El-Gorashi, Taisir E.H., Pranggono, Bernardi, Mehmood, Rashid and Elmirghani, Jaafar M.H.

A Mirroring Strategy for SANs in a Metro WDM Sectioned Ring Architecture under Different Traffic Scenarios

Original Citation

El-Gorashi, Taisir E.H., Pranggono, Bernardi, Mehmood, Rashid and Elmirghani, Jaafar M.H. (2008) A Mirroring Strategy for SANs in a Metro WDM Sectioned Ring Architecture under Different Traffic Scenarios. Journal of Optical Communications, 29 (2). pp. 89-97. ISSN 0173-4911

J. Opt. Commun. 29 (2008), 89-97

A Mirroring Strategy for SANs in a Metro WDM Sectioned Ring Architecture under Different Traffic Scenarios

Taisir E.H. El-Gorashi, Bernardi Pranggono, Rashid Mehmood[1], Jaafar M. H. Elmirghani

Summary

In this work, a novel data mirroring technique for storage area networks (SANs) is introduced in a metropolitan area wavelength division multiplexing (WDM) sectioned ring scenario. Simulation is carried out to evaluate the performance of the network with the proposed mirroring method for 16 and 24 nodes under two different traffic models– *Poisson* and *self-similar*. Simulation results show four network parameters including node throughput, queuing delay, transmission buffer packet dropping probability and receiver packet dropping probability. A modified version of the MAC protocol that improves mirroring time and bandwidth usage is also introduced and evaluated.

1 Introduction

Today, storage area networks (SANs) [1, 2], the most common storage management structure, are an important part of the modern networking infrastructure. The day when storage systems were expected only to store and retrieve randomly accessible data is long gone. Today storage systems are expected to play an integral role in supporting high levels of data availability in the event of disasters such as flood, fire, earthquake, power outage, and terrorist attacks. The concept of backup is vital to improve reliability of data storage systems. Backup methods are usually based on *data replication*, also known as *data mirroring* [3-5]. In data mirroring, exact replicas of the original data are created to be sent to secondary storage systems in far locations. In the event of data loss or corruption in the primary storage site, data is retrieved from the secondary storage sites. This becomes possible because of storage protocols that operate over extended distances such as Internet SCSI (iSCSI), Fibre Channel over TCP/IP (FCIP), and Internet Fibre Channel Protocol (iFCP) [6]. Storage data can travel over metropolitan area networks (MANs) or wide area networks (WANs) for large distances. The challenge is to design a network which provides reliability, scalability in terms of distance and number of nodes, high throughput, and full accessibility [7].

Remote mirroring is usually implemented by one of two strategies - synchronous and asynchronous [8]. Mirroring is considered to be *synchronous* if the data is transmitted from the transmitting node to the two SAN nodes at the same time; i.e. the states of the two SAN nodes are synchronized. However, in addition to the high bandwidth requirements, synchronization can introduce significant delays for the network if the SANs are separated by large distances. On the other hand, mirroring is considered to be *asynchronous* if initially data is only transmitted to the primary storage location, and then the primary location replicates it to the secondary storage location. Usually asynchronous mirroring is scheduled to run after peak hours when the network bandwidth is idle to save peak hour bandwidth. Therefore asynchronous backup is efficient and cost effective. However asynchronous mirroring cannot always keep the secondary location updated with the most recent state of the data which makes it unsuitable for critical applications.

In an earlier work [9], a metropolitan area WDM sectioned ring architecture with a single SAN node was proposed and evaluated. In this paper, a novel mirroring technique for the SAN node, which is essential for reliability, is introduced and evaluated. In this technique, a secondary SAN node with the same capacity as the primary SAN node is introduced to the network where packets are automatically transmitted to the two SAN nodes. A modified version of the MAC protocol is also introduced and evaluated. In this work, the focus is to provide support for storage replication in MAN ring network scenarios, and therefore we are not concerned here if the two SANs are synchronized with each other in real time. Inevitably transmission and queuing delay are expected in the process and therefore the SANs are synchronized in a fashion that is close to real time as is possible.

2 Network architecture

The network architecture considered, illustrated in Fig. 1, is a metropolitan area ring network with a unidirectional (clock wise) multi-channel slotted single fiber. The network connects a number of access nodes within a circum-

Address of authors:

University of Leeds
School of Electronic and Electrical Engineering
Leeds LS2 9JT, UK
Email: eltemh@leeds.ac.uk

[1] University of Wales Swansea
Institute of Advanced Telecommunications
Swansea, SA2 8PP, U.K.

Received 31 October 2007; accepted 5 December 2007

ference of 138 km (a typical metro reach, and a circumference results in an integer number of Ethernet MTU slots at the data rates considered under typical propagation delays). There are two types of nodes in the network: *ordinary* nodes and SAN nodes. Each node has add-and-drop capabilities to access the ring slots and is used to connect a local area network (or access network) to the ring via Gigabit Ethernet (GbE) links as the vast majority of LANs are based on Ethernet. However, the network can equally use other medium access technologies. These links operate at 1 Gbit/s for ordinary nodes and at 5 Gbit/s for SAN nodes. A fixed size slot equal to the Ethernet maximum transfer unit (MTU) frame size (1,500 bytes) is used in the ring. The ring is sectioned into two parts with equal number of nodes using a point-to-point communication link (44 km) passing through the center of the ring. This link consists of a pair of fibers with opposite propagation directions and directly connects two nodes. The section provides a shorter path for some source-destination pairs to communicate instead of going through the entire ring and helps deal with traffic asymmetry and the hot-node scenario created by the SAN node. Without mirroring, there is only one SAN node. With the introduced mirroring technique, a secondary SAN node is used as backup for the primary SAN.

Figure 2 illustrates the logical topology for the sectioned ring architecture. The architecture could be considered as three logical rings: an outer ring, connecting all the nodes, and two inner rings, one connecting nodes in the upper part of the network, and the other connecting nodes in the lower part of the network. The logical ring for each source-destination pair over the time slotted ring is chosen according to a shortest path algorithm. To demonstrate the effects of network loading, asymmetric traffic, hot node scenario and the impact of the ring section, a small number of wavelengths is used. Two wavelengths are dedicated for each ring. Each wavelength operates at a 2.5 Gbit/s transmission rate and the total number of nodes considered is 16. A higher number of nodes could be considered, however 32 and 64 nodes in a metro setting may be high..

