The AWG||PSC Network: A Performance-Enhanced Single-Hop WDM Network With Heterogeneous Protection

Chun Fan, Martin Maier, and Martin Reisslein, *Senior Member, IEEE*

***Abstract*—Single-hop wavelength-division-multiplexing (WDM) networks based on a central passive star coupler (PSC) or arrayed-waveguide grating (AWG) hub have received a great deal of attention as promising solutions for the quickly increasing traffic in metropolitan and local area networks. These single-hop networks suffer from a single point of failure: if the central hub fails, then all network connectivity is lost. To address this single point of failure in an efficient manner, we propose a novel single-hop WDM network, the AWG||PSC network. The AWG||PSC network consists of an AWG *in parallel* with a PSC. The AWG and PSC provide *heterogeneous protection* for each other; the AWG||PSC network remains functional when either the AWG or the PSC fails. If both AWG and PSC are functional, the AWG||PSC network uniquely combines the respective strengths of the two devices. By means of analysis and verifying simulations we find that the throughput of the AWG||PSC network is significantly larger than the total throughput obtained by combining the throughput of a stand-alone AWG network with the throughput of a stand-alone PSC network. We also find that the AWG||PSC network gives over a wide operating range a better throughput-delay performance than a network consisting of either two load sharing PSCs in parallel or two load sharing AWGs in parallel.**

***Index Terms*—Arrayed-waveguide grating (AWG), medium access control, passive star coupler, protection, single-hop networks, throughput-delay performance, wavelength-division multiplexing (WDM).**

I. Introduction

Single-hop wavelength-division-multiplexing (WDM) networks have attracted a great deal of attention due to theirminimum hop distance, high bandwidth efficiency (no bandwidth is wasted due to packet forwarding as opposed to their multihop counterparts), and inherent transparency. Single-hop networks come in two flavors: *broadcast* networks and *switched* networks. In the 1990s much research was focused on the design and evaluation of media access control (MAC) protocols for single-hop WDM networks that are based on a passive star coupler (PSC), see, for instance, [1]. These networks form broadcast networks in which each wavelength is distributed to all destination nodes. Recently, arrayed-waveguide grating (AWG)-based single-hop networks have attracted much interest [2]–[8]. By using a wavelength-routing AWG instead of a PSC as central hub each wavelength is not broadcast but routed to a different AWG output port resulting in *switched* single-hop networks. These switched single-hop networks allow each wavelength to be used at all AWG input ports simultaneously without resulting in channel collisions at the AWG output ports. The resulting spatial wavelength reuse dramatically improves the throughput-delay performance of single-hop networks [9], [10].

Given the ever-increasing traffic amount due to higher line rates, larger wavelength counts, and spatial wavelength reuse, protection becomes paramount. Specifically, single-hop network operation is immune from node failures since nodes do not have to forward traffic. But all single-hop networks—either PSC or AWG based—suffer from a *single point of failure*: if the central hub fails the network connectivity is entirely lost due to missing alternate paths. Note that this holds also for all multihop networks whose logical topology is embedded on a physical single-hop network. Therefore, protection of (physical) single-hop networks is required to ensure survivability.

Protection of single-hop networks has received only little attention so far [11], [12]. While the passive nature of the PSC and AWG makes the network fairly reliable, it does not eliminate the inherent single point of failure. Clearly, two protection options that come to mind are conventional $1 + 1$ or $1:1$ protection. In these cases, the network would consist of two PSCs or two AWGs in parallel. This type of (homogeneous) protection is rather inefficient: while in the $1 + 1$ protection the backup device is used to carry duplicate data traffic, in the $1:1$ protection the backup device is not used at all during normal operation. To improve network efficiency we propose a novel protection scheme for single-hop WDM networks in this paper. The proposed network consists of one AWG and one PSC in parallel, which we subsequently call the AWG||PSC network. Under normal operation, i.e., both AWG and PSC are functional, the AWG||PSC network uniquely combines the respective strengths of both devices and provides *heterogeneous protection* in case either device fails. The AWG||PSC network enables highly efficient data transport by: 1) spatially reusing all wavelengths at all AWG ports and 2) using those wavelengths *continuously* for data transmission. As discussed shortly, nodes are attached to the central AWG with one tunable transmitter and one tunable receiver. Both

Manuscript received September 30, 2003; revised January 22, 2004. This work was supported in part by the National Science Foundation under Grant Career ANI-0133252. A shorter version of this paper appears in *Proc. IEEE INFOCOM'03*, San Francisco, CA, April 2003.

C. Fan and M. Reisslein are with the Department of Electrical Engineering, Goldwater Center, Arizona State University, Tempe AZ 85287-5706 USA (e-mail: chun.fan@asu.edu; reisslein@asu.edu; web: http://www.fulton.asu.edu/~mre).

