Quality of service in an optical burst switching ring [‡]

Vishwas S Puttasubbappa and Harry G Perros
{vsputtas,hp}@csc.ncsu.edu

Computer Science Department
North Carolina State University
Box 8206, Office 460 EGRC
1010 Main Campus Drive
Raleigh, NC 27695-7534
Phone: (919) 412 9201, Fax: (919) 515 7925

Abstract

Several access protocols are proposed to support different service classes in an optical burst switched ring. Their performance is evaluated through simulation. Various performance metrics such as throughput, utilization, burst loss rate, end-to-end delay and fairness are used to analyze the behaviour of each protocol.

Keywords - Optical burst switching, MAN, access protocols, service classes

1 Introduction

Optical Burst Switching (OBS) is a novel method currently under study that can be used to transport data over a Wavelength Division Multiplexing (WDM) optical network. Battestelli and Perros [1] provide a detailed survey on OBS and its variations. There is not much work done in the field of OBS over metropolitan-area rings. Xu et al [6] investigated access protocols for OBS rings based on the Just Enough Time (JET) scheme and a new scheme called the Only Destination Delay (ODD). Jong [4] proposed several access protocols for multicasting in such an environment. A new architecture called the LightRing has been proposed by Fumagalli and Krishnamoorthy [3] with multi-token protocol to prevent contention among bursts. Each node can transmit on any of the wavelengths as long as it has the token associated with that particular wavelength. Several Burst Assembly and Transmission (BAT) strategies which deal with simultaneous assembly and scheduling of bursts are proposed. Packets from different flows can

[‡]This work was supported by ARDA under contracts MDA904-00-C-2133 and MDA904-02-C-0482

be assembled into the same burst so as to achieve lower latency of real-time packets. Bouabdallah et al [2] proposed a collision avoidance MAC protocol for a metropolitan bus-based optical access network. Analytical models were developed to calculate the mean access delay of each node in such a shared-medium system. Fairness issues were also investigated.

The work done so far on OBS rings considered traffic to be best effort except in Fumagalli and Krishnamoorthy [3], where real-time and best-effort were considered. The aim of this paper is to investigate how an OBS ring can support different classes of traffic. In this study, we extend the OBS ring architecture proposed in Xu et al [6] in order to consider the following three different classes of traffic. The first class of traffic (Class 1) is a variable bit rate traffic with stringent end-to-end delay constraints, the second class (Class 2) is variable bit rate with no delay constraints, and the third class (Class 3) is non-real time variable bit rate best effort traffic. Several access protocols are proposed and their performance evaluated through simulation.

The paper is organized as follows: Section 2 presents the system architecture including the structure of the metro ring and the OBS nodes. The proposed protocols are presented in section 3, and in section 4 we describe the simulation model and the arrival processes. The results of the simulation are given in section 5, and finally section 6 presents the inferences from the study.

2 System Architecture

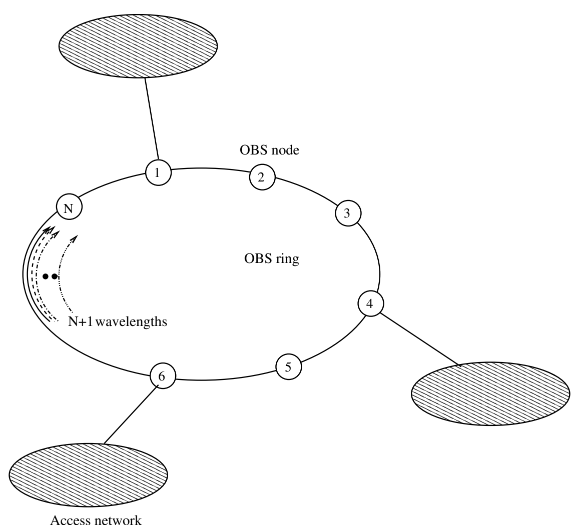

Figure 1: Architecture of the ring

The metro ring shown in figure 1, uses the WDM architecture and it consists of several nodes which serve as concentration points to incoming and outgoing traffic from several access

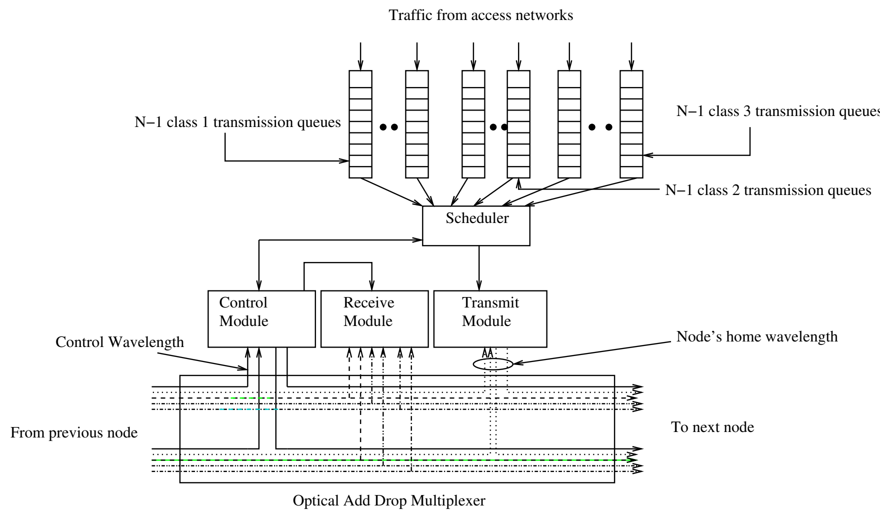

Figure 2: OBS node architecture

networks. The ring is divided into two co-existing rings as follows: Some of the wavelengths are used to host SONET/SDH rings, and the remaining wavelengths are used for optical burst switching. The metro ring has to carry all types of traffic, such as circuit switched traffic, ATM traffic and IP traffic. SONET/SDH rings can cater to circuit switched traffic and data traffic can be transported through the OBS ring. In this paper, we only investigate the OBS part of the metro ring.

The wavelengths allocated for optical burst switching are divided into S sets of N+1 wavelengths. Within a set of wavelengths, each of the N wavelengths is allocated to a different node. This wavelength is referred as the *home* wavelength of the node. Since there are S sets of wavelengths, each node is allocated to S home wavelengths. A node can only transmit bursts on its home wavelengths. The (N+1)th wavelength within each set of wavelengths is used as the control channel, and it is accessed by all the nodes. The control wavelength carries control frames around the ring. The control frames implement the signalling necessary for OBS. Each set of home wavelengths has its own control channel. Scalability of the ring with respect to this concept of S wavelengths per node is not an issue since the metro ring has limited number of nodes and the number of wavelengths per fiber is continuously increasing.

The architecture of an OBS node in the ring is shown in figure 2. Each node in the ring has S transmitters each fixed-tuned to one of the S home wavelengths, and S tunable receivers one per wavelength set. These S pairs of transceivers are used for transmitting and receiving bursts. A node can transmit a burst on any free home wavelength. A free receiver can tune to receive bursts arriving on any wavelength in its corresponding wavelength set. Each node is additionally equipped with S transceivers, one set per control wavelength.

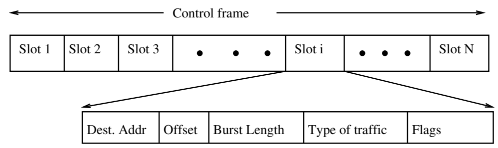

Figure 3: Control Frame structure

The node is equipped with a control module that reads each control frame that passes through it in the ring. Based on the information that the control frame carries, the control module decides what to do next. For instance, it might decide to transmit a burst following the frame, tune a receiver to a particular wavelength to receive a burst, etc.

