Information Sciences 149 (2003) 75–81

INFORMATION SCIENCES
AN INTERNATIONAL JOURNAL

www.elsevier.com/locate/ins

Access protocols for optical burst-switched ring networks

Lisong Xu, Harry G. Perros, George N. Rouskas *

Department of Computer Science, College of Engineering, North Carolina State University, P.O. Box 7534, Raleigh, NC 27695-7534, USA

Received 27 September 2001; accepted 8 May 2002

Abstract

In this paper, we consider a WDM metro ring architecture with optical burst switching. Several access protocols are proposed and their performance is analyzed by simulation.

Keywords: Optical burst switching; Ring; MAN; Access protocols

1. Introduction

Wavelength division multiplexing (WDM) appears to be the solution of choice for providing a faster networking infrastructure that can meet the explosive growth of the Internet. Since this growth is mainly fueled by IP data traffic, wavelength-routed optical networks [2] which employ circuit switching may not be the most appropriate for the emerging optical Internet. Optical packet switching [5,6] is an alternative technology that appears to be the optimum choice. However, at this moment the technology is not mature

* Corresponding author. Tel.: +1-919-515-3860; fax:+1-919-515-7896.

E-mail addresses: lxu2@csc.ncsu.edu (L. Xu), hp@csc.ncsu.edu (H.G. Perros), rouskas@csc.ncsu.edu (G.N. Rouskas).

0020-0255/02/$ - see front matter
PII: S0020-0255(02)00247-5

enough to provide a viable solution. Optical burst switching (OBS) [3] is a switching technique that occupies the middle of the spectrum between the circuit switching and packet switching paradigms. The unit of transmission is a *burst*, which may consist of several packets. The transmission of each burst is preceded by the transmission of a control packet, which usually takes place on a separate signaling channel. We will refer to the interval of time between the transmission by the source node of the first bit of the control packet and the transmission of the first bit of the data burst as the *offset*.

There are several variants of burst switching, mainly differing on the length of the offset. In the burst switching scheme called *Tell And Go* (TAG) [4], the burst is transmitted immediately after the control packet. This scheme is practical only when the switch configuration time and the switch processing time of a control packet are very short. At the other extreme, the *Tell and Wait* (TAW) [4] scheme requires the offset to be at least equal to the time required to receive an acknowledgement from the destination. An intermediate burst switching scheme, known as *Just Enough Time* (JET) [3], selects the offset to be long enough to account for the processing time of the control packet at the intermediate nodes and the destination plus the switch setup time at the destination.

One issue that arises in computing the offset under JET is determining the number of intermediate switching nodes (hops) between the source and destination. However, the number of hops in a path may not, in general, be readily available. Given the recent advances in hardware implementation of communication protocols, it is reasonable to assume that the processing time at the intermediate nodes will be very short for most common functions of the signaling protocol. In this case, fiber delay lines of reasonable length may be used at intermediate nodes to delay each incoming burst by an amount of time equal to the processing time. Given such fiber delays, the offset could be only longer than the processing time of the control packet at the destination plus the switch setup time at the destination. We call this new scheme the *Only Destination Delay* (ODD). Furthermore, instead of using destination-specific values for the processing and switching delays, one may use a constant offset value by taking the maximum of these values over all destinations.

In this paper we study burst switching protocols for ring networks. To the best of our knowledge, this is the first study of burst switching protocols specifically for ring networks.

This paper is organized as follows. Section 2 describes the ring network we consider and the basic operation of burst switching in such an environment. Section 3 provides a detailed description of the various burst switching access protocols studied in this paper. Section 4 presents the simulation results on the performance of these burst switching access protocols, and finally Section 5 provides some concluding remarks.

2. The network under study

2.1. Ring and node architecture

We consider N OBS nodes organized in a unidirectional ring. The ring can be a metropolitan area network (MAN) serving as the backbone that interconnects a number of access networks. Each fiber link between two consecutive OBS nodes in the ring can support $N + 1$ wavelengths. Of these, N wavelengths are used to transmit bursts, and the $(N + 1)$th wavelength is used as the control channel.

Each OBS node is attached to one or more access networks. In the direction from the access networks to the ring, the OBS node collects and buffers electronically data packets, transmitted by users over the access networks. Buffered packets are subsequently grouped together and transmitted in a burst to the destination OBS node. A burst can be of any size between a minimum and maximum value. Bursts travel as optical signals along the ring, without undergoing any electro-optic conversion at intermediate nodes. In the other direction from the ring to the access networks, an OBS node terminates optical bursts destined to itself, electronically processes the data packets contained therein, and delivers them to users in its attached access networks.

