Switched Optical Star-Topology Network with Edge Electronic Buffering and Centralized Control

Dr. Xianhui Che
School of Applied Computing
Faculty of Applied Design and Engineering
Swansea Metropolitan University, Swansea, UK
xianhui.che@sihe.ac.uk

Dr. David K. Hunter, Prof. Ian D. Henning
Department of Computing and Electronic Systems
University of Essex
Colchester, UK
dkhunter@essex.ac.uk

***Abstract*—A star-topology local area network is designed to carry optical packet traffic, employing collision-free switching. Data forwarding is under centralized control based on scheduled resource reservation, eliminating the need for optical buffers in the central switch. Packets are stored in electronic buffers within the end stations prior to their electronic-to-optical conversion and transmission in the network. With the combination of central control and the edge buffering, the system forms a stable, efficient and low-cost network with smoothed traffic. Modeling and simulations have been carried out using OPNET modeler to evaluate network performance with TCP/IP traffic.**

Optical Packet Switching; Local Area Networks

I. INTRODUCTION

With the rapidly growing demand for capacity and speed, the study of local area networks (LANs) has become more and more relevant and widespread. Optical LANs show appealing merits by permitting end users to access the network via the entire optical path capacity. OPS (*Optical Packet Switching*) technology utilizes the tremendous capabilities of optical networks for flexible, fast and capacity-efficient data transmission [1]. Most research towards applying OPS to local area networks is based on DWDM (*Dense Wavelength Division Multiplexing*) which involves issues such as expensive devices (e.g. AWG, couplers, etc.), split loss, the constraint of limited number of wavelength etc. [2, 3]. Currently packet switching is not well matched with state-of-the-art optical technology, mainly due to the lack of optical memories and optical processing techniques. To avoid the problems just mentioned and in order to realize optical networking with the components that will be available in the short term, we combine optics and electronics to design a centralized control protocol for optical packet transmission in the data link layer, which eliminates the need for optical buffers. Unlike DWDM, where each individual wavelength works as a channel or path to transport data, every data packet is spread over the entire link capacity; statistical multiplexing in the time domain is employed.

The protocol design takes advantage of the concept underlying IEEE 802.11 [4], but also differs from it in a number of aspects. In both cases, RTS (*Request-to-Send*) and CTS (*Clear-to-Send*) messages are used. Prior to data transmission, end stations need to send RTS messages to request resources; once the data channels are established successfully, CTS messages are returned to the end stations as confirmations for them to send the actual data. In IEEE 802.11 wireless LANs, the control messages and data share the same transmission media, and broadcast is necessary, which is not the case with the protocol described in this paper. In this optical network, a separate channel (i.e. wavelength) is designated for the control messages. Moreover, instead of sending only one RTS every time, we allow end stations to send several at once, yielding a message that contains the queue of requests in time order. The corresponding CTS messages are sent back one by one to the end stations after successful resource reservation for each request. In the interests of simplicity, segmentation is not implemented in our simulations, although it could be implemented in a real system. Each data packet delivered from the network layer to the MAC layer occupies one slot, even if it is not completely filled. With network resources reserved ahead of time, the network offers collision-free packet switching.

In this paper, Section II briefly describes the system design; Section III discusses the issues surrounding traffic smoothing in the proposed network; Section IV describes the performance evaluation based on the network modeling and simulations; Section V concludes the paper.

II. SYSTEM DESIGN

The network's end stations form a star topology (Figure 1), connecting via edge interfaces to the concentration point (the hub switch – a black-box switch controlled by an arbiter).

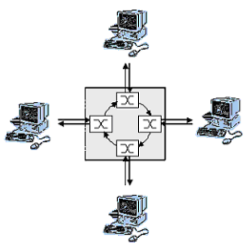

Figure 1. Network scenario

In reality the edge interface can be implemented in line cards within the workstations, which embed the designated protocol. The electronic buffer is integrated in the edge interface. The network control plane in implemented by the edge interface and the arbiter.

The hub consists of a group of 2×2 switches interconnected to form a ring topology. Each 2×2 switch has seven possible states that can be indicated by three binary digits, as shown in Figure 2. The upper ports in the figure are used for the connections to other switches in the hub; the lower ports are linked to the end stations. The term “add” means that the end station transmits data; and “drop” means that the end station receives data. “Though” indicates that data just passes through the switch. “Loopback” enables the end station to talk to itself; for example, sometimes an end station needs to send a message to the network for operation and maintenance purposes. In states 011, 110, 100, 001 and 000, some ports are disabled. It is necessary for idle lines to be disconnected in these cases, otherwise noise and crosstalk could interfere with ongoing communications, especially when optical amplifiers are part of the optical network infrastructure and idle connections may lead to increased noise.