Each control wavelength carries back to back control frames. The structure of a control frame is as shown in figure 3. The control frame comprises of slots equalling the number of nodes in the ring. Each node has its own slot into which it can write information during transmission. Each slot has the following fields: (a) destination address, (b) offset value, (c) burst length, (d) type of traffic and (e) flags. The flags field carries information which is protocol specific and will be explained later on. The control frames on the S control wavelengths travel around the ring in a synchronous manner. That is, the control frames in the (i+1)st control wavelength lag behind their corresponding ones in the ith control wavelength by the time the control module requires to process them. This arrangement ensures efficient usage of control frames for burst transmission. For instance, if a node cannot transmit on the first control channel, it has an opportunity to transmit immediately using the control frame in the second control channel without having to wait even for a small amount of time.

Each node serves a number of access networks. The incoming data from these networks is queued in the transmission queues of a node. Specifically, each node maintains N-1 transmission queues of each class, where N is the total number of nodes in the ring. Since we consider three classes, each node maintains 3(N-1) queues. In this paper we assume that class i has non-preemptive priority over class i+1, i=1,2. That is, the transmitter will always transmit a burst from a class 1 queue. If there is no class 1 traffic, it will transmit a burst from a class 2 queue, and if there is no class 1 or 2 traffic, it will transmit a burst from class 3. Bursts in the N-1 class 1 transmission queues are served in a FIFO manner. That is, these N-1 queues are equivalent to single class 1 queue. In other words, bursts arrive and are served through a single class 1 queue. The N-1 class 2 transmission queues are served in a round-robin fashion. Likewise, round robin is used to serve the N-1 class 3 transmission queues.

3 The Access Protocols

In this paper, we assume that the class 1 traffic consists of multiple HDTV streams. Each HDTV frame constitutes a single burst. In the case of class 2 and class 3 traffic, a burst is comprised of several data packets which may be IP packets, ATM cells etc. A class 2 and class 3 transmission queue is eligible to be served if there are enough packets in the queue whose aggregate size is greater than MinBurstSize_CLA2 and MinBurstSize_CLA3 correspondingly. The burst size is limited by MaxBurstSize_CLA2 and MaxBurstSize_CLA3. Both the minimum and maximum burst size values are set equal to those used in Xu et al [6].

We have defined and analyzed the following five protocols. *Destination-Reservation Free* which provides no guaranteed delivery to any traffic class, *Ack* and *Token* which provide guaranteed delivery to class 1 traffic and *Token-Token* and *Ack-Ack* which provide guaranteed service to both class 1 and class 2 traffic. All these protocols transmit class 3 bursts when bandwidth is available. These access protocols fall into two categories: *collision-free* and *collision* protocols. Collision-free protocols result is no burst loss in the ring. Before the source transmits a burst, the source makes sure that the destination is free to receive the burst and that there are no collisions at the destination. Typically, the source achieves this by using some sort of a reservation scheme to reserve resources at the destination. None of these five protocols are collision-free for class 3 traffic. Token-Token and Ack-Ack are class 1 and class 2 collision-free protocols since bursts of either classes are not lost. Token-Token uses tokens and Ack-Ack makes use of acknowledgements to ensure a collision-free reception at the destination. Ack and Token are class 1 collision free protocol as no class 1 bursts are lost. Ack and Token protocols use acknowledgements and tokens respectively to reserve resources at the destination but only for class 1 traffic. No burst-loss guarantee is provided for class 2 traffic. Finally, the Destination-Reservation-Free protocol is not collision-free for any of the three classes. Each of the protocols is explained in detail in the following sub-sections.

3.1 Destination-Reservation Free Protocol (Dest-Resv-Free)

Nodes transmit bursts without making any reservations at the receiver node. Hence there is no guaranteed acceptance of transmitted bursts. This results in multiple bursts arriving at the destination at the same time and hence collisions. This is the Tell-and-Go protocol currently used in OBS mesh networks.

Single transceiver case: Upon arrival of a control frame, the control module checks whether the transmitter is busy. If it is not, then a burst can be transmitted. The bursts are transmitted with the priority scheme as stated in section 2.

Since there is no reservation of resources at the destination, multiple bursts arrive at the same time and thus collisions can occur. Priority for receiving a burst is given to class 1 bursts.

If multiple class 1 bursts arrive at the same time, one of them is randomly selected. Pre-emption of class 2 and class 3 bursts is allowed if a class 1 burst arrives. Class 2 bursts are given the next priority and in case many of them arrive at the same time, one of them is randomly chosen. Class 3 bursts are given the least priority and if many of them arrive at the same time, one of them is randomly selected.

Multiple transceivers case: In the case of multiple home wavelengths, the control module can choose any free home wavelength to transmit. The reception mechanism is identical in each of the different sets of wavelengths.

3.2 Token Protocol

The Token protocol uses the concept of tokens to resolve receiver collisions. Tokens are used only for class 1 bursts. Class 2 and class 3 is serviced through best-effort. Every node has a token circulating around the ring. If a source wants to transmit a class 1 burst to a particular destination, it has to have the token for that destination. All the nodes maintain a queue to hold tokens. The token is released after the transmission is completed. Since only the node that has possession of the token can transmit a burst to the appropriate destination, the destination can only receive a single burst at a time, and therefore the token protocol is a collision-free protocol.

Single transceiver case: Tokens exist only for class 1 bursts. Each node monitors the control frames for tokens. If a control frame is carrying a token for destination k, then the flag field of slot k has the value 1. The node takes the token out of the control frame and queues it in its FIFO queue, provided that there is traffic for that destination node, otherwise, it lets the token pass through. The token has to be queued because the node may not be able to transmit the burst immediately because of busy transmitter. The token is released after the transmission is completed. This guarantees no collision at the destination between class 1 bursts. Class 2 and class 3 bursts do not have a collision-free reception mechanism. Bursts are simply transmitted following the Dest-Resv-Free protocol.

On the receiver's side, class 1 bursts are given the highest priority. There is pre-emption of class 2 and class 3 bursts, in case the receiver is busy receiving either of them and a class 1 burst arrives. Since this is a class 1 collision-free protocol, there are no collisions among class 1 bursts. But collisions can occur involving multiple class 2 and class 3 bursts. If no class 1 burst arrives, priority is given to a class 2 burst and in case of multiple class 2 bursts, one of them is randomly chosen. Likewise, if no class 1 and class 2 bursts arrive, priority is given to class 3 bursts, one of which is randomly selected.

Multiple transceivers case: Each node maintains a separate token queue for each home wavelength. A node cannot use more than one home wavelength simultaneously to transmit bursts to a particular destination. This is achieved by making sure that only one token to a particular destination can be held in any of the multiple token queues. A class 1 burst can be transmitted on any home wavelength as long as the corresponding token queue has the token for the appropriate destination. Reception in any one set of wavelengths is independent and identical to the other sets.

3.3 Ack Protocol

This protocol also ensures guaranteed reception of class 1 bursts. It is different to the Token protocol and is based on the Tell-And-Wait (TAW) protocol proposed in OBS networks. The protocol uses a request and acknowledgement mechanism. A minimum of a round-trip delay is required for the node to transmit a burst after it can be formed. The acknowledgement mechanism is available only for class 1 traffic. Class 2 and class 3 bursts are transmitted as in the case of the Token protocol using the Dest-Resv-Free protocol.