Each node is equipped with two pairs of optical transceivers. The first pair consists of a receiver and a transmitter fixed tuned to the control wavelength. The second pair of transceivers consists of a transmitter that is fixed tuned to the node's *home wavelength*, and an agile receiver (or a receiver array) that can receive from all N wavelengths that transmit bursts.

To support ODD, an extra fiber delay line is added into the node to delay outgoing data on all wavelengths except the control wavelength and the node's home wavelength.

Packets waiting for transmission are organized into (logical) transmit queues according to their destination. In this paper, the transmit queues are served in a Round-Robin manner.

2.2. Control wavelength operation

The control wavelength is used for the transmission of control slots. In a ring with N nodes, N control slots, one for each node, are grouped together in a *control frame* which continuously circulates around the ring. Depending on the length of the circumference of the ring, there may be several control frames circulating simultaneously. In this case, control frames are transmitted back-to-back on the control wavelength.

Each node is the owner of one control slot in each control frame. Each control slot contains several fields. The format and type of the fields depend on the OBS protocol used. In general, however, each control slot includes fields for the destination address, the offset, and the burst size. Other fields, such as a token field, may be included for some of the protocols, as necessary.

3. Access protocols

In this paper, we proposed a number of different access protocols that differ mainly in the way that receiver conflicts are resolved. Our emphasis is on protocols that use few rules, are simple to implement in hardware and are distributed in nature.

3.1. Round-robin with random selection (RR/R)

At the transmitting side, node i visits all transmit queues in a round-robin fashion. Suppose that, transmit queue j is selected for service. When the next frame arrives, node i writes the burst information and destination address j in its own control slot. After a delay equal to the offset value, node i transmits the burst on its home wavelength. At the receiving side, when a control frame arrives at node i, it scans the control slots of the control frame, checking for any slot that has i in the destination address field. If there are collisions, one burst is selected randomly.

3.2. Round-robin with persistent service (RR/P)

The *Round-Robin with Persistent Service* (*RR/P*) protocol is similar to the RR/R protocol, but it is designed to eliminate receiver conflicts that can be detected prior to the transmission of a burst. At the transmitting side, node i maintains a variable for each destination node j, which specifies the earliest time at which the receiver of node j would be free. This variable is updated by monitoring the burst information in control slots that have j in the destination address field. Node i also computes the time that its burst would arrive at node j. If node i determines that transmitting a burst to j would result in a collision, then node i does not transmit the burst; instead it waits for the next control frame and repeats the process of transmitting the burst to node j. This is the *persistent* feature of the protocol. However, RR/P does not altogether eliminate receiver collisions, since two nodes may simultaneously determine (based on information they read in *different* control frames) that it is safe to send a burst to some destination. At the receiving side, the operation of the protocol is identical to RR/R.

3.3. Round-robin with non-persistent service (RR/NP)

The operation of the *Round-Robin with Non-Persistent Service* (*RR/NP*) protocol is identical to the operation of the RR/P protocol with one exception.

If node i determines that transmitting a burst to j would result in a collision, then instead of continuing its attempt to serve transmit queue j, it proceeds to serve the next transmit queue upon arrival of the next control frame. As in RR/P, RR/NP does not completely eliminate receiver collisions.

3.4. Round-Robin with tokens (RR/Token)

This protocol uses tokens to resolve receiver collisions at the receivers. There are N tokens, one for each destination node. A token may be either available or in use. The status of token j is indicated in the token field of the jth control slot. A node can only transmit to a destination node j, if it captures the jth token. Because of the token operation, there will be at most one burst transmission arriving at a destination node at any time. That is, RR/Token is a receiver collision free protocol.

3.5. Round-Robin with acknowledgement (RR/ACK)

The *Round-Robin with Acknowledgment* (*RR/ACK*) protocol is based on the Tell and Wait (TAW) scheme [4]. A source node i first sends a request (including destination and size) to transmit a burst to the destination node j. When node j receives the request, it calculates an offset value, and sends it back to node i in the offset field of control slot i. After node i receives the acknowledgment from node j one round-trip time later, it instructs its transmitter to send out the burst at the time specified by node j in the acknowledgement. We note that a source node is not allowed to have more than one outstanding request. RR/ACK is a receiver collision free protocol.

4. Numerical results

In our simulation study we consider a ring network with 10 nodes, each with an electronic buffer of 10 MBytes. The distance between two successive nodes in the ring is taken to be 5 km. The control wavelength runs at 622 Mbps, while each burst wavelength runs at 2.5 Gbps. Each control slot in a control frame is 100 bytes long regardless of the protocol used in the ring. The processing time of a control frame at both the intermediate and destination nodes is set to be 10 slot times, and the setup time at the destination nodes is 1 μs.