Each ordinary node is equipped with two fixed transmitters to connect it to the two logical rings available to it, A & B or A & C depending on its location. SAN nodes placed at the sectioning points require an additional transmitter as they have access to the three logical rings. As the number of nodes is greater than the number of wavelengths, each wavelength is shared by a number of nodes for transmission. This can be accommodated since the wavelength rate is greater than the node transmission rate. Each node is equipped with four fixed receivers to share wavelengths for reception which results in higher scalability compared to the single fixed receiver architecture where the number of wavelengths is equal to the number of nodes. The use of a tunable receiver is possible but results in receiver collisions where multiple wavelengths carry packets in a given time slot all destined to the same node. Receiver collision mitigation by receiving one packet and allowing the others to circulate the ring until received will be considered in future work. Tunable transmitters can also be introduced, however tuning latencies can degrade the performance and are more expensive that two fixed transmitters.

Each node has three FIFO (First-In-First-Out) buffers to queue packets arriving from access networks for transmission on the ring. The first buffer is for packets for ring A, the second is for packets for ring B (or ring C depending on the node location), and the third buffer is for packets where ring A or ring B (or C) results in the same distance to the destination node. Upon transmission, each node checks its buffers. If a packet is found, the wavelength associated with the packet ring is sensed. Each buffer is assumed to accommodate up to 1000 packets. A packet is discarded if upon its arrival the buffer supposed to accommodate it is full or if the packet spent more than 1 ms (which is a typical upper bound on the delay for real time traffic [10]) waiting in the buffer. It is expected that higher layers (e.g. IP) will do the recovering task.

As nodes in this architecture can receive packets on any wavelengths, each node is assigned a different subcarrier multiplexed tone. The Destination's subcarrier tone, representing the packet's destination address, is multiplexed by the source node into the packet. At the same time, nodes constantly monitor all wavelengths in parallel. If a node detects its own subcarrier tone, it receives the packet. A physical implementation of such a mechanism was proposed in [11] where sub-carrier multiplexed tones are transmitted over the wavelength along with the data to identify the destination.

Destination stripping is applied, i.e. marking the slot empty after receiving the packet is the responsibility of the destination node. This empty slot can be used by the same node if the node has a packet to send. We also introduced and evaluated a fairness mechanism where the node that marks the slot empty is not able to use the slot immediately. The results indicated that this restriction reduces the bandwidth utilization efficiency of the ring while introducing little change in the throughput difference between nodes. Therefore this restriction is not used here. A modified version of the MAC protocol is introduced in section 3.

3 Mirroring technique for SAN node

A simple technique is used to mirror the primary SAN node (node 0) to a secondary SAN node introduced to the network at the other sectioning point (i.e., node 8). These SAN nodes are located at the sectioning points to make use of the section link and to ensure that the two SAN nodes are separated by the maximum distance to survive disasters in a limited distance scenario. The secondary SAN does not send any traffic to the primary SAN node and its transmission rate is set at 5 Gbit/s. Ordinary nodes do not send to the secondary SAN node. However, the secondary SAN node ultimately receives all the traffic addressed to the primary SAN node as the primary SAN node only marks a slot empty after receiving a packet if the secondary SAN node has already received this packet. Otherwise it will let this packet remains in the ring to go to the secondary SAN node.

From the discussion above it is clear that the main difference mirroring will introduce to the network is that some packets, those which pass by the primary SAN node first, will travel further in the network to be mirrored to the sec-

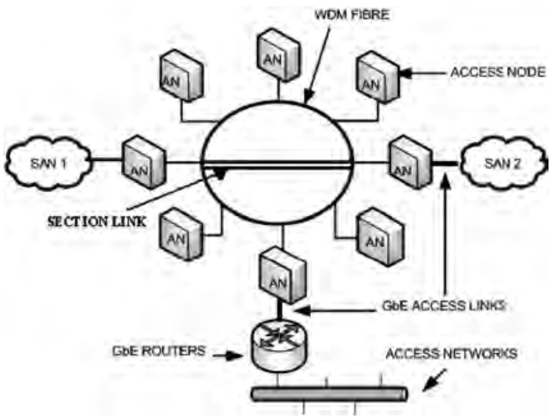

Fig. 1: Network architecture

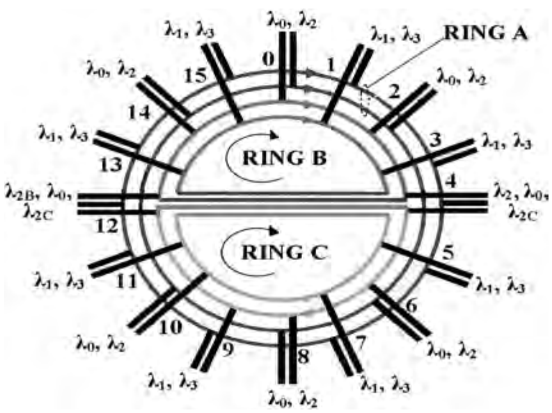

Fig. 2: Logical topology of the architecture

ondary SAN node which means extra bandwidth is used. However, on average this proposed mirroring scheme saves bandwidth and introduces efficiency in that separate transmissions are not needed to synchronise the SAN and its mirror. The Two remain synchronized at all time subject to the ring propagation delay.

According to the original MAC protocol, nodes in the upper part of the network could use either ring B or ring A to send to the primary SAN node. Although both rings result in the same distance to the primary SAN node, ring A results in a longer distance to the secondary SAN node. This extra distance increases mirroring time and leads to inefficient bandwidth usage. To reduce bandwidth usage

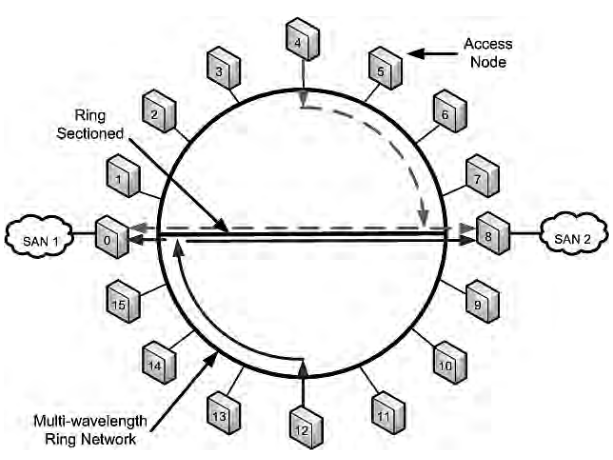

Fig. 3: Mirroring strategy with the modified MAC protocol

and the mirroring time for the upper nodes, a modification could be introduced to the MAC protocol. In this modified version of the protocol, the upper nodes have to use ring B to send to the primary SAN node. To overcome the extra load introduced to ring B, a wavelength could be taken from ring A and assigned to rings B and C. Ring A can accommodate its traffic in a single wavelength as less traffic travels through it (40% to 60% is assumed to be destined to the SAN node in the asymmetric scenario). Figure 3 shows the mirroring strategy under the modified MAC protocol.