Single transceiver case: As soon as a HDTV frame arrives at any of the class 1 queues of the node, a *Request* is sent out to the destination node requesting it to return an *Acknowledgement* in which it indicates the earliest time it is free to receive this burst. The source is not allowed to send out a request to any destination as long as it has an outstanding request. This makes sure that there are no transmitter conflicts. To send a request, the source node catches the next arriving control frame and marks the flag field in its slot as *Request*. When the destination node sends an acknowledgement, the offset field in the destination's slot will be the earliest time by which the source can start transmitting. This is to make sure that by the time the source starts transmitting the burst, the destination is free to receive it.

The transmission of class 2 bursts is without any acknowledgement mechanism. If the transmitter is not busy transmitting class 1 traffic, then, class 2 bursts are transmitted. If neither class 1 nor class 2 queues are ready for transmission, a class 3 burst can be transmitted.

A destination node cannot send out an acknowledgement for a request that it has received until it starts receiving the burst from the previous acknowledgement that it sent. This is necessary because the source node to which the last acknowledgement was sent may not be able to transmit the burst immediately because of a busy transmitter. Upon starting to receive the previous burst, the receiver can calculate the time until which it is busy, and it can then send out an acknowledgement to the next request in the queue as to when the source can start transmitting. This prevents receiver conflicts between class 1 bursts. Collisions do occur between a class 1 burst and multiple class 2 and class 3 bursts, in which case the single class 1 burst is chosen. Pre-emption of class 2 and class 3 bursts occurs in case the receiver has to receive a class 1 burst while it is busy receiving either of them. The receiver has a queue to hold requests and the request with the earliest time-stamp will be served first. This makes sure that the frame that arrived the earliest at a source node would be served first irrespective of where the source node is with respect to the destination and the order in which the requests were received at the destination. Class 2 bursts are given higher priority than class 3 bursts.

Multiple transceivers case: Similar to the Token protocol, a node is not allowed to use multiple home wavelengths to transmit bursts to the same destination. This is to make sure that a particular destination does not get held up serving requests by the same destination. While sending a request on any of the control channels, the source node makes sure that no other home wavelength is being used for the same destination. If this is the case, other transmission queues are served which send out requests to different destinations. During reception, each receiver looks at its control channel and tunes to a particular wavelength in its wavelength set using the regular priority scheme.

3.4 Token-Token Protocol

This is a collision free protocol for class 1 and class 2 bursts. Nodes use the token mechanism to ensure that the bursts belonging to class 1 or class 2 categories are received without receiver conflict.

Single transceiver case: The token for a particular destination is captured by the node only if it has bursts of class 1 or class 2 to be sent to that particular destination node. In case both the class 1 and class 2 transmission queues have traffic for the destination, priority is given to class 1 traffic. Unlike the token protocol, wherein class 2 bursts were sent whenever the transmitter was free, this protocol transmits class 2 bursts when it captures a token and there is no class 1 traffic for any destination. Thus, this protocol ensures that neither class 1 nor class 2 bursts are dropped at the receiver due to collisions. Class 3 queues are served whenever the transmitter is free and the bursts can be formed.

Reception is simpler in this case since there is no collision between class 1 and class 2 bursts and hence no priority between them. Class 3 bursts are subject to pre-emption in case either of the other classes' burst arrives.

Multiple transceivers case: The operation is similar to that of the Token protocol.

3.5 Ack-Ack Protocol

Guaranteed reception both for class 1 and class 2 bursts is provided by extending the acknowledgement scheme to cover both traffic classes. Class 3 is served through best effort.

Single transceivers case: The operation of the transmitter is similar to the Ack protocol, except that class 2 bursts are sent only after receiving an acknowledgement from the destination node. When the node does not have any outstanding request, a request is sent to the destination for either type of traffic. Priority is given to class 1 traffic. Unlike the Ack protocol, a class 2 burst is not transmitted whenever the transmitter is free. A request to transmit a class 2 burst is sent only if all the class 1 transmission queues are empty. Thus, this is a collision-free protocol for class 1 and class 2. Additional details that are sent in the control frame carrying the request include the time of arrival of the frame and the type of traffic the request is for. Class 3 bursts are transmitted whenever the transmitter is free.

The receiver exercises its decision to send out acknowledgements not in the FCFS manner, but based on priority. Once the receiver starts receiving a burst for the acknowledgement it last sent, it scans all the requests in its queue and sends out the next acknowledgement according to the following rules: (1). Requests which are for class 1 traffic are given priority. If there are none, then one of the class 2 requests is arbitrarily picked and served (2). If there are multiple class 1 requests, then the request that has the earliest time-stamp will be served first. Clearly, the above mechanism is designed in view of the delay constraints of class 1 traffic. Rule 2 makes sure that requests are served not in their order of arrival (which may be biased towards nodes close by) but by the earliest time-wise arrival of frames at their respective source nodes.

Since acknowledgements make sure that there are no multiple bursts following the same control frame, it is clear that there is no receiver conflict for class 1 and class 2 bursts. Pre-emption of class 3 bursts takes place in case a burst of either class arrives when the receiver is busy receiving a class 3 burst.

Multiple transceivers case: The operation is similar to that of the Ack protocol.

4 The Simulation Model

An event-based simulation model was constructed with a view to analyzing the performance of the above ve proposed protocols. This simulation model was not constructed following the typical event-based simulation techniques, see Perros [5], where all the pending events are kept in a sorted linked list. Rather, the nodes of the OBS ring are simulated one at a time following the path of each control frame. For instance, for the single transceiver case, the N nodes are simulated in sequence as visited by the first control frame, then they are simulated for the second control frame etc. For each control frame that arrives at a node, the node runs the arrival process module and updates the transmission queues, and then the transmitter and receiver modules are executed. The transmitter module has the following functions depending on the protocol that is being simulated. (1) Determine whether to transmit a class 1, class 2 or a class 3 burst by looking at their respective queues. (2) Transmit a request to the destination to send back an acknowledgement (Ack protocol) (3) Capture a token for a particular destination (Token protocol). The receiver module has functions such as: (1) Receive the bursts that follow the control frame by an offset. (2) Give higher priority to class 1 bursts over class 2 bursts and class 3 bursts. This may involve pre-emption of class 2 and class 3 bursts. (3) Receive requests sent by source nodes for acknowledgements. (4) To stack all the requests and send out acknowledgements one by one.

The transmission queues of a node are updated whenever a control frame arrives at a node. Class 2 and class 3 packets that arrived between the last control frame and the next control frame are accumulated in their respective transmission queues. Real time class 1 frames are placed in their corresponding real time transmission queues. The arrivals are taken care of before examining the control frame.

In the simulation model the actual duration of a burst is not simulated. Since the bursts follow the control frame after an offset amount of time, the receiver is kept in the busy state from the time the control frame passes the destination node plus offset till the time it takes for the burst to be received completely. Each node maintains its local time through the control frames. Since the control frames have the same length and are equally spaced out, the local clock of a node can be updated whenever the next control frame arrives. This simulation model permits very fast execution times.