M. Maier is with the Centre Tecnològic de Telecomunicacions de Catalunya (CTTC), Barcelona 08034, Spain (e-mail: martin.maier@cttc.es).

Digital Object Identifier 10.1109/JLT.2004.827657

transmitter and receiver are tunable in order to guarantee any-to-any connectivity in one single hop. In such a highly flexible environment, where both transmitter and receiver are tunable, wavelength access is typically controlled by reservation protocols (see the survey [13] and references therein). That is, prior to transmitting a given data packet the source node sends a control packet to inform the corresponding destination node. To do this efficiently, in the proposed network, each node is equipped with an additional transmitter/receiver pair which is attached to the PSC and broadcasts control packets (reservation requests) over the PSC. After one end-to-end propagation delay (i.e., *half* the round-trip time) each node knows the outcome of its reservation and also acquires global knowledge, which is used in a distributed common scheduling algorithm. Besides broadcasting control information, the PSC is used to transport "overflow" data traffic that cannot be accommodated on the AWG.

In this paper, we develop and analyze MAC protocols for the proposed AWG||PSC network. The presented MAC protocols are devised for the three different operating modes: 1) "both AWG and PSC functional" (*AWG-PSC mode*); 2) "PSC failed" (*AWG-only mode*); and 3) "AWG failed" (*PSC-only mode*). We find that the throughput of a stand-alone AWG network plus the throughput of a stand-alone PSC network is significantly smaller than the throughput of the AWG||PSC network in the AWG-PSC mode. Moreover, over a wide operating range, the AWG||PSC network achieves a better throughput-delay performance than a network consisting of either two-load sharing PSCs in parallel or two-load sharing AWGs in parallel.

This paper is organized as follows. In the following subsection, we review related work. In Section II, we briefly describe the properties of the AWG and the PSC. In Section III, we describe the architecture of the AWG||PSC network. In Section IV, we develop MAC protocols for the three operating modes of the AWG||PSC network. In Section V, we develop a probabilistic model of the network and analyze the throughput and delay performance of the three operating modes. In Section VI, we use our analytical results to conduct numerical investigations. We also verify our analytical results with simulations. We summarize our conclusions in Section VII.

A. Related Work

Single-hop networks based on one PSC as the central broadcasting device have been studied extensively since WDM technology was first proposed for optical networks. The studies [1], [14]–[25] represent a sample of the numerous proposals of MAC protocols and analysis of throughput-delay performance associated with various PSC based network architectures. The main constraint of using one PSC is that each wavelength provides only one communication channel between a pair of nodes at any one instance in time. However, wavelengths are precious in metropolitan and local area networks due to cost considerations and tunable transceiver limitations.

One of the ways to increase the transmission efficiency, i.e., to increase capacity without increasing the number of wavelengths, is to reuse the same set of wavelengths in the network. A number of strategies have been examined over the years. Kannan *et al.* [26] introduce a two-level PSC star so that the same set of wavelengths can be reused in each star cluster. Janoska and Todd [27] propose a hierarchical arrangement of linking multiple local optical networks to a remote optical network. Chae *et al.* [28] use an AWG to link multiple PSC networks in series. Again the same set of wavelengths are reused in each star cluster. General design principles for network architectures based on AWG routers for wavelength reuse are studied in [29]–[36]. Bengi [37] studies the scheduling in LAN architectures based on a single AWG or a single PSC.

We introduce the AWG||PSC network to address the single point of failure in single-hop WDM networks. To our knowledge this issue has so far only been considered by Hill *et al.* [11] and Sakai *et al.* [12]. In [11], the central hub of the single-hop WDM network consists of r working AWGs which are protected by n identical standby AWGs. These standby wavelength routers are activated only in case of failure, thus implementing a conventional homogeneous $n : r$ protection scheme. Sakai *et al.* [12] study a dual-star structure where 2 AWGs back up each other in 1 : 1 fashion. Our work differs from [11], [12] in that we propose a heterogeneous protection scheme which efficiently benefits from the respective strengths of AWG and PSC and uses both devices under normal operation.

The operation of our network is different from the parallel processing network described by Arthurs *et al.* [38], which consists of two PSCs. In [38], one PSC is used for data transmission and the other PSC is used for data reception. In case of PSC failure, data transmission or/and reception is impossible due to missing protection. In terms of network architecture, we do not divide the nodes into subnetworks as proposed in [26]–[28]. In the proposed network architecture, all of the nodes are connected directly to the AWG as one network, similar to [3], [6], [10], and [31]. The difference is that all of the nodes are also connected to a PSC, which provides effective broadