Metro WDM Networks: Comparison of Ring and Star Topologies

Hyo–Sik Yang and Martin Reisslein
Arizona State University
Email: {yangkoon, reisslein}@asu.edu

Martin Herzog, Martin Maier, and Adam Wolisz
Technical University Berlin
Email: {herzog, maier, wolisz}@ft.ee.tu–berlin.de

Abstract— **Both WDM networks with a ring architecture and WDM networks with a star architecture have been extensively studied as solutions to the ever increasing amount of traffic in the metropolitan area. Studies typically focus on either the ring or the star and significant advances have been made in the protocol design and performance optimization for the WDM ring and the WDM star, respectively. However, very little is known about the relative performance comparisons of ring and star networks. In this study we conduct a comprehensive comparison of a state–of–the–art WDM ring network with a state–of–the–art WDM star network. In particular, we compare a time–slotted WDM ring network with tunable–transmitter and fixed–receiver (TT–FR) nodes and an AWG based star network (TT–TR). We evaluate mean aggregate throughput, relative packet loss, and mean delay by means of simulation for Bernoulli and self–similar traffic. Our results quantify the fundamental performance characteristics of ring networks vs. star networks and vice versa, as well as their respective performance limiting bottlenecks, and thus provide guidance for directing future research efforts.**

I. Introduction

Today's metropolitan area networks (MANs) are mostly SONET/SDH ring networks which suffer from a number of drawbacks. Due to their voice–centric TDM operation and symmetric circuit provisioning bursty asymmetric data traffic is supported only very inefficiently. Furthermore, SONET/SDH equipment is quite expensive and significantly decreases the margins in the cost–sensitive metro market. This prevents new companies from entering the metro market. The inefficiencies of SONET/SDH ring networks create a severe bandwidth bottleneck at the metro level. The resultant so–called *metro gap* prevents high–speed clients, e.g., Gigabit Ethernet, from tapping into the vast amounts of bandwidth available in the backbone [1]. In order to (i) bridge this bandwidth abyss between high–speed clients and backbone, (ii) enable new applications benefitting from the huge amounts of bandwidth available in the backbone, and (iii) stimulate revenue growth, more efficient and cost–effective metro architectures and protocols are needed.

Wavelength division multiplexing (WDM) networks have been extensively investigated as solutions to the metro gap. Studies typically focus on either the ring [2] or the star topology [3] and significant advances have been made in the medium access control (MAC) protocol design and performance optimization of the WDM ring and the WDM star, respectively. However, very little is known about the relative performance comparison of ring and star networks. In this paper, we conduct a comprehensive comparison of a state–of–the–art ring and a state–of–the–art star metro WDM network. Our findings reveal the respective strengths and weaknesses of ring networks and star networks. We also identify the bottlenecks that limit the ring/star performance.

II. Related Work

With respect to the node structure, ring WDM networks come in different flavors. Nodes can be equipped with either one fixed–tuned transmitter and an array of fixed tuned receivers (FT–FR$^{\Lambda}$) [2], an array of fixed–tuned transmitters and one fixed–tuned receiver (FT$^{\Lambda}$–FR) [4], or two arrays of fixed–tuned transmitters and receivers (FT$^{\Lambda}$–FR$^{\Lambda}$) [5], where Λ denotes the number of wavelengths in the system. Alternatively, the array of fixed–tuned transceivers can be replaced with one tunable device, e.g., (TT–FR) [6]. For cost and scalability reasons it is generally desirable to deploy a small number of transceivers at each node and decouple the number of wavelengths from the number of nodes. Therefore, we consider a ring network in which each node is equipped with one single transceiver and each wavelength is shared among multiple nodes [7] [8], as described in greater detail in Section III.

Most reported star WDM networks are based on the broadcast–and–select passive star coupler (PSC) [3]. It was shown in [9] that arrayed–waveguide grating (AWG) based single–hop networks clearly outperform their PSC based counterparts in terms of throughput, delay, and packet loss due to spatial wavelength reuse. Therefore, in our comparison we consider a single–hop star WDM network that is based on a wavelength routing AWG, as outlined in Section IV.