4.1 Arrival Processes

The class 1 traffic is assumed to be HDTV traffic. To this effect, for each node i, i=1,2,...,N, a number of HDTV streams are setup at the beginning of the simulation. Each of these streams originate at node i and terminate at destination node j. In each stream, frames are generated at a rate of 60 frames per second giving an inter-frame arrival time of 16.667 milliseconds. We assume that the frames follow the MPEG 2 Group Of Pictures (GOP) structure of IBB PBB PBB PBB. The size of each frame is generated using the auto-regressive model, see Bragg [8], S(t) - S(t-12) = e(t) - 0.69748 $\times$ e(t-3), where S(t) is the size of frame t, and e(t) $\sim$ N(0,σ^2) with σ^2=4849.5. A trace of frame sizes generated over a period of time is plotted in figure 4. This autoregressive model produces an average bit rate of 20 Mbps.

A class 2 source in our simulation is a variable bit rate source with no end-to-end time constraints. In our simulation experiments, we assume that the packets are generated from a storage area network (SAN), with the following packet-size distribution: 44 % of 64Kbytes, 18 % of 56K, 21 % of 40K, 4 % of 32K, 4 % of 24K and 6 % of 8Kbytes, see Trevitt [7]. The arrival process consists of packets arriving in succession with an exponentially distributed inter-packet delay. The time it takes for each packet to arrive is taken into account.

Finally the class 3 traffic arrival process is best effort traffic and it is modelled as in Xu et al [6] by a modified Interrupted Poisson Process. The ON and OFF periods are exponentially distributed. Packets arrive back to back during the ON period at the rate of 2.5 Gbps. The last packet that arrives when the ON period ends is truncated. During the OFF period, no packets are generated. The mean packet size is 500 bytes and any packet size above 5000 bytes is truncated to 5000 bytes. To calculate the ON and OFF periods, we use the coefficient of

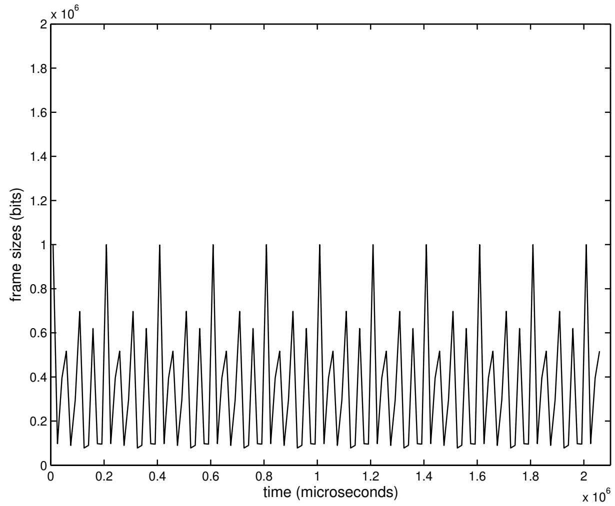

Figure 4: Frame sizes of stream 8 generated in node 5

variation c^2, defined as the ratio of the variance of the packet inter-arrival time divided by the squared mean of the packet inter-arrival time. c^2 indicates the burstiness of the arrival process.

$$c^2_{IPP} = 1 + \frac{2\lambda\mu_1}{(\mu_1 + \mu_2)^2} \tag{4.1}$$

where $\frac{1}{\lambda} = (500\text{bytes})/(2.5\text{Gbps}) = 1.6\mu\text{s}$, and $\frac{1}{\mu_1}$ and $\frac{1}{\mu_2}$ are the mean times of the ON and OFF periods. The arrival process of class 3 traffic is completed by the following equation:

$$\text{Average Arrival Rate} = 2.5 \text{ Gbps} \times \frac{\mu_1}{\mu_1 + \mu_2} \tag{4.2}$$

5 Discussion of Simulation Results

We simulated a ring consisting of 10 nodes and each node is separated by a distance of 5 km. Each wavelength was assumed to have a bandwidth of 2.5 Gbps and the control wavelength works at a rate of 622 Mbps. For each class i, i=1,2,3, the transmission queue in a node was assumed to have a buffer size of 1 MB. In the single transceiver case, 11 wavelengths are required for the OBS network. In the multiple transceivers case, the number of wavelengths used is an integral multiple of 11.

Each slot in the control frame is assumed to be of size 100 bytes. Assuming that there are 10 slots in the control frame, it takes 12.86 seconds for a frame to arrive. In view of this, we assume that each node takes 12.86 microseconds to read one control frame. Given that a receiver has 1

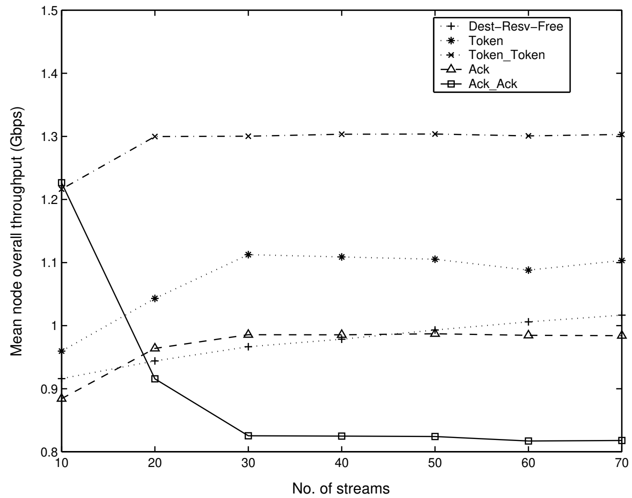

Figure 5: Mean node overall throughput (1 transceiver)

microsecond tuning delay, the offset value is 13.86 microseconds. The burst is delayed by this value and then transmitted behind the control frame.

The simulation results are plotted with 95% confidence interval estimated by the method of batch means, see Perros [5]. Each batch is completed when each node generates 10,000 bursts. The confidence intervals are very tight and are not discernible in the graphs.

The simulation model was used to evaluate the performance of each of the protocols discussed in section 3. Several characteristics of the system such as throughput per node, % bandwidth utilization, hit ratio, % burst loss, mean frame delay, delay fairness and throughput fairness were measured assuming 1, 2 and 3 transceivers. For all the results obtained, the class 2 average arrival rate at each node was fixed to 0.8 Gbps and the average arrival rate of class 3 traffic to 0.5 Gbps. The x axis is always the number of HDTV streams originating at each node. Specifically, in each simulation experiment, the same number of HDTV streams originate at each node, and the destination node of each stream is randomly selected. Each stream contributes an average of 20 Mbps of the total traffic. The total average arrival rate is the sum of the average arrival rates of the 3 traffic classes. In all the experiments, the overall traffic a node transmits is less than the bandwidth of the home wavelength(s).

Figure 5 plots the *Mean node overall throughput* versus the number of HDTV streams per node. The *mean node overall throughput* is defined as the average number of bits received (class 1, class 2 and class 3) by all the nodes in a unit time divided by the number of nodes in the ring. The Token-Token protocol has the highest mean node overall throughput followed by the Token protocol. All the protocols show an upward trend with an increase in the arrival rate

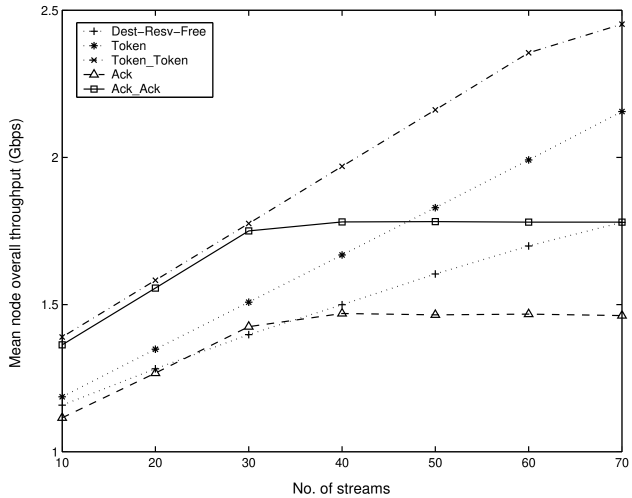

Figure 6: Mean node overall throughput (3 transceivers)

except for the Ack-Ack protocol. With the Ack-Ack protocol, it decreases with increasing class 1 traffic because of the lack of service provided to class 2 traffic. Simulation results showed that Ack-Ack protocol is the only class 1 and class 2 collision free protocol wherein class 1 traffic is insensitive to class 2 traffic. Thus, as class 1 traffic increases in intensity, the overall throughput is only due to class 1 and class 3 traffic, and as a result the mean node overall throughput drops. This can be confirmed by looking at figure 9 which plots the % bandwidth utilization of class 2 traffic versus the number of HDTV frames per node. As class 1 arrival rate increases, the overall throughput becomes largely a factor of class 1 traffic since priority is given for class 1 over class 2. Token-Token is not as biased towards class 1 traffic as Ack-Ack. This can be confirmed from figure 7 where we note that even though the class 1 arrival rate increases, Token-Token cannot increase the service provided to class 1 traffic. The throughput saturates at a very early point. Token-Token provides better service to class 2 traffic than Ack-Ack. The Token protocol performs better than Ack protocol.

The mean node overall throughput for the three transceivers case is plotted in figure 6. A notable feature is a better performance by Ack-Ack. With the addition of two more home wavelengths, Ack-Ack has a comparable throughput to others because class 2 traffic gets serviced better, thus contributing to a better throughput. It can be seen that Token-Token scales well.

Figure 7 plots the *Mean node class 1 throughput* versus the number of HDTV streams per node. The *mean node class 1 throughput* is defined as the average number of class 1 bits received by all the nodes in a unit time divided by the number of nodes in the ring. The Token protocol has the highest mean node class 1 throughput followed by Ack. Ack-Ack performs very closely to Token and Ack despite the fact that it provides acknowledgement services to both class 1 and

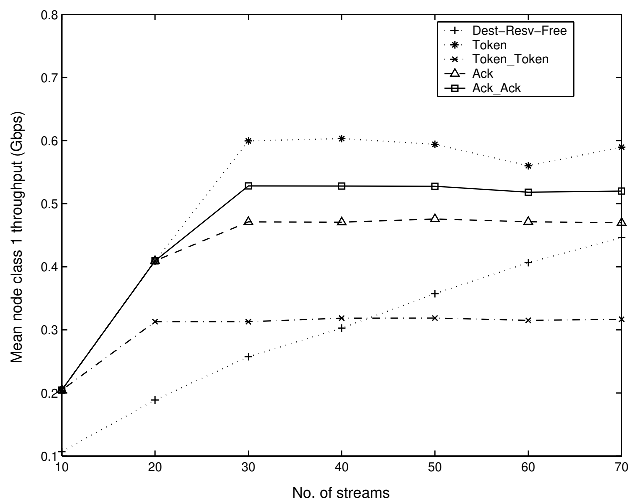

Figure 7: Mean node class 1 throughput (1 transceiver)

class 2 traffic. In Token-Token, there is a noticeable lack of service provided for class 1 traffic since collision free delivery is provided for both types of traffic.

Figure 8 plots the mean node class 1 throughput when there are three transceivers per node. We note that Token and Token-Token scale well. With more than adequate bandwidth available for class 1, Token-Token performs on par with Token even after providing collision-free reception for both class 1 and class 2 traffic. Ack and Ack-Ack do not scale up well. The protocols that make use of tokens scale up well with more control channels because each node gets a token more frequently. With the acknowledgements-based protocols, the source nodes may now send out the requests early cutting down the queueing delay but the propagation delay remains the same. This limits the scalability of such protocols with additional transceivers.

The % *Bandwidth utilization* is defined as the amount of time a home wavelength is busy transmitting bursts. Two graphs for the utilization of the home wavelength for node 1 are given: one depicting % bandwidth utilization for class 2 traffic and the other for class 1 traffic. Figure 9 plots the % bandwidth utilization for class 2 traffic for the 1 transceiver case. It is clear that class 2 traffic suffers in the Ack-Ack protocol because of the reasons explained earlier. Token-Token is a more fair protocol when class 1 and class 2 traffic is considered. The Ack, Token and Dest-Resv-Free protocols have a constant % bandwidth utilization for class 2 traffic, firstly, because of constant class 2 average arrival rate of 0.8 Gbps and secondly, because all of them provide best-effort transmission of bursts. The % bandwidth utilization of these three protocols coincide with each other as shown in figure 9. As figure 10 depicts, Ack-Ack has higher utilization for class 2 traffic because of higher bandwidth available in the 3 transceiver case.

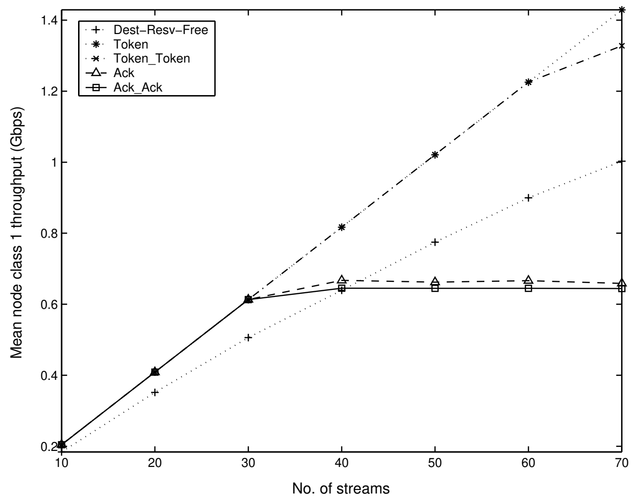

Figure 8: Mean node class 1 throughput (3 transceivers)

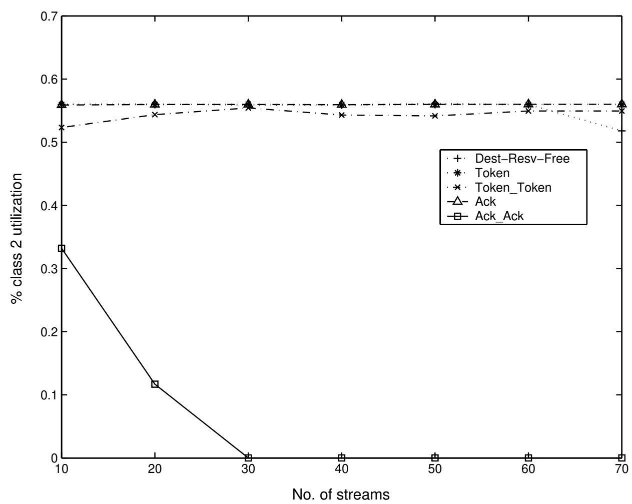

Figure 9: % Bandwidth utilization of node 1 transmit wavelength for class 2 traffic (1 transceiver)

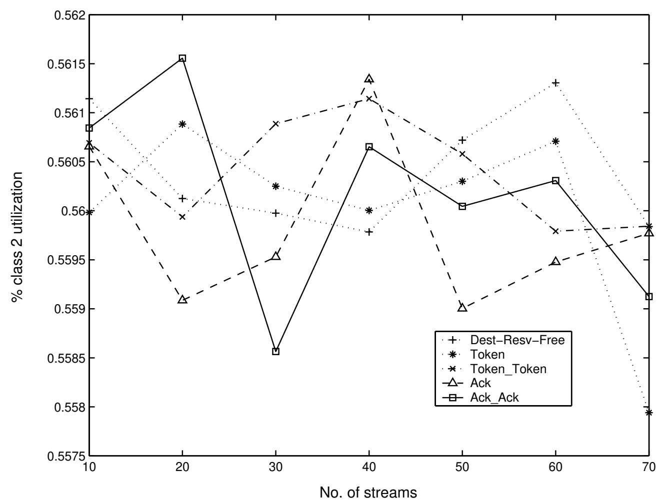

Figure 10: % Bandwidth utilization of node 1 transmit wavelength for class 2 traffic (3 transceivers)

Figure 11 plots the % bandwidth utilization of transmit wavelength of node 1 for class 1 traffic for the 1 transceiver case. From these plots, it can be seen that although Dest-Resv-Free protocol has a high % bandwidth utilization, it has a low class 1 throughput because of collisions. Figure 12 plots class 1 % bandwidth utilization for the 3 transceivers case. It can be seen that Token and Token-Token scale up for higher class 1 % bandwidth utilization than Ack and Ack-Ack.

Overall burst loss rate is defined as the total number of bursts (class 1, class 2 and class 3) lost because of receiver collisions divided by the total number of bursts transmitted by all the nodes. Figure 13 plots the *overall burst loss rate* with varying number of HDTV streams per node. Ack-Ack and Token-Token are class 1 and class 2 collision-free protocols and hence, bursts lost in the system are only due to class 3. Ack-Ack has lesser overall burst loss than Token-Token. It is surprising to see that Dest-Resv-Free has lesser overall burst loss than Token-Token, Ack and Token. This is because class 2 burst loss forms the majority of burst losses, since more class 2 bursts are transmitted because of smaller burst sizes than class 1 bursts. Protocols which provide guaranteed delivery to class 1 bursts have a higher class 2 burst loss rate. This can be confirmed from figure 15 which plots the *class 2 burst loss rate* with class 1 arrival rate. Since Ack-Ack restricts the number of class 2 bursts sent, overall burst loss is less. Dest-Resv-Free has a lower overall burst loss rate because of fewer class 2 bursts dropped. Figure 14 shows the overall burst loss rate for the case of three transceivers.

Figure 15 plots the class 2 burst loss rate with class 1 arrival rate. It can be seen that Token-Token and Ack-Ack have zero bursts lost because they are class 2 collision-free protocols.

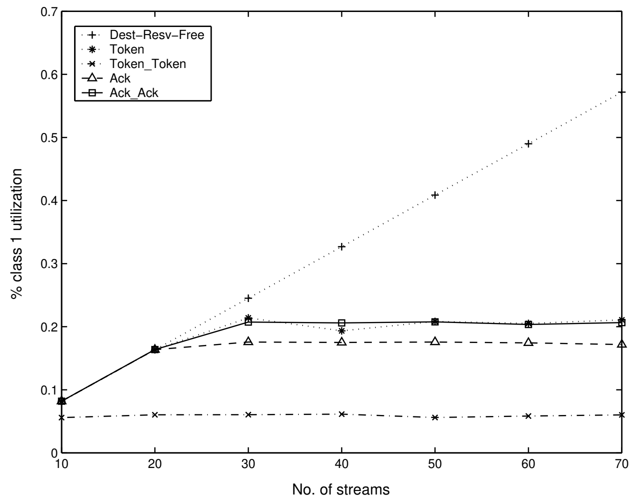

Figure 11: % Bandwidth utilization of node 1 transmit wavelength for class 1 traffic (1 transceiver)

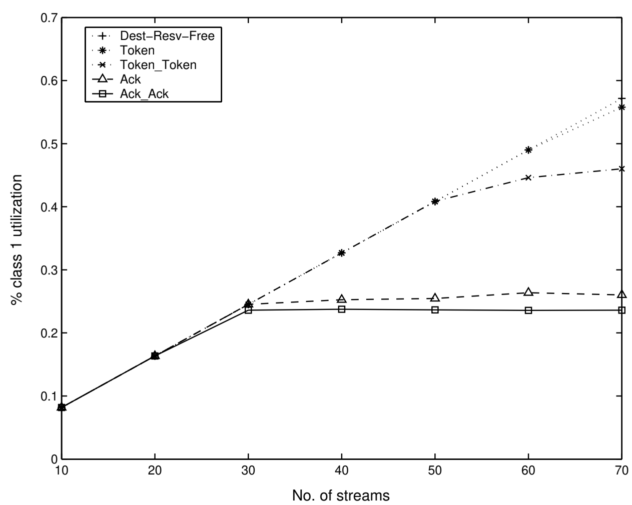

Figure 12: % Bandwidth utilization of node 1 transmit wavelength for class 1 traffic (3 transceivers)

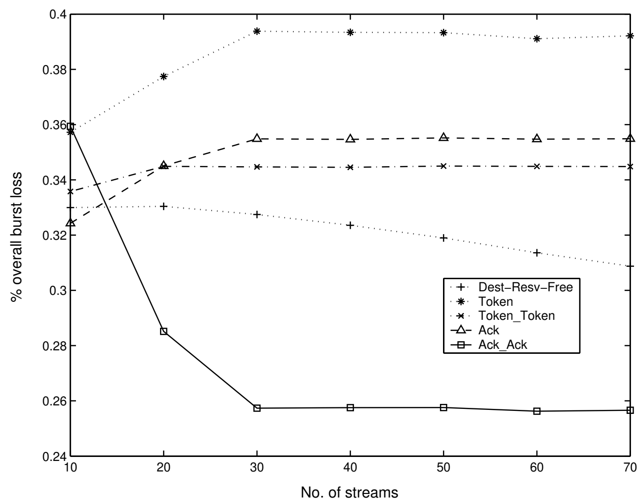

Figure 13: Overall Burst loss rate (1 transceiver)

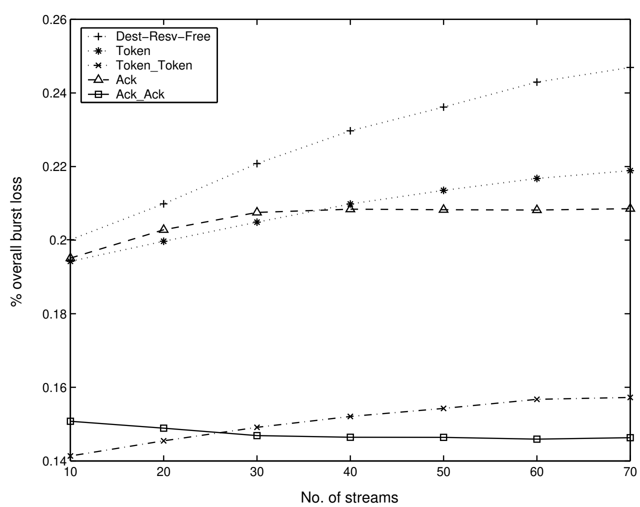

Figure 14: Overall Burst loss rate (3 transceivers)

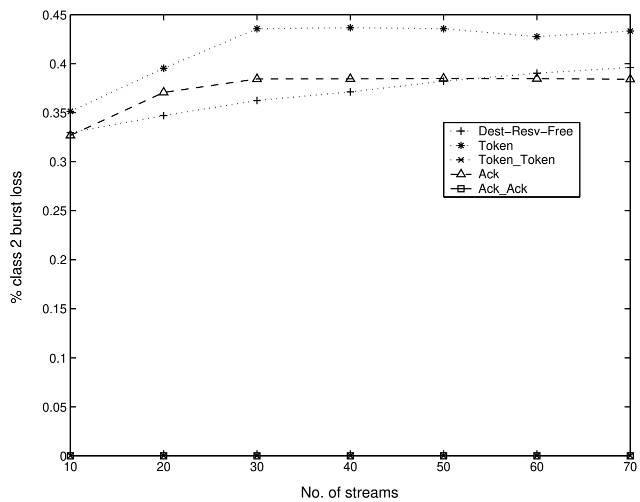

Figure 15: Class 2 Burst loss rate (1 transceiver)

The other protocols exhibit similar burst loss rates as each other.