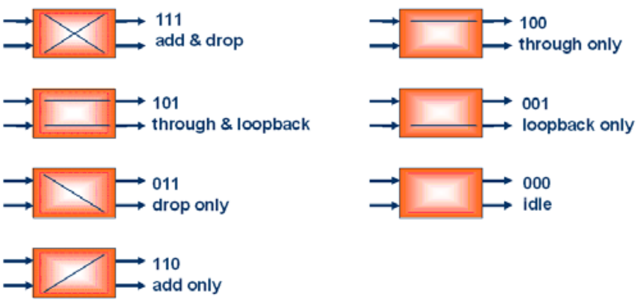

Figure 2. States of the switches

New arriving data are stored in the electronic buffers inside the interfaces until the corresponding CTS messages are received. Then they are converted into optical packets and are transmitted into the network. All the request queues that arrive at the arbiter comprise a matrix called the Request Array. Assume the request queue length is R and that the number of nodes is N. Figure 3 shows an example of the request array for a LAN with four nodes (A, B, C, and D). Each row of the request array stands for the requests from one node, which lists the destination addresses. In this case, each node can send up to five requests at once. If there are fewer than five requests during the period, the remaining fields of request array are filled with zeroes.

Request Array ($t \rightarrow$)

0	C	C	B	D
A	C	A	D	C
0	D	A	A	B
0	0	0	B	A

NAV ($t \rightarrow$)

110	000	100	011	…	100
100	110	011	000	…	100
101	011	110	000	…	011
011	000	100	110	…	110

Figure 3. Request array and NAV (R = 5, N = 4)

There is a NAV (*Network Allocation Vector*) inside the hub which contains the state codes of all the switches, and is updated by the arbiter based on the request array. The NAV columns correspond to the requests in time order of being served. The arbiter periodically extracts the information from the NAV, sends out CTS messages, and configures each optical switch with the respective state codes, so that they are all configured properly to form a clear end-to-end path when the data packet arrives.

III. TRAFFIC SMOOTHING

The packet arrival process cannot be captured by Markovian traffic models [5] due to the properties of self-similarity and long-range dependence. Traffic smoothing on the timescale of individual packets involves shaping the traffic fluctuation on the packet level. When there is congestion in the core network, traffic smoothing should endeavor to maximize resource allocation on the packet level. To do this for unslotted traffic, each wavelength must carry consecutive packets. On the other hand, for slotted traffic, when congestion happens, the network should fill every slot with a packet to maximize the resource utilization. Figure 4 [6] depicts under-utilized capacity in both cases; shaded blocks stand for actual packets and unshaded blocks represent idle capacity or slots.

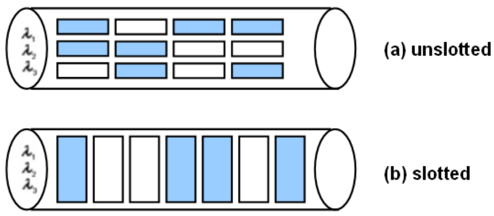

Figure 4. Slotted/unslotted traffic within under-utilized links

When the traffic demand becomes intensive, traffic smoothing re-distributes the packets evenly over time in order to obtain the maximum occupancy, implying maximum resource utilization. When smoothing unslotted traffic and using a conventional WDM network, apart from the problems of complexity and expense, resource allocation involves many issues, such as limited available wavelengths, synchronization, wavelength conversion, etc. Fortunately the protocol proposed here avoids these problems by employing slotted traffic. The smoothing buffer at the edge accommodates and schedules the stochastic packet flows in response to the resource allocation status in the hub switch.

The design of the packet format and networking protocol adopted here facilitates traffic smoothing. Upon the successful exchange of control messages (i.e. RTS and CTS), the data packets transmitted in the network form a slotted traffic stream. In this format, slots over a fiber link are allocated serially. As discussed, the request array may not be fully filled with requests, so when the arbiter updates the NAV based on the request array, there may be several columns left empty, which imply empty slots in the traffic stream on the fiber link. This problem can be easily solved through minor modifications to the system. The changes can be made in two ways. If the arbiter can skip the processing of empty columns in the NAV, there are no empty slots left in the traffic stream as shown in Figure 5. Alternatively, when the NAV is being updated by the arbiter based on the request array, it can simply leave out the empty fields so that there is no empty column created in the NAV. Either way, the network carries smoothed slot traffic. Unlike traditional TDM (*Time Division Multiplexing*), this approach takes advantage of the concept of statistical multiplexing [7].

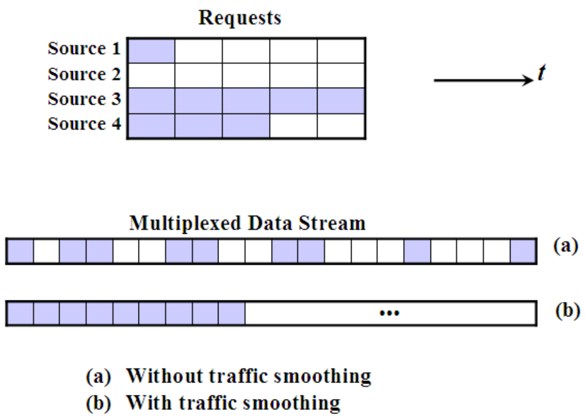

Figure 5. Traffic smoothing for slotted traffic through statiscal muliplexing

IV. PERFORMANCE EVALUATION

A 100 m four-node LAN is modeled with OPNET [8] providing FTP over TCP. The application inter-request times are uniformly distributed between 0 and 5 ms, and the file sizes are all set to 50 kbytes. Based on the principle of this protocol, there are three parameters that can affect performance for the same type of LAN, i.e. request queue length (*R*), arbiter timer granularity, and edge buffer length. Figure 6 shows the relationship between these three parameters. When the arbiter processing speed increases, there are fewer packets queuing in the edge buffer. Longer request queues result in faster servicing, hence shorter queues in the edge buffer.

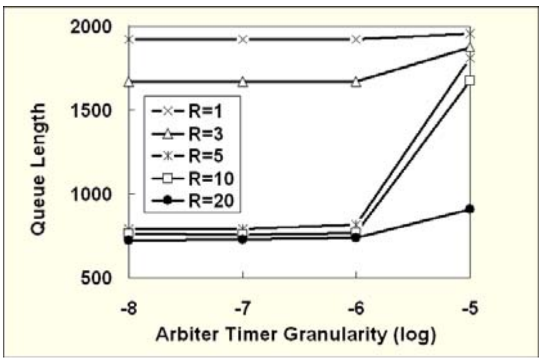

Figure 6. Buffer occupancy, request queue and arbiter timer

In TCP, retransmission is triggered by lost segments. In this collision-free switching network, the only place for packet loss is the edge buffer. Figure 7 and Figure 8 show the TCP retransmission count over a period of network simulation, with different lengths of edge buffers and different length of request queue respectively. With the same traffic intensity, bigger buffers can accommodate more packets to be scheduled, so there is less packet loss, hence fewer retransmissions. Figure 8 indicates that the probability of an edge buffer overflowing is greatly decreased if more requests can be sent at once, as expected. The handshake implemented by RTS/CTS imposes a major overhead on actual data transmission. Longer request queues reduce the redundancy of data transmission, hence improving network efficiency.

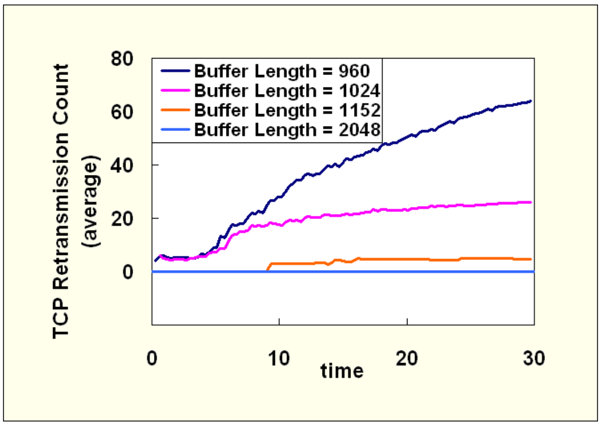

Figure 7. TCP retransmission count with different buffer capcities

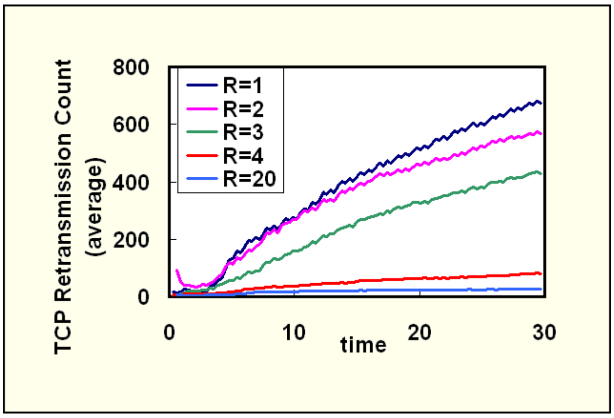

Figure 8. TCP retransmission count with different request queue lengths

To permit better evaluation of the protocol, a reference model with the same network topology and scale is also evaluated by simulation. The reference model has a hub much like a conventional router or switch but utilizes the optical capacity. In the reality this model does not exist, as issues like optical buffering, O-E-O conversion and optical processing are ignored. The reference model statistics represent what the ideal performance of the network would be.