4 Performance evaluation

Simulation is carried out to evaluate the performance of the sectioned WDM metro ring network with the proposed mirroring technique. Simulation results show three network parameters including node throughput, queuing delay and transmission buffer packet dropping probability. Two different models of the traffic are used – *Poisson* and *self-similar*. It has been shown in [12] that LAN and WAN traffic is better modeled using statistically self-similar processes. However, Poisson models are still used because they are analytically tractable and can be modeled easily compared to the self-similar processes; see e.g. [13, 14]. The self-similar traffic sources were simulated using aggregated ON–OFF Pareto-distributed sources. The Hurst parameter used for the aggregated traffic is H=0.8 which represents a high degree of burstiness.

The performance of the network is evaluated under varying levels of traffic loads. Results are shown against the normalized network load, denoted by L, ranging between zero and one. A network of 16 nodes is evaluated. In this case, an aggregated data rate of 20 Gbit/s traffic (5 Gbit/s from the SAN node and 15 Gbit/s from other ordinary nodes) is generated by nodes without the mirroring technique and 24 Gbit/s (10 Gbit/s from the SAN nodes and 14 Gbit/s from other ordinary nodes) with the mirroring technique. This maximum traffic represents a normalized load of 1, L=1. It worth mentioning that the total bandwidth capacity of the WDM ring is around 15 Gbit/s (2.5 Gbit/s × 6, where for rough guidance only we assume that the introduction of the section is similar to the introduction of two extra wavelengths) and therefore a normalized load of 1 will create more traffic on the ring than the total carrying capacity of an unslotted WDM ring network, however the slotted regime introduces a further spatial multiplexing gain.

To reflect the effect of traffic asymmetry and mirroring in the performance of the network, average results of nodes in the upper part of the network are shown separately from those of nodes in the lower part of the network. Also results of the primary SAN node are shown separately from those of the secondary SAN node. In the following discussions, we present the performance under different cases. First the performance of the network without mirroring under both 40% and 60% asymmetric traffic is shown. Then the performance is shown under the mirroring technique with the original and the modified MAC protocols for both 40% and 60% asymmetric traffic. All these cases are evaluated under Poisson traf-

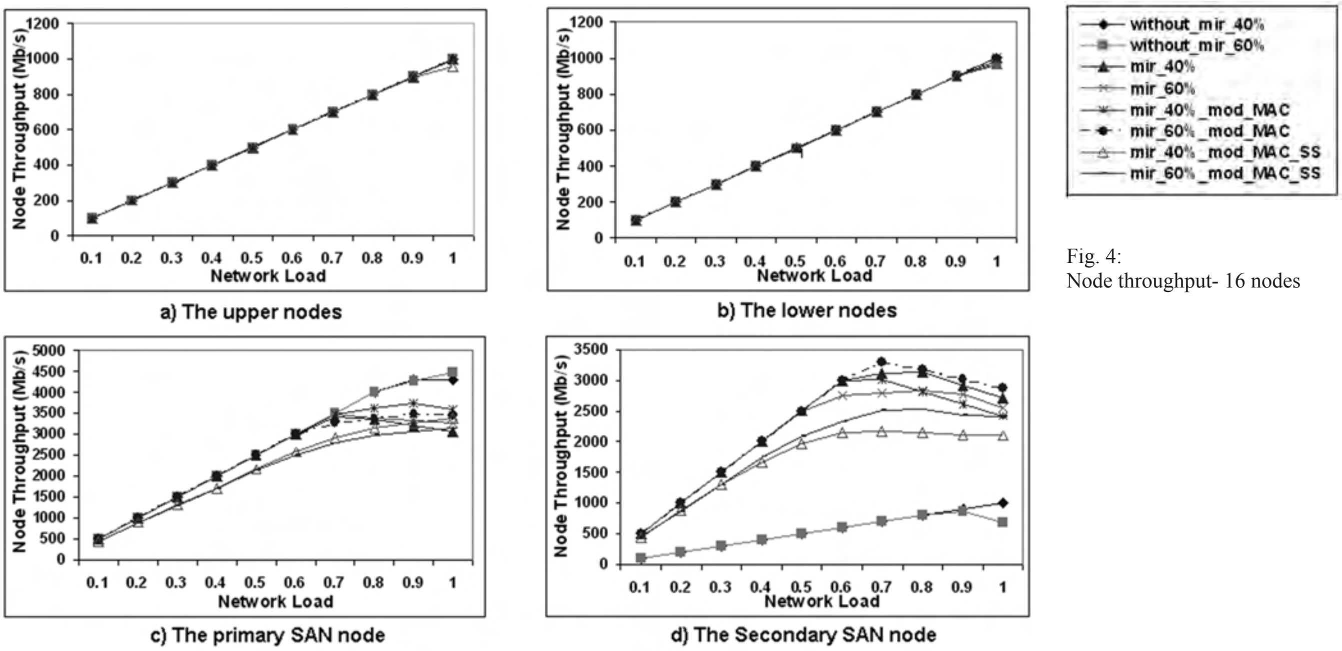

Fig. 4:
Node throughput- 16 nodes

fic. For comparison, the performance of the network with mirroring under the modified MAC protocol is also evaluated under self-similar traffic.

Figure 4-a shows the node throughput results of the upper nodes. They manage to achieve the maximum throughput with and without mirroring for both 40% and 60% asymmetric traffic. However, this is expected as under both the original and modified MAC protocol, the upper nodes are not affected by mirroring as the bandwidth available to them will not be used by any of the mirrored packets. Mirrored packets from the lower nodes will always go through ring C. The figure also shows that the use of self-similar traffic slightly reduces the achievable throughput, under 40% asymmetric traffic when compared with the case where Poisson traffic is used. This is expected due to the burstiness of the traffic produced by self-similar sources.