Figure 16 plots the hit ratio for the five protocols. This is an important performance metric that describes how good a protocol is to support class 1 traffic. Hit ratio is defined as the total number of class 1 bursts (i.e. HDTV frames) received in time by all the nodes, divided by the total number of frames sent by all the nodes. A frame is received in time if it arrives within 17 milliseconds of receiving the previous frame from the same stream. If a frame arrives after 17 milliseconds of receiving the last frame, then the frame is late and is a miss. A higher hit ratio for a larger number of streams per node is the desired performance criterion. For the single transceiver case, the Ack protocol has the best hit ratio followed by Token and Ack-Ack. All the protocols which provide guaranteed delivery for class 1 traffic have a sudden fall in the hit ratio values after a particular number of streams per node. In these protocols, since each frame is transmitted after getting an assurance that they will be received without collision, the delay has a cascading effect on subsequent frames and the end-to-end delay makes it highly improbable to have a decent hit ratio. Dest-Resv-Free protocol does not provide such a guarantee and hence has a slowly decreasing curve. Token-Token performs badly as well.

Token and Token-Token perform remarkably well when more bandwidth is available as can be seen in figure 17. Even though Ack has an initial advantage, the Token and Token-Token protocols scale well to provide better service to class 1 traffic. This is due to the availability of more tokens which leads to more frequent availability and thus faster service of class 1 bursts.

Another performance metric that can give a good idea about the number of streams each node can support is the *mean frame delay*. The mean frame delay is defined as the sum of

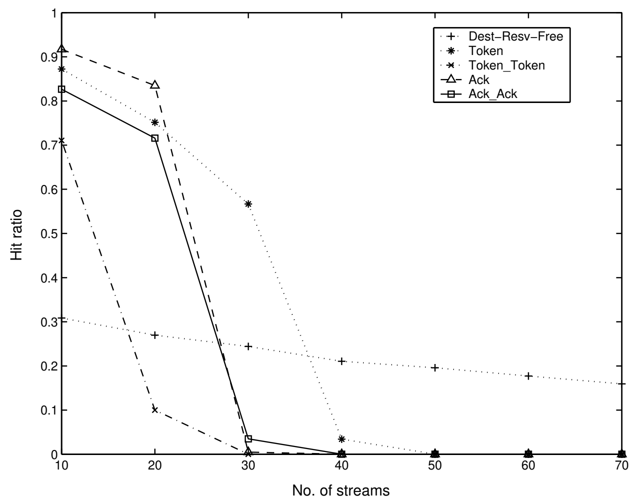

Figure 16: Hit ratio (1 transceiver)

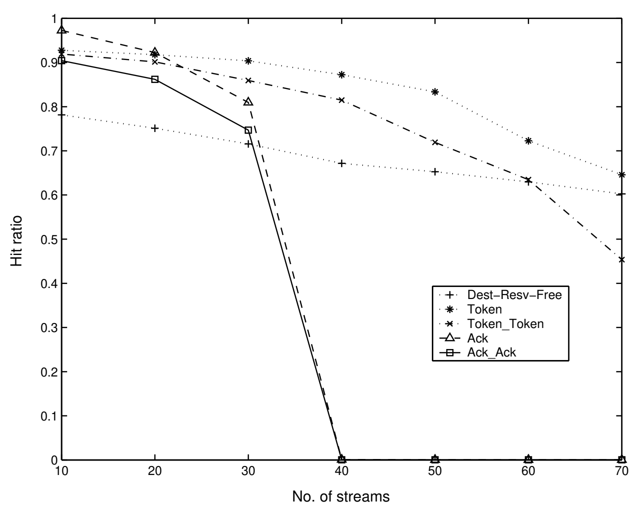

Figure 17: Hit ratio (3 transceivers)

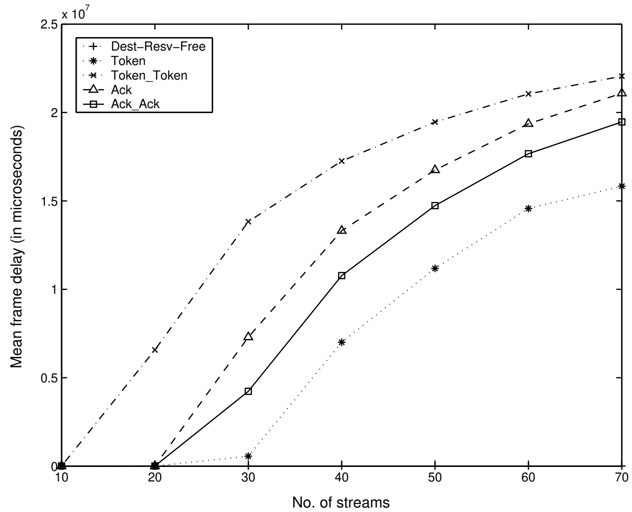

Figure 18: Mean frame delay (1 transceiver)

the end-to-end delay experienced by all the frames received, divided by the number of frames received by all the nodes. The end-to-end delay comprises of the queueing delay that a frame experiences at the transmitting node from the moment it arrives to the moment it is transmitted out, plus the propagation delay. Figure 18 plots the mean frame delay versus class 1 traffic. The Token protocol steeps later than all the other protocols. This implies that the Token protocol can support more streams per node than any other protocol. Dest-Resv-Free protocol has a small variation because the queueing delay component does not vary highly. Figure 19 plots the mean frame delay for the 3 transceivers case.

The *fairness* of a protocol is an important criterion in ring networks, since it shows whether the positioning of a node in the ring has any effect on the protocol's performance metrics. Two types of fairness indices were calculated. The *throughput fairness index* of a protocol which determines how fair a protocol is with respect to the throughput of individual nodes, and the *delay fairness index* of a protocol which determines how fair a protocol is with respect to the delay in a node. The definition of the two indices metrics are the same as those used in Xu et al[6]. The throughput fairness index of a node i is defined as the c^2 of the throughput from node i to all other nodes.

$$\text{Throughput Fairness Index of Node i} = \left(\sum_{j=1, j \neq i}^{10} (H_{ij} - \overline{H_i})^2 \right) \times \frac{1}{\overline{H_i}^2} \tag{5.3}$$

where H_{ij} is the throughput from node i to node j, and $\overline{H_i} = (\sum_{j=1,j\neq i}^{10} H_{ij})/9$. The throughput fairness index of the protocol is defined as the average of the throughput fairness indices of all

