}{k!}$. Therefore, the probability that all the bursts whose `setup` messages arrive in the period $(0, t)$ would have their first bit arrive before time s is:

$$\sum_{k=0}^{\infty} \frac{\lambda}{W} e^{-\lambda t/W} \frac{(\lambda t/W)^k}{k!} \left[\frac{1}{t}\int_0^t G(s-x)\, dx\right]^k \quad = \quad \frac{\lambda}{W} e^{-\lambda/W\left[t - \int_0^t G(s-x)dx\right]} \tag{6}$$

Now, the probability that burst $i+1$ (whose `setup` message arrives at time t) has an offset greater than $s+u$ is $1 - G(s+u-t)$, and the probability that no burst arrives during the interval $(s, s+u)$ is:

$$Prob^{noburst}(u) \quad = \quad \int_{s=0}^{\infty}\int_{t=0}^{\infty} \frac{\lambda}{W} e^{-\lambda\left[t - \int_0^t G(s-x)dx\right]/W} \left[1 - G(s+u-t)\right] dt\; dH(s) \tag{7}$$

The CDF of u is $P(u) = 1 - Prob^{noburst}(u)$, and we obtain the expected value of u as: $\overline{u} = \int_0^{\infty} u dP(u) = \int_0^{\infty}(1 - P(u))du$. Given the CDF $G(z)$ and $H(s)$, it is possible to compute $\overline{u}$ numerically. We then let $\Delta = \frac{\overline{u}}{W} - \frac{1}{\lambda}$ in the expression (5) for the traffic intensity of Horizon.