The relative performance comparison of WDM networks with different topologies has received very little attention so far. We are only aware of the delay comparison between ring and bus networks [10].

III. Slotted Ring WDM Network

Due to space constraints we review the considered slotted ring WDM network only briefly. For more details on the network architecture and MAC protocol we refer to [7] and [8], respectively, as well as [11].

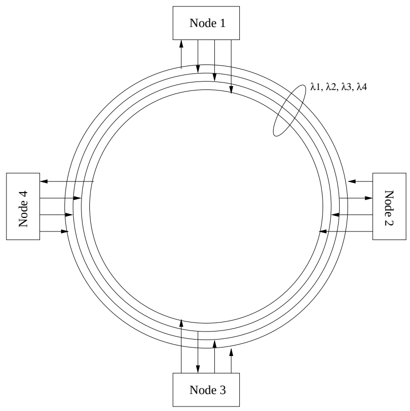

Fig. 1. Architecture of slotted ring WDM network with $N = 4$ nodes and $\Lambda = 4$ wavelengths

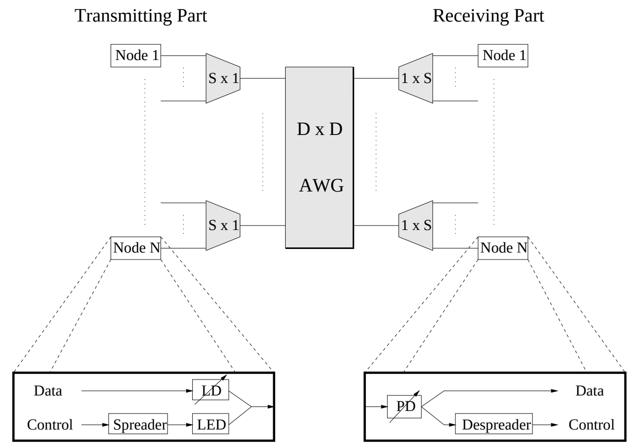

Fig. 2. Architecture of AWG based star WDM network

A. Network Architecture

The ring network interconnects N nodes via a single *unidirectional* fiber whose bandwidth is divided into Λ wavelength channels. Each channel is divided into fixed–length time slots whose boundaries are synchronized across all wavelengths. The slot duration equals the transmission time of a *fixed–size* packet. On all wavelengths each slot consists of a payload field and a *subcarrier multiplexed* (SCM) header. The header on each wavelength contains the slot availability status (empty or occupied) and the destination address of the packet transmitted in the slot. By tapping off some optical power each node is able to monitor all wavelengths simultaneously and detect the presence/absence of packets in every slot.

Each node is equipped with one tunable transmitter and one fixed–tuned receiver (TT–FR). A node can send packets on any wavelength, while it is able to receive packets only on a preassigned *drop wavelength*. For $N = \Lambda$ each node has its own separate *home channel* for reception, as shown in Fig. 1 for $N = \Lambda = 4$. For $N > \Lambda$ each wavelength is equally shared by several nodes for the reception of packets. Specifically, the destination nodes $j = i + n \cdot \Lambda$ with $i \in \{1, 2, \ldots, \Lambda\}$, $n \in \{0, 1, \ldots, \lceil \frac{N}{\Lambda} \rceil - 1\}$, and $j \in \{1, 2, \ldots, N\}$ share the same drop wavelength i. Consequently, nodes sharing the same drop wavelength have to forward packets toward the destination node (*multihopping*). The destination node takes the packet from the ring (destination stripping). With this destination stripping, wavelengths can be *spatially reused* by downstream nodes, leading to an increased network capacity. To avoid head–of–line (HOL) blocking each node deploys $(N - 1)$ *virtual output queues* (VOQs), one for each destination node. Each VOQ holds up to B packets.

B. MAC Protocol

Nodes use the so–called *a posteriori* access strategy, i.e., first a node checks the availability status of each slot by detecting the headers of all wavelength channels and then selects the appropriate VOQ. When an arriving slot is empty the node can use this slot to transmit a packet from the corresponding VOQ. Otherwise, it has to wait until an empty slot arrives. When two or more channels carry an empty slot in the current slot period, *buffer selection* is necessary since each node can transmit only one packet at any given time with its single transmitter. Among various buffer selection schemes we choose the *longest queue (LQ) selection scheme* [8]. With the LQ scheme, the longest VOQ is chosen. When more than one longest VOQ exist, the queue with the lowest index $i \in \{1, 2, \ldots, (N-1)\}$ is chosen. The motivation behind this scheme is load balancing among the queues in the system, thereby increasing the node and network throughput.