Without the mirroring strategy, the lower nodes (Fig. 4-b) manage to achieve the maximum throughput for 40% asymmetric traffic while for 60% asymmetric traffic the throughput is slightly less than the maximum. This is understood as higher proportion of traffic sent to the primary SAN node means more load on ring C. It could be seen also from the figure that while under Poisson traffic the modified MAC protocol achieves the maximum throughput, the throughput very slightly decreases under self-similar traffic as with the upper nodes.

For the primary SAN node (Fig. 4-c), it could be seen that mirroring with the original MAC protocol decreases the achieved throughput as half of the packets (those from the upper nodes), which used to be emptied and possibly reused by the primary SAN node, continue their way to the secondary SAN node to be emptied and possibly reused. Therefore the bandwidth available to the primary SAN node decreases and the bandwidth available to the secondary SAN node increases. It is noticed from the figure that without mirroring the original MAC protocol performs better under 60% asymmetric traffic compared to 40% asymmetric traffic as increasing the proportion of traffic going to the primary SAN node means more slots will be emptied and possibly reused by it. It could be seen from the figure that the performance of the primary SAN node has improved under the modified MAC protocol as the extra bandwidth introduced by the added wavelength is greater than the extra traffic added to ring B. However, this extra traffic is more in the case of 60% asymmetric traffic which makes performance under 60% asymmetric traffic worse than 40% asymmetric traffic although in the first case more packets are emptied and possibly reused by the primary SAN node. The figure also shows that the throughput of the modified MAC protocol decreases under self-similar traffic. It is noticed from the figure that for the mirroring cases the network becomes heavily loaded at load of 0.7 and therefore the throughput is almost constant for higher loads.

For the secondary SAN node (Fig. 4-d), without the mirroring technique the throughput achieved under 40% asymmetric traffic is higher than 60% asymmetric traffic as under 60% asymmetric traffic ring C is more loaded. The throughput reaches its maximum at a load of 0.9 and then decreases as the network becomes heavily loaded. As mentioned before, under the mirroring technique the maximum transmission rate of the secondary SAN node increases to 5 Gbit/s. Although under mirroring more packets will be emptied and possibly reused by the secondary SAN under 60% asymmetric traffic, still the high load on ring C makes the performance of the original MAC protocol better under 40% asymmetric traffic. It could be seen from the figure that the secondary SAN node under the original MAC protocol achieved its maximum throughput at a load of 0.6 then the through appears to be almost constant and then decreases as the network gets heavily loaded. It could be seen from the figure that while the modified MAC protocol manages to increase the throughput under 60% asymmetric traffic, the throughput decreases under 40% asymmetric traffic. This is understood as under 40% asymmetric traffic more traffic is going through ring A compared to 60% asym-

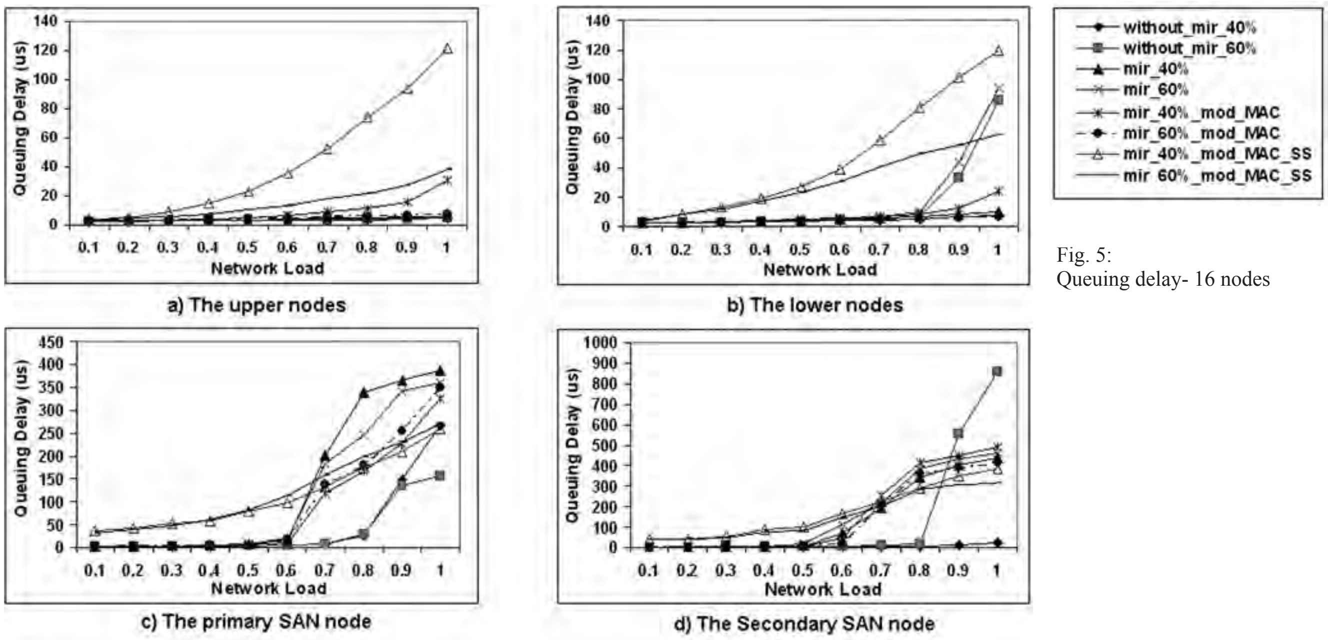

Fig. 5:
Queuing delay- 16 nodes

metric traffic. Therefore taking a wavelength from ring A and giving it to ring C would result in decreasing the throughput. Again under self-similar traffic the throughput is reduced.

Figure 5-a shows the upper nodes average queuing delay. The increase in queuing delay introduced by mirroring is due to the increase in the transmission rate of the secondary SAN node which reduces the bandwidth available to the upper nodes. It could be seen from the figure that under mirroring the original MAC protocol outperforms the modified protocol. This is understood if we remember that under the original MAC protocol, traffic to the primary SAN node could go through either ring A or B, each having 2 wavelengths while under the modified MAC protocol only ring B is available with 3 wavelengths. However, it should be clear that the modified MAC protocol will result in decreasing the propagation time of mirrored packets of the upper nodes by an amount greater than the increase in the queuing delay. It is also noticed that without mirroring and with mirroring under the original MAC protocol, the upper nodes perform better under 60% asymmetric traffic compared to 40% asymmetric traffic as in the case of 60% asymmetric traffic more traffic will have the chance to be sent through either ring A or ring B. Under the modified MAC protocol, 60% asymmetric traffic still outperforms 40% asymmetric traffic but this time because under 40% asymmetric traffic ring A, which has just a single wavelength, is more loaded. However, under the original MAC protocol as expected the difference between 40% and 60% asymmetric traffic is much smaller.

For the lower nodes (Fig. 5-b), mirroring under the original MAC protocol increases the queuing delay slightly compared to the "without mirroring" case as a proportion of the bandwidth available to the lower nodes in ring A will be used by mirrored packets from the upper nodes as these packets, as mentioned before, could go through either ring B or ring A. However, this effect is eliminated under the modified MAC protocol. Also, it is shown in the figure that without mirroring and with mirroring under the original MAC protocol, 40% asymmetric traffic has a lower queuing delay than 60% asymmetric traffic as with the node throughput. It could be seen from the figure that for 60% asymmetric traffic, with and without the mirroring technique, the queuing delay significantly increases as the network gets heavily loaded ($L>0.8$). For the modified MAC protocol, the figure shows that while the queuing delay is reduced under 60% asymmetric traffic, it increases under 40% asymmetric traffic. This could be explained if we remembered that under 40% asymmetric traffic more traffic is going through ring A under 40% asymmetric traffic compared to 60% asymmetric traffic. Therefore taking a wavelength from ring A and giving it to ring C would result in deteriorating the performance.

Similar trends to those of node throughput are observed for the queuing delay of the primary SAN node (Fig. 5-c) under Poisson traffic. While the network gets heavily loaded and the queuing delay significantly increases at loads higher than 0.7 without the mirroring technique, it gets heavily loaded and the queuing delay significantly increases at loads higher than 0.6 with the mirroring technique. It could be seen from the figure that without the mirroring technique under 60% asymmetric traffic and with the mirroring technique under the original MAC protocol that after a certain load the queuing delay appears to be almost constant as the buffers become full. The figure also shows that the use of self-similar traffic gives higher queuing delay than with Poisson traffic when load is below 0.8. For higher loads, the performance under self-similar traffic is better than under Poisson traffic. This is understood from the nature of Poisson traffic whose transmission rate, at high loads, becomes more constant as the packet interarrival duration decreases. Under load of 1, i.e. the maximum transmission capacity of the network is reached; the congestion state is immediately deteriorated by the almost nonvariable, constant arrival of packets and the nodes transmission buffer are constantly filled with packets. Therefore, the queuing delay also in-

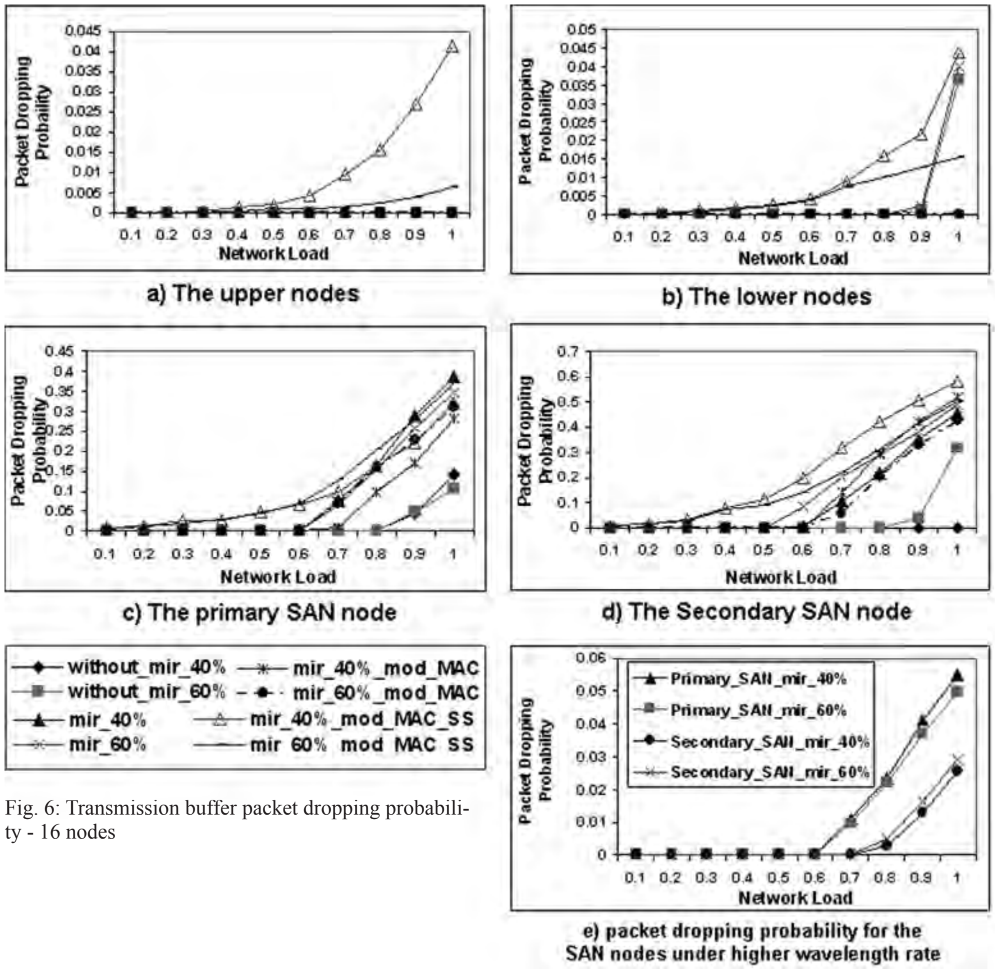

Fig. 6: Transmission buffer packet dropping probability - 16 nodes

creases rapidly when the congestion state is reached. On the other hand, under self-similar traffic, packets bursts tend to frequently increase the buffer loads and the average queuing delay value is seen to be higher than with Poisson traffic when L<0.8. However, as the maximum transmission capacity of the network is reached, the buffers are filled in an intermittent way and the congested state does not induce an abrupt decline in networking performance. Therefore, the queuing delay appears to be almost constant for higher loads.

Figure 5-d shows the queuing delay of the secondary SAN node. Due to the higher transmission rate, the queuing delay increases with mirroring at loads less than 0.9. For higher load the queuing delay for the secondary SAN node without the mirroring technique is higher as ring C becomes highly loaded and fewer packets are emptied by the node compared to the mirroring technique. Applying self-similar traffic results in reducing the queuing delay for loads higher than 0.8 as with the primary SAN node. Also the queuing delay after a certain load appears to be almost constant as with the primary SAN node. Other trends in the figure are similar to those of the node throughput.