We model the packet arrival process to each node by a modified Interrupted Poisson Process (IPP) [1], which is an ON/OFF process. Packets arrive back to back during the ON period at the rate of 2.5 Gbps. The packet size is assumed to follow a truncated exponential distribution with an average size of 500 bytes and a maximum size of 5000 bytes. We use the squared coefficient of variation, c^2, of the packet inter-arrival time to measure the burstiness of the arrival process. In the simulation, c^2 is set to 20. Packets arriving at a node are assigned a destination node following the uniform distribution.

In the simulation, we vary three variables: `AvgArrivalRate` of IPP which is the average arrival rate during both the ON and OFF periods, and `MaxBurstSize` and `MinBurstSize` which specifiy the maximum and minimum burst size that can be transmitted on the ring, respectively. For each possible combination of access protocols and offset calculation methods, we ran three simulations. First, we varied the `AvgArrivalRate` to each node from 0.5 to 2.0 Gbps with an increment of 0.3 Gbps, and set `MaxBurstSize` and `MinBurstSize` to 112 and 16 KB, respectively. Second, we varied `MaxBurstSize` from 32 to 112 KB with an increment of 16 KB, and set `MinBurstSize` to 16 KB, and set the `AvgArrivalRate` to 1.7 Gbps. Third, we varied `MinBurstSize` from 16 to 96 KB with an increment of 16 KB, and set `MaxBurstSize` to 112 KB, and set the `AvgArrivalRate` to 1.7 Gbps.

In Section 4.1 we present a comparison of the performance of the RR/R, RR/P, RR/NP, and RR/Token protocols with ODD offsets. In Section 4.2 we discuss the impact of JET versus ODD on the performance of those four protocols. Finally, in Section 4.3 we compare RR/ACK that uses the TAW scheme with RR/Token that uses the ODD offset.

4.1. Performance of protocols with ODD offset

Simulation results show that, with ODD, RR/Token achieves the highest mean node throughput, followed by RR/P, RR/NP and RR/R. RR/R has the smallest mean packet delay, followed by RR/NP, RR/P and RR/Token. RR/R also has the smallest mean buffer requirement, followed by RR/NP, RR/P and RR/Token. The burst loss rate due to receiver collisions for the protocols which are not receiver collision free depends on the squared coefficient of variation of the burst size. Only RR/R is a delay fair protocol, which means that the queueing delay of a packet is insensitive to the relative position of the source and destination nodes in the ring. Both RR/R and RR/Token are throughput fair protocols, that is, for a specific source node, the throughput to each destination node does not depend on the location of the destination node in the ring.

4.2. JET vs. ODD

Simulation results show that, compared to ODD, JET leads to a longer mean packet delay for all protocols, which in turn leads to a larger mean buffer requirement, and to a larger packet loss rate due to buffer overflow. Therefore, as a receiver collision protocol, RR/Token has a lower mean node throughput with JET than with ODD. Moreover, JET naturally leads to delay unfair protocols, but does not change the throughput fairness property of the protocols.

4.3. TAW vs. ODD

Simulation results show that when the `MaxBurstSize` is small, RR/Token with ODD gets both a higher mean node throughput and a lower mean packet delay than RR/ACK with TAW. When the `MaxBurstSize` is large, in most cases, both protocols have similar mean node throughput and mean packet delay. However, when both the `AvgArrivalRate` and the `MaxBurstSize` are very large, RR/ACK gives a higher mean node throughput and lower mean packet delay than RR/Token. As for fairness, RR/ACK is a throughput fair protocol, but not a delay fair protocol.

5. Concluding remarks

This paper described a WDM metro ring architecture with optical burst switching. Several access protocols are proposed and their performance is analyzed by simulation. Depending on the specific requirement, different protocols may be used. For example, if we want a protocol with a small burst loss rate due to receiver collisions and a mean packet delay which is not very large, we may choose RR/P. If a short optical delay line is available, then ODD is a good choice.

References

[1] W. Fischer, K. Meier-Hellstern, The Markove-Modulated Poisson Process (MMPP) cookbook, Performance Evaluation 18 (1992) 149–171.
[2] B. Mukherjee, Optical Communication Networking, McGraw-Hill, New York, 1997.
[3] C. Qiao, M. Yoo, Optical burst switching (OBS) – a new paradigm for an optical Internet, J. High Speed Networks 8 (1) (1999) 69–84.
[4] I. Widjaja, Performance analysis of burst admission control protocols, IEEE Proc. Commun. 142 (February) (1995) 7–14.
[5] L. Xu, H.G. Perros, G.N. Rouskas, Techniques for optical packet switching and optical burst switching, IEEE Commun. 39 (1) (2001) 136–142.
[6] S. Yao, S. Dixit, B. Mukherjee, Advances in photonic packet switching: an overview, IEEE Commun. 38 (2) (2000) 84–94.