IV. AWG Based Star WDM Network

Again, to save space we review the AWG based star WDM network architecture and MAC protocol only briefly. For a detailed description we refer to [11], [12].

A. Network Architecture

As shown in Fig. 2, the network is based on a $D \times D$ AWG. A wavelength–insensitive $S \times 1$ combiner is attached to each AWG input port and a wavelength–insensitive $1 \times S$ splitter is attached to each AWG output port. Thus, the network connects $N = D \cdot S$ nodes. Each node is equipped with a laser diode (LD) and a photodiode (PD) for data transmission and reception, respectively. Both data transmitter and receiver are tunable over Λ wavelengths which are not preassigned to nodes (TT–TR). In addition, each node deploys a broadband light source, say, a light emitting diode (LED), for broadcasting control. The LED signal is *spectrally sliced* by the wavelength–routing AWG such that all receivers obtain the control information. The signaling is done *in–band*, i.e., LED and LD signals overlap spectrally. In order to distinguish data and control information we employ *direct sequence spread spectrum* (DSSS) techniques. Similar to the ring network, each node has $(N-1)$ VOQs, one for each destination. Again, each VOQ holds up to B packets.

B. MAC Protocol

In the considered star WDM network wavelengths are dynamically assigned *on demand* such that any pair of nodes