Figure 6-a shows the transmission buffer packet dropping probability for the upper nodes. It could be seen from the figure that the upper nodes suffer no packet dropping under Poisson traffic. Under self-similar traffic the upper nodes starts to drop packets from loads as low as 0.4. The dropping probability at full load reaches 0.006 and 0.04 under 40% and 60% asymmetric traffic, respectively.

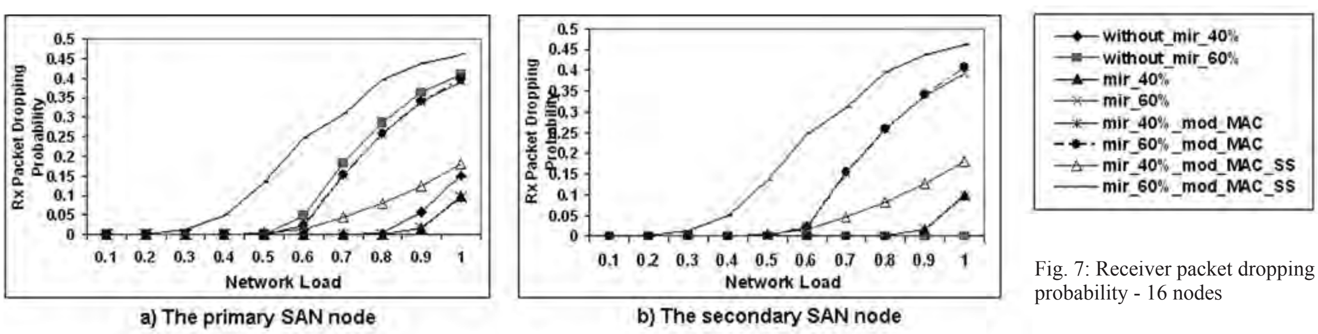

Fig. 7: Receiver packet dropping probability - 16 nodes

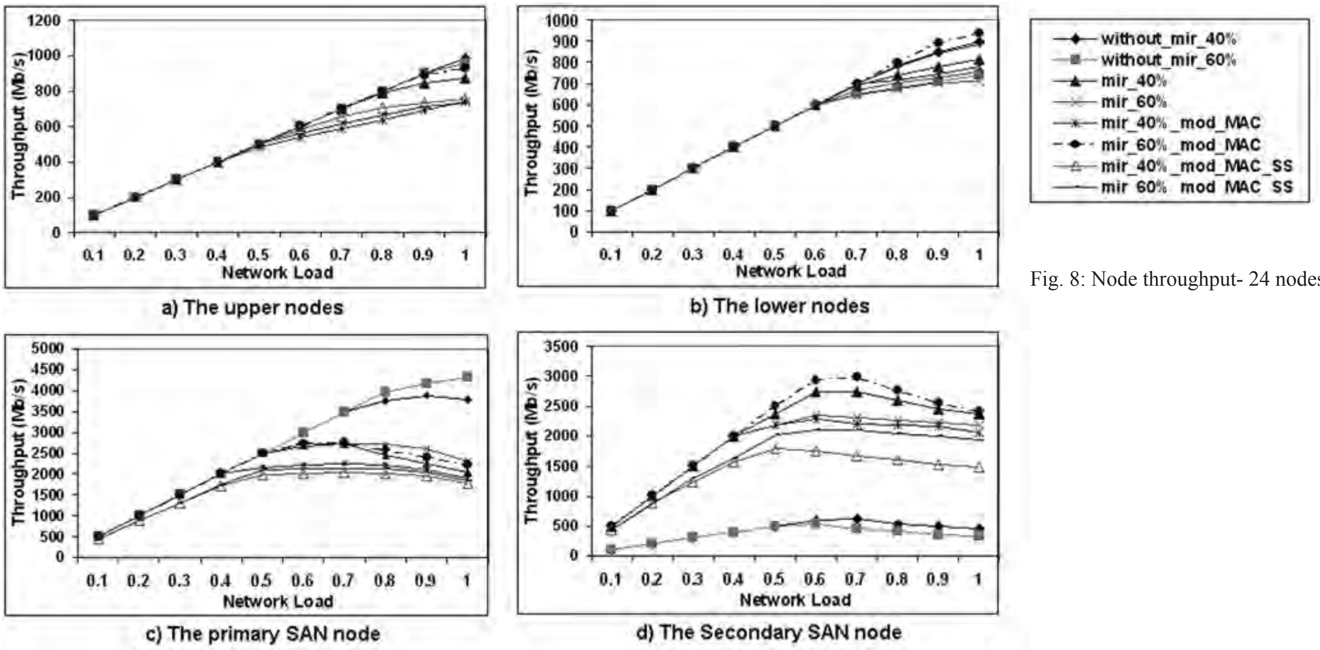

Fig. 8: Node throughput- 24 nodes

Similar trends to those of the queuing delay could be seen for the lower nodes buffer packet dropping probability in Fig. 6-b.

The transmission buffer packet dropping probability of the primary SAN node is shown in Fig. 6-c. Results are shown to be worse under self- similar traffic as the traffic burst quickly fill up the transmission buffers. The same trend is seen in Fig. 6-d for the secondary SAN node buffer packet dropping probability. It is clear from the figures that SAN nodes suffer from a very high packet dropping probability with the mirroring technique.

Figure 6-e shows that the dropping probability can be significantly decreased under the original MAC protocol by increasing the transmission rate of the wavelength on ring C assigned to the SAN nodes to 5 Gbit/s. It is seen in the figure that the dropping probability of the primary SAN node has decreased to about 0.05 and to about 0.03 for the secondary SAN node.

As the two SAN nodes are receiving 40%-60% of the traffic and their receiving rate is limited to 5 Gbit/s, they are expected to drop packets at high loads. Figure 7-a and Fig.7-b show the reception packet dropping probability for the primary and the secondary SAN node respectively. It is clear that the dropping probability under 60% mirroring for the primary SAN node is slightly worse than with mirroring as in the case of mirroring the secondary SAN node does not send to the primary SAN node.

Figure 7 also shows that reception packet dropping probability increases under self-similar traffic. As mentioned before this is due to the burstiness of the traffic produced by self-similar sources.