TCP segment delay is the end-to-end delay of segments received by the TCP layer across all nodes. Figure 9 shows the mean TCP segment delay with increasing request queue length and with different buffer lengths. Longer request queues improve efficiency, hence the delay is lower. Although larger buffers can accommodate more packets so as to reduce the packet drop rate as shown in Figure 7, they also imply greater queuing delay; hence the increased edge buffer capacity increases network latency.

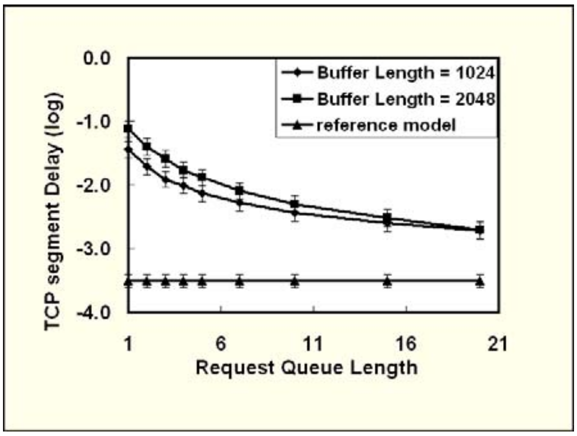

Figure 9. Mean TCP segment delay vs. request queue length

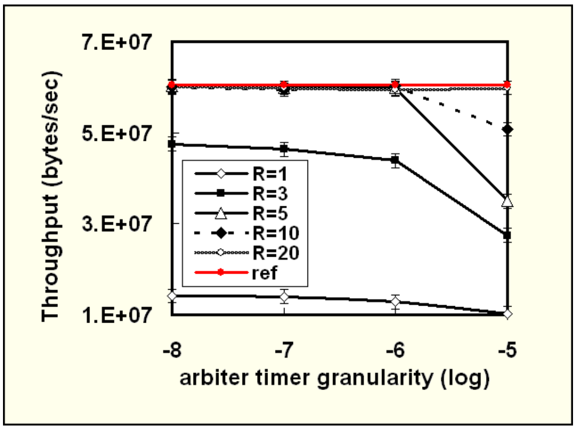

Figure 10. TCP Throughput with different request queue lengths and arbiter timer interval

Figure 10 demonstrates the impact of the arbiter timer interval as well as the request queue length on TCP throughput. Figure 9 and Figure 10 conclude that network has steady performance when the arbiter interval granularity is under 1 μs and request queue length is greater than four slots. Any variation beyond this would lead to visible performance change.

V. CONCLUSIONS

A star-topology local area network with a specially designed protocol has shown to be promising for the future implementation of optical switched LANs with low cost, simplicity, and high performance. Successful exchange of control messages in a reserved channel guarantees collision-free packet switching, which eliminates the need for optical buffers in the central switch. The use of edge electronic buffering combined with centralized control facilitates effective traffic smoothing. Network modeling and simulations have evaluated network performance when varying different parameters, e.g. request queue length, edge buffer length, network arbiter timer interval, etc. The reference model delivers ideal performance, with which the designated protocol's performance can be compared. The simulation results are all as expected, analyzing the network from different perspectives, while indicating that the network is stable.

REFERENCES

[1] L. Xu, H. Perros, "Techniques for Optical Packet Switching and Optical Burst Switching", *IEEE Communications Magazine*, pp136-142, January 2001.

[2] H. Woesner, "All-optical LAN architectures based on Arrayed Waveguide Grating Multiplexers", *Proceedings of All-Optical Networking*, pp158-163, November 1998.

[3] A. Grah, "Packet-Switched Local Area Networks Using Wavelength-Selective Station Couplers", *IEEE/ACM Transactions on Networking*, Vol 8, Issue 2, pp251-264, April 2000.

[4] *Wireless LAN MAC and PHY specifications*, IEEE Std 802.11b-1999.

[5] P. Salvador, A. Nogueira, R. Valadas, "Modeling Local Area Network Traffic with Markovian Traffic Models", *ConfTele*, Figueira da Fez, Portugal, 23-24 April 2001.

[6] Z. Lu, D. K. Hunter, I. D. Henning, "Contention Reduction in Core Optical Packet Switches through Electronic Traffic Smoothing and Scheduling at the Network Edge", *IEEE/OSA Journal of Lightwave Technology*, vol. 24, no. 12, pp4828-4837, December 2006.

[7] B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson, J. D. Robbins, "Performance Models of Statistical Multiplexing in Packet Videocommunications", *IEEE Transactions on Communications*, vol.36, no.7, pp834-844, July 1988.

[8] http://www.opnet.com – OPNET modeler