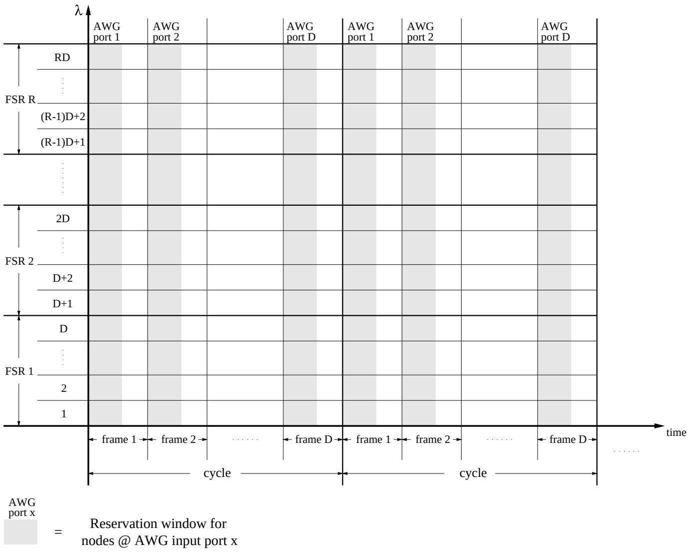

Fig. 3. Timing structure and wavelength assignment

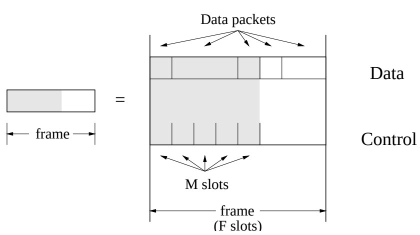

Fig. 4. Frame format

is able to communicate in one *single hop*. As depicted in Fig. 3, the number of available wavelengths Λ span R adjacent free spectral ranges (FSRs) of the underlying AWG, each comprising D contiguous wavelengths, i.e., $\Lambda = R \cdot D$. The applied MAC protocol is an *attempt–and–defer reservation* protocol, i.e., data is only sent after the corresponding control packet has been successful. In the MAC protocol time is divided into cycles, as shown in Fig. 3. Each cycle consists of D frames. Each frame contains F slots, as illustrated in Fig. 4. Each frame is partitioned into the first M, $1 \leq M < F$, slots (shaded region) and the remaining $(F - M)$ slots. In the first M slots, control packets are transmitted and all nodes tune their receivers to one of the corresponding LED slices (channels) in order to obtain the control information. To avoid receiver collisions of control packets, all nodes attached to AWG input port i (via the same combiner) send their control packets in frame i, $1 \leq i \leq D$, of a given cycle. The M slots are not fixed assigned. Instead, control packets are sent on a contention basis using slotted ALOHA. A node is permitted to send up to M control packets in its frame i. Similar to the ring network, each node deploys the LQ scheme, i.e., it first chooses the longest VOQ and randomly selects one of the M reservation slots for sending the corresponding control packet. (Unlike the ring network, however, transmitted data packets do not necessarily have to be of fixed size, as illustrated in Fig. 4. Instead, a control packet attempts to make a reservation for as many fixed–size data packets as possible, thus forming *variable–size* data aggregates with a length of up to F slots.) Next, a control packet is transmitted for the second longest VOQ in a randomly selected slot out of the remaining $(M - 1)$ reservation slots. This procedure is repeated until control packets for all but no more than M non–empty VOQs are sent. Control and data are transmitted simultaneously, but only from nodes at a given AWG input port i. In the first M slots of frame i, the nodes at the other AWG input ports can not transmit. In the last $(F - M)$ slots of each frame no control packets are sent, allowing receivers to be tuned to any arbitrary wavelength. This freedom enables transmissions between any pair of nodes. Thus, data from any AWG input port can be received, allowing for *spatial wavelength reuse*.

After half the end–to–end round–trip time every node (including the sending node) collects all control packets by locking its receiver to one of the LED slices carrying the control information during the first M slots of every frame. Thus, each node maintains *global knowledge* of all the other nodes' activities and also learns whether its own control packets collided in the control packet contention or not. All nodes process the successfully received control packets by executing the same first–come–first–served and first–fit scheduling algorithm, which we adopt since scheduling in very–high–speed optical networks must have low complexity. The scheduling algorithm tries to schedule the variable–size data packets within the scheduling window of pre–specified length. If the scheduling fails, the source node retransmits the control packet in the next cycle, provided the corresponding VOQ is still among the M longest VOQs.

V. Numerical Results

In this section we compare the mean aggregate throughput (mean number of transmitting nodes), relative packet loss, and mean packet delay (time between creation and complete reception of a packet) of the ring and the star network by means of discrete event simulation. Each simulation was run for 10^6 time slots (including 10% warm–up), where a time slot is equal to the transmission time of one fixed–size data packet. The resulting 95% confidence intervals are too small to show up in the plots. The parameters used in both networks are summarized in Table I. In the star, each cycle contains 1024 control slots that cover the first halves of the frames. The duration of a cycle minus a half frame is set equal to the propagation delay, so that the control packets arrive just before the next cycle begins. This results in a frame size of 64 (28) data packets for $D = 2$ ($D = 4$).

Bernoulli and self–similar traffic are considered. In both cases, the arrival rate at each node equals $\sigma \in [0, 1]$, given in packet/slot. More precisely, in each time slot with probability $\sigma/(N - 1)$ a new packet is placed independently into each VOQ (or dropped if the VOQ is full). Self–similar traffic with Hurst parameter 0.75 is generated from ON/OFF–processes with Pareto distributed on–duration and geometrically distributed off–duration. Note that for both types of traffic the total amount of offered traffic is equal to $N \cdot \sigma$.

TABLE I

NETWORK PARAMETERS: DEFAULT VALUES

Description	Symbol	Default Value
Network Diameter	Δ	91.67 km
Number of Wavelengths	Λ	8
Link Speed	C	2.5 Gbit/sec
Propagation Speed	c	$2 \cdot 10^8$ m/sec
Packet Size	L	1500 Bytes
VOQ Size	B	64 Packets

A. Uniform (Balanced) Traffic Scenario

First, we consider uniform traffic where a given source node sends a generated packet to any one of the remaining $(N-1)$ destination nodes with equal probability $1/(N-1)$.