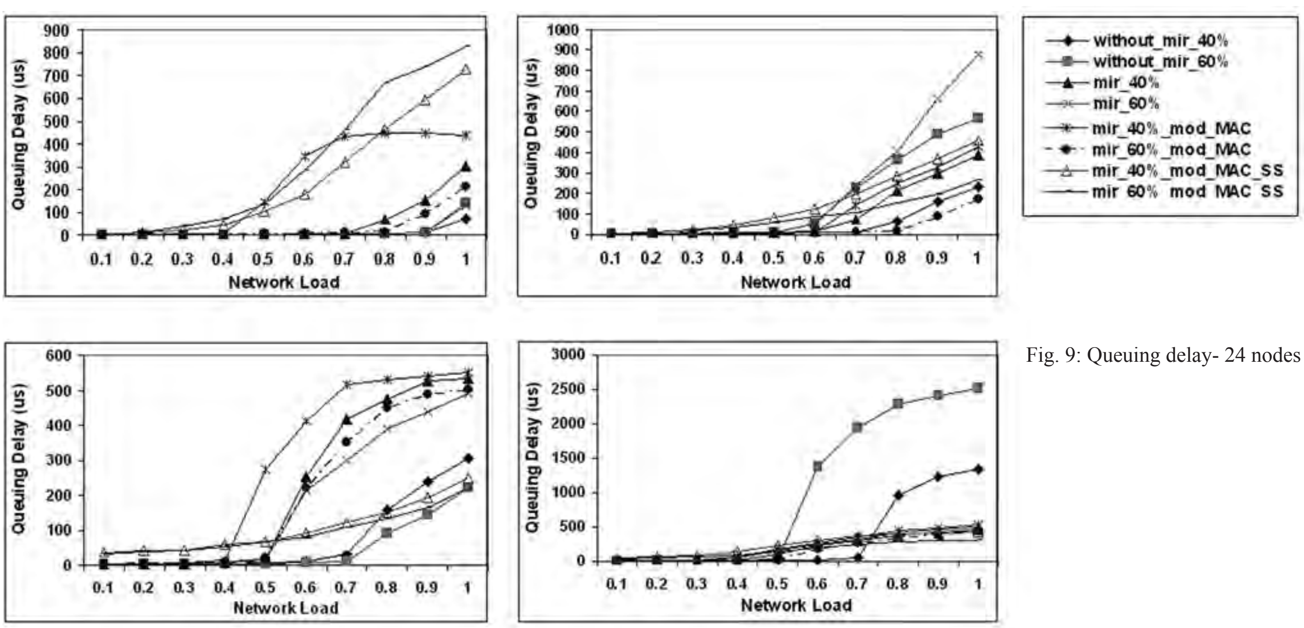

Fig. 9: Queuing delay- 24 nodes

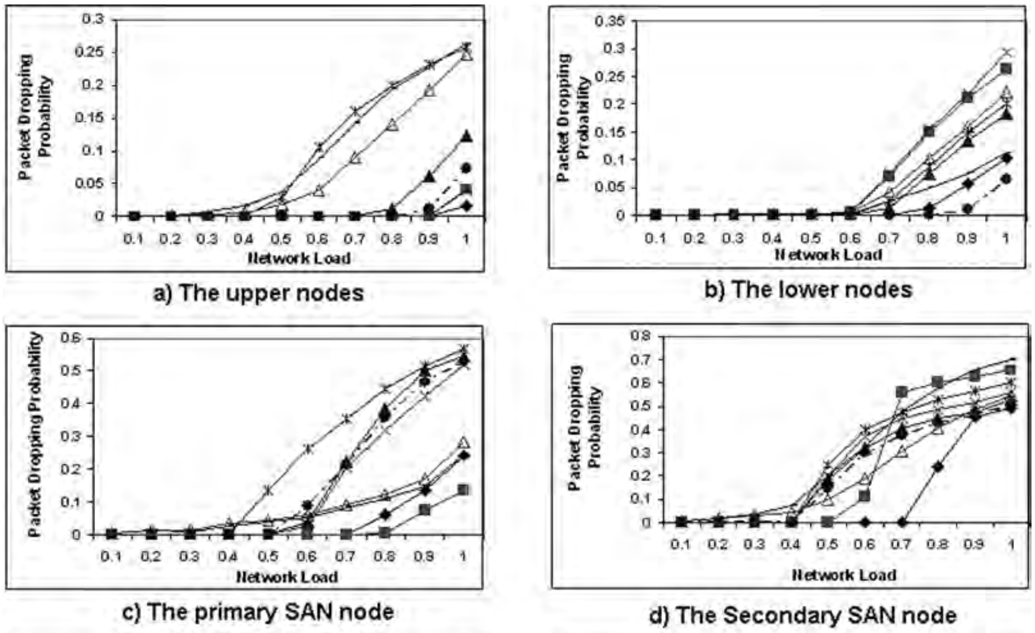

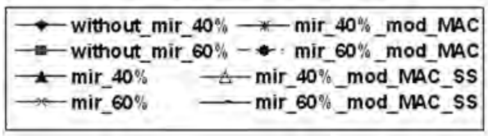

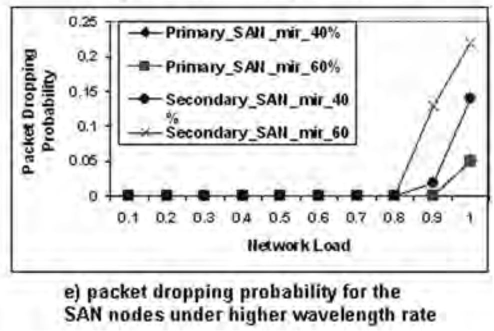

Fig. 10: Transmission buffer packet dropping probability - 24 nodes

Simulation was also carried out for a network of 24 nodes with the same number of wavelengths. In this case, a maximum total of 28 Gbit/s traffic (5 Gbit/s from the SAN node and 23 Gbit/s from other ordinary nodes) is generated by nodes without the mirroring technique and 32 Gbit/s traffic (10 Gbit/s from the SAN node and 22 Gbit/s from other ordinary nodes) with the mirroring technique.

The average throughput of the upper nodes, average throughput of the lower nodes, and throughput of secondary SAN node of a 24 nodes network are given in Fig. 8-a, 8-b and 8-d, respectively. Similar trends to those of a 16 nodes network results could be seen.