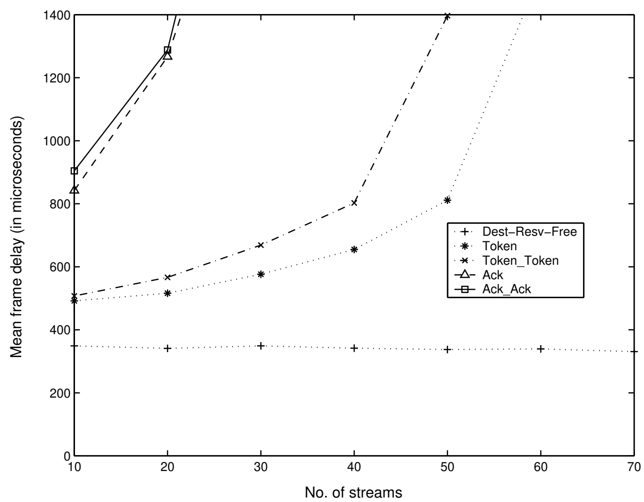

Figure 19: Mean frame delay (3 transceivers)

the nodes. The throughput fairness index of the protocol was computed only for the class 1 traffic. We note that the number of HDTV streams from node i to the other N-1 nodes may not be the same, since the destination of each stream is randomly selected. In view of this, the term H_{ij} is normalized by dividing it by the number of HDTV streams between node i and node j.

Figure 20 plots the throughput fairness index of the protocols (considering only class 1 traffic) versus the number of HDTV streams per node. Since all the protocols have a value very close to zero, all of them are throughput fair. If the throughput fairness index is calculated considering all the three classes, the protocols would still be fair because the protocols themselves do not distinguish between closer and farther nodes.

The *delay fairness index* of a node is defined as follows:

$$\text{Delay Fairness Index of Node i} = \left(\sum_{j=1, j\neq i}^{10} (W_{ij} - \overline{W_i})^2 \right) \times \frac{1}{\overline{W_i}^2} \tag{5.4}$$

where W_{ij} is the mean frame delay from node i to node j, and $\overline{W_i} = (\sum_{j=1,j\neq i}^{10} W_{ij})/9$. Here, the mean frame delay counts only the queueing delay that all frames in queue j of node i experience. The delay fairness does not include the propagation delay. The delay fairness index of the protocol is defined as the average of the delay fairness indices of all the nodes.

Figure 21 plots the delay fairness index of the protocols (considering only class 1 traffic) versus the number of HDTV streams per node. It shows that none of the protocols are delay fair. This is not due to the physical positioning of the nodes around the ring, because as mentioned earlier, the protocols do not distinguish between closer and farther nodes. Due to

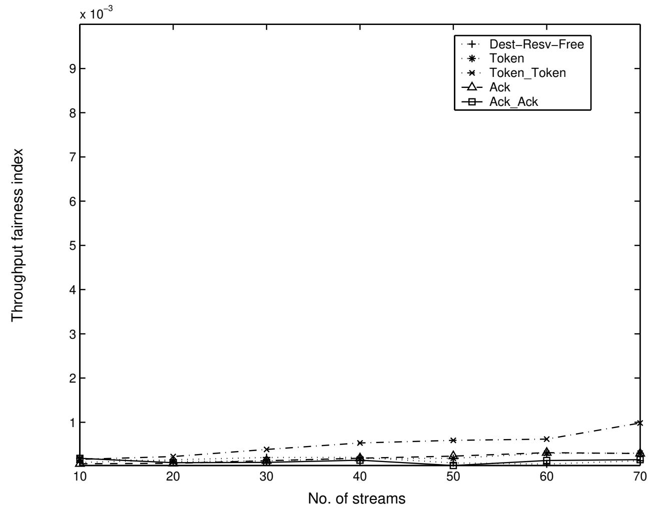

Figure 20: Throughput fairness index of the protocols (1 transceiver)

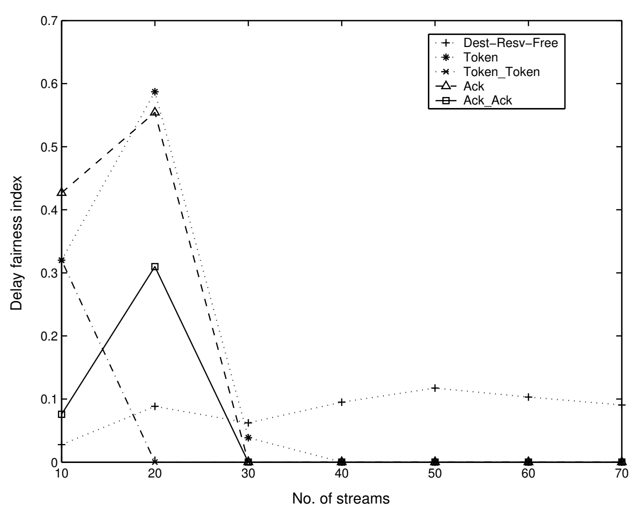

Figure 21: Delay fairness index of the protocols

the asymmetric traffic pattern of class 1, node i may have more streams arriving than node j. In this case node i would take more time sending the acknowledgements than node j in case of acknowledgement based protocols. In token based protocols, since a token can be used to transmit a single burst and is then released, nodes with more number of streams to a particular destination translates to frames of these streams experiencing higher queueing delays because of increased waiting time to obtain the token to that particular destination. Thus, difference in delays experienced by class 1 bursts are due to the traffic pattern and not due to the physical positioning of the nodes with respect to each other.

6 Conclusion

Five different protocols for the support of different service classes on an optical burst switched ring network were proposed and their performance was evaluated through simulation. Dest-Resv-Free protocol was based on the best-effort type of service. Ack and Token provided guaranteed delivery to class 1 bursts but do not ensure zero burst loss for classes 2 and 3. These two protocols can support more HDTV streams per node than other protocols because of guaranteed delivery only to class 1 bursts. Ack-Ack and Token-Token provided zero burst loss delivery for both classes 1 and 2 of traffic and they ensured that class 1 bursts are received in time and without too much jitter. Class 3 received best-effort service. For the single home wavelength case, Ack-Ack provides better service to class 1 traffic than Token-Token. But, the Token-Token and Token protocols perform better than their acknowledgement-based counterparts when multiple home wavelengths are available for transmission. Additional simulations experiments performed indicate that token based protocols are more scalable with ring size and number of nodes in the ring.

References

[1] T. Battestilli and H. G. Perros. An introduction to optical burst switching. *IEEE communications optical magazine*, 41:510–515, 2003.

[2] N. Bouabdallah, A-L. Beylot and G. Pujolle. Collision avoidance and fairness issues in metropolitan bus-based optical access networks. Technical report, Paris 6 university, 2003.

[3] A. Fumagalli and P. Krishnamoorthy. A low-latency and bandwidth-efficient distributed optical burst switching architecture for metro ring. *In Proceedings of IEEE ICC*, 2:1340–1344, 2003.

[4] James Jong. Multicast access protocols in an optical burst switched WDM ring network. Master's thesis, NC State University, 2002.

[5] H. G. Perros. *Computer Simulation Techniques: The definitive introduction.* Available for free download from http://www.csc.ncsu.edu/faculty/perros/hp.html.

[6] L. Xu, H. G. Perros and G. N. Rouskas. A simulation study of optical burst switching access protocols for WDM ring networks. *Computer Networks*, 41(2):143–160, January 2003.

[7] S. Trevitt. Traffic patterns in fiber channel fabrics, 2002. Available from World Wide Web: http://www.mcdata.com/downloads/mkt/wpaper/TrafficPatterns.pdf.

[8] Bragg, A. W. *MPEG-4 Analysis.* private e-mail communication, March 2003.