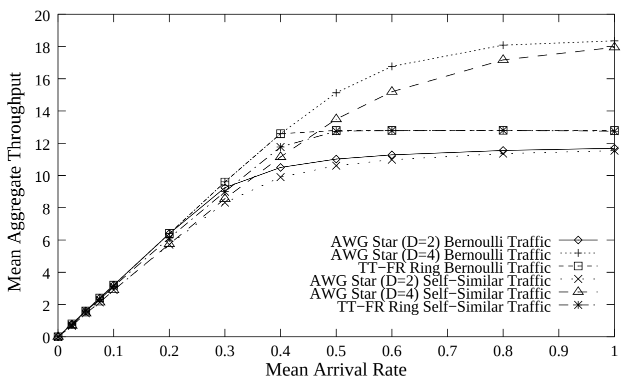

Fig. 5. Mean aggregate throughput for uniform traffic with $N = 32$ (fixed)

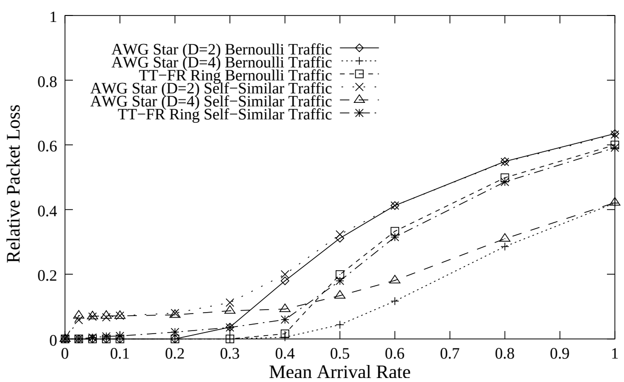

Fig. 6. Relative packet loss for uniform traffic with $N = 32$ (fixed)

Fig. 5 depicts the mean aggregate throughput (mean number of transmitting nodes) vs. σ for $N = 32$. Note that in both ring and star networks the maximum throughput is larger than eight due to spatial wavelength reuse. However, as opposed to the ring network the degree of spatial wavelength reuse can be controlled in the star network by varying the AWG degree D. We observe that for $D = 4$ the AWG based star network clearly outperforms the ring network in terms of throughput. We observe that in both networks for self–similar traffic the throughput is smaller than for Bernoulli traffic. Due to the bursty nature of self–similar traffic, more packets are dropped, resulting in a larger relative packet loss, as depicted in Fig. 6, and a decreased throughput. Moreover, fewer packets are buffered and contribute to the mean delay resulting in a smaller mean delay, as shown in Fig. 7. In terms of delay the ring generally performs better than the star. This is due to the pretransmission coordination overhead of the reservation MAC protocol applied in the star.

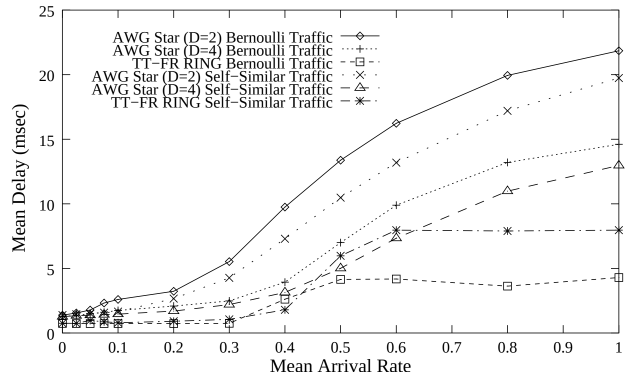

Fig. 7. Mean delay for uniform traffic with $N = 32$ (fixed)

B. Non–uniform (Unbalanced) Traffic Scenario

In the following, we focus on self–similar traffic and investigate uniform traffic in conjunction with client–server traffic. Specifically, we assume to have one hot spot (either server or point of presence (POP)), while the remaining $(N-1)$ nodes act as identical clients. A client sends a fraction h of the traffic to the hot spot, while the remaining fraction $(1-h)$ of the traffic is equally distributed among the other $(N-2)$ clients. Note that $h = 1/(N-1)$ corresponds to uniform traffic only as discussed above. We assume that the server generates as much traffic as all $(N-1)$ clients together and set $\sigma = 0.4$. Fig. 8 illustrates that for increasing h the throughput decreases in both networks due to hot spot congestion, while increased delays occur only in the star network, as shown in Fig. 9.

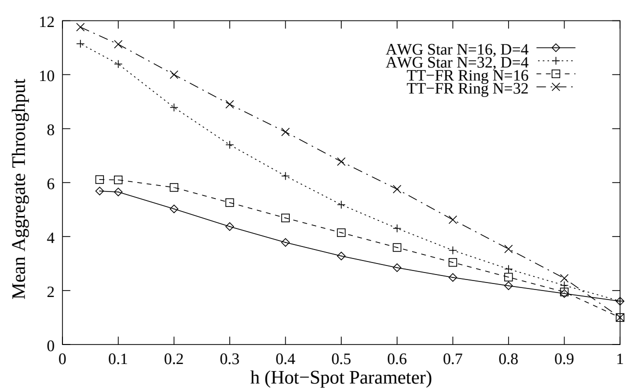

Fig. 8. Mean aggregate throughput as a function of the fraction of hot–spot traffic h with $N \in \{16, 32\}$ and $\sigma = 0.4$, fixed

Next, we fix $h = 0.3$ and $N = 32$ and study fairness among the individual source–destination node pairs, with node 1 functioning as server. As shown in Fig. 10, the star network provides throughput fairness among all clients due to the

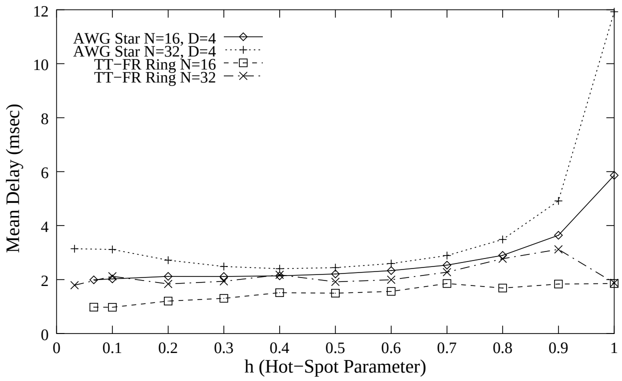

Fig. 9. Mean delay as a function of the fraction of hot–spot traffic h with $N \in \{16, 32\}$ and $\sigma = 0.4$, fixed

round–robin reservation windows and the applied random access of reservation slots and first–come–first–serve scheduling policy. The server achieves a larger mean throughput which is desirable since it has much more data to send than the clients. In contrast, Fig. 11 re–confirms the fairness problems in ring networks.

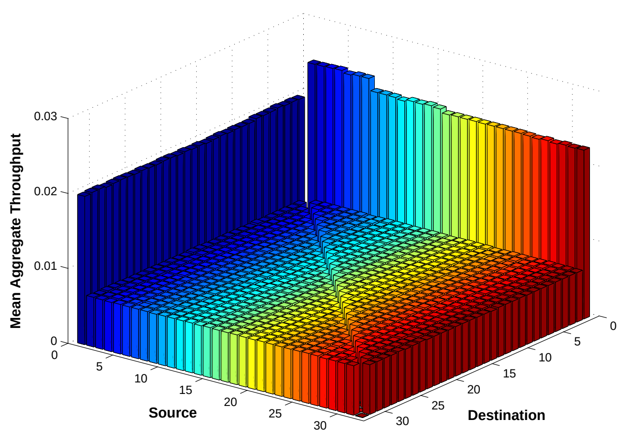

Fig. 10. Pairwise mean aggregate throughput of AWG star network for self–similar hot–spot traffic scenario, $\sigma = 0.4$, $h = 0.3$, and $N = 32$

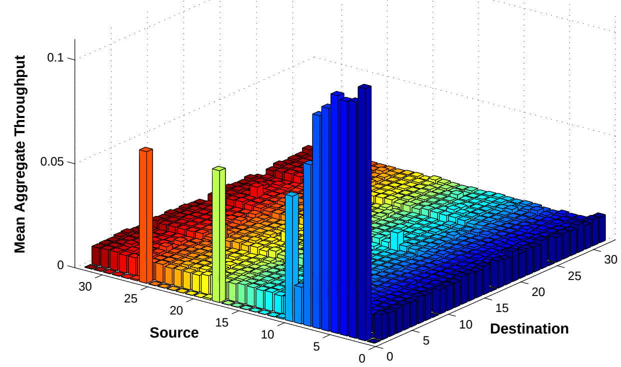

Fig. 11. Pairwise mean aggregate throughput of TT–FR ring network for self–similar hot–spot traffic scenario, $\sigma = 0.4$, $h = 0.3$, and $N = 32$

VI. Conclusion

We have compared the performance of a TT–FR ring and a TT–TR AWG–based star WDM metro network by means of discrete event simulation and have made a number of observations that suggest the following areas for further research and improvement:

- For star networks, efforts should focus on reducing delay (especially when employed in hot–spot settings), while maintaining the high throughput (due to spatial reuse) and fairness. In particular, the pretransmission coordination bottleneck needs to be addressed, see [11] for details.
- For ring networks, research needs to address the throughput limitations and fairness problems, while preserving the small delays. Strategies for increased spatial wavelength reuse appear promising in this regard [11].
- With either network topology, hot–spots (e.g., a POP connecting the MAN to the WAN) require special attention both in the design of network/node architecture and MAC protocol.

References

[1] P.-H. Ho and H. T. Mouftah. A Framework of Scalable Optical Metropolitan Networks for Improving Survivability and Class of Service. *IEEE Network*, 16(4):29–35, July/Aug. 2002.

[2] C. S. Jelger and J. M. H. Elmirghani. Photonic Packet WDM Ring Networks Architecture and Performance. *IEEE Communications Magazine*, 40(11):110–115, November 2002.

[3] B. Mukherjee. WDM–Based Local Lightwave Networks Part I: Single–Hop Systems. *IEEE Network*, 6(3):12–27, May 1992.

[4] J. Fransson, M. Johansson, M. Roughan, L. Andrew, and M. A. Summerfield. Design of a Medium Access Control Protocol for a WDMA/TDMA Photonic Ring Network. In *Proc., IEEE Globecom'98*, volume 1, pages 307–312, November 1998.

[5] J. Cai, A. Fumagalli, and I. Chlamtac. The Multitoken Interarrival Time (MTIT) Access Protocol for Supporting Variable Size Packets Over WDM Ring Network. *IEEE Journal on Selected Areas in Communications*, 18(10):2094–2104, October 2000.

[6] M. A. Marsan, A. Bianco, E. Leonardi, M. Meo, and F. Neri. MAC protocols and fairness control in WDM multirings with tunable transmitters and fixed receivers. *IEEE Journal of Lightwave Technology*, 14(6):1230–1244, June 1996.

[7] K. V. Shrikhande, I. M. White, D.-R. Wonglumsom, S. M. Gemelos, M. S. Rogge, Y. Fukashiro, M. Avenarius, and L. G. Kazovsky. HORNET: A Packet–Over–WDM Multiple Access Metropolitan Area Ring Network. *IEEE Journal on Selected Areas in Communications*, 18(10):2004–2016, October 2000.

[8] K. Bengi and H. R. van As. Efficient QoS Support in a Slotted Multihop WDM Metro Ring. *IEEE Journal on Selected Areas in Communications*, 20(1):216–227, January 2002.

[9] M. Maier and A. Wolisz. Demonstrating the Potential of Arrayed–Waveguide Grating Based Single–Hop WDM Networks. *Optical Networks Magazine*, 2(5):75–85, September 2001.

[10] I. Rubin and J. Ling. Delay analysis of all–optical packet–switching ring and bus communications networks. In *Proc. of IEEE Globecom 2001*, volume 3, pages 1585–1589, November 2001.

[11] H.-S. Yang, M. Reisslein, M. Herzog, M. Maier, and A. Wolisz. Metro WDM Networks: Comparison of Ring and Star Topologies. Technical report, Arizona State University, Telecommunications Research Center, available at `http://www.eas.asu.edu/~mre`, March 2003.

[12] M. Maier, M. Reisslein, and A. Wolisz. A Hybrid MAC Protocol for a Metro WDM Network Using Multiple Free Spectral Ranges of an Arrayed-Waveguide Grating. *Computer Networks*, 41(4):407–433, March 2003.