However, due to extra traffic, the throughput achieved by nodes is reduced. For the primary SAN node (Fig. 8-c), the modified MAC protocol reduces the achieved throughput as the extra bandwidth introduced by the added wavelength is less than the extra traffic added to ring B which was not the case for the 16 nodes network.

Also the queuing delay and packet dropping probability increase due to the extra traffic. Similar trends to those of a 16 nodes network results could be seen for the upper nodes (Fig. 9-a and Fig. 10-a, respectively), the lower nodes (Fig. 9- b and Fig. 10-b, respectively). The primary SAN node queuing delay and transmission buffer packet dropping probabilities (Fig. 9-c and Fig. 10-c, respectively) increase under the modified MAC protocol. The reason is clear from above discussion. For secondary SAN node, it could be seen that the queuing delay (Fig. 9-d) without mirroring is higher than with mirroring which was not the case for 16 nodes results. However, this is ex-

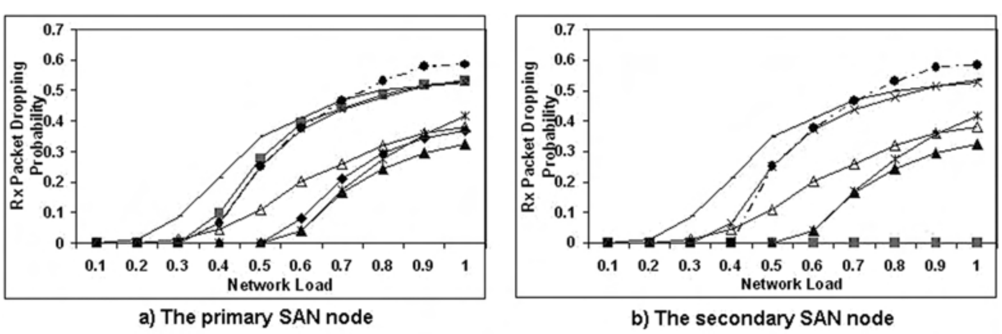

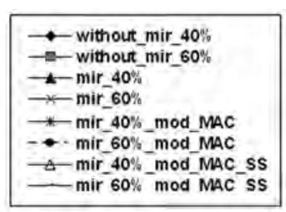

Fig. 11: Receiver packet dropping probability - 24 nodes

pected as the extra bandwidth available to the secondary SAN node with the mirroring is greater than the increase in its transmission rate which was not the case for the 16 nodes network. The same trend could be seen for the packet dropping probability in Fig.10-d It is clear from the figures that all nodes suffer from a very high packet dropping probability with the mirroring technique. As with the 16 nodes network, Fig. 10-e shows that the dropping probability of SAN nodes is significantly decreased under the original MAC protocol by increasing the transmission rate of wavelengths on rings B and C assigned to the SAN nodes to 10 Gbit/s. The packet dropping probability of the upper and lower nodes dropped to zero.

The receiving packet dropping probability of the primary and secondary SAN nodes (Fig.11-a and Fig. 11-b, respectively) shows the same trends as with the 16 nodes network.

5 Conclusion

This paper introduced a novel data mirroring technique for storage area networks (SANs) in a metropolitan area wavelength division multiplexing (WDM) sectioned ring scenario. Performance is evaluated through the simulation of a network with 16 and 24 nodes under different traffic asymmetries of both Poisson and self-similar traffic. Results of average node throughput, queuing delay and packet dropping probability are presented and analyzed. The results showed that applying the proposed technique with the original MAC protocol has different effects on different parts of the ring. While the performance of upper nodes is not affected by mirroring, the performance of the lower nodes and the primary SAN node performance deteriorate. The secondary SAN node has more available bandwidth with mirroring. Also a modified version of the MAC protocol was introduced to decrease mirroring time and improve bandwidth utilization. Simulation results showed that the modified MAC protocol is more effective when a higher proportion of traffic goes to the SAN node. Although the performance of the upper nodes drops under the modified MAC protocol, the lower nodes performance improves significantly which creates better fairness between the upper and the lower nodes.

References

[1] Stanislav Milanovic, Zoran Petrovic: “Building the enterprise-wide storage area network”; EUOROCON 1 (2001), 136 – 139

[2] R. Griswold: “Storage topologies”; Computer 35 (2002) 12, 56 – 63

[3] A. C. Azagury, M. E. Factor, W. F. Micka: “Advanced Functions for Storage Subsystems: Supporting Continuous Availability”; IBM Systems Journal 42 (2003)

[4] Minwen Ji, Alistair Veitch, John Wilkes: “ Seneca: remote mirroring done write”; Proceedings of USENIX Technical Conference, San Antonio, TX (2003), 253–268

[5] T. C. Jepson: “The basics of reliable distributed storage networks”; IT Professional 6 (2004) 3, 18-24

[6] P. Sarkar, et al.: “Internet protocol storage area networks”; IBM Systems Journal 42 (2003) 2, 218-231

[7] X. Qiu, et al..: “Reliability and availability assessment of storage area network extension solutions”; IEEE Commun Mag 43 (2005), 80-85

[8] R. Yan, J. Shu, D. Wen: “An implementation of semi-synchronous remote mirroring system for SANs”; in Proc. of the Workshops on Grid and Cooperative Computing (GCC), Wuhan (2004), 229-237

[9] T. El-Gorashi, B. Pranggono, J. M. H. Elmirghani: “WDM metropolitan sectioned ring for Storage Area Networks extension with symmetrical and asymmetrical traffic”; Proceedings of IEEE International Conference on Communications (ICC), Istanbul (2006)

[10] M. Baldi, Y. Ofek: “End-to-end delay analysis of video conferencing over packet-switched networks”; IEEE/ACM Trans. Networking 8 (2000), 479–492

[11] K. V. Shrikhande, I. M. White, D. Wonglumsom, et al.:”HORNET: a packet-over-WDM multiple access metropolitan area ring network”; IEEE JSAC 18 (2000) 10, 2004-2016

[12] V. Paxson, S. Floyd: “Wide area traffic: The failure of Poisson modeling”; IEEE/ACM Trans. Network 3 (1995) 3

[13] M. A. Marsan, A. Bianco, E. Leonardi, et al.: “All-optical WDM multi-rings with differentiated QoS”; IEEE Communications Magazine 37 (1999) 2, 58-66

[14] J. Cai, A. Fumagalli, I. Chlamtac:”The multitoken interarrival time (MTIT) access protocol for supporting variable size packets over WDM ring network”; IEEE JSAC 18 (2000) 10, 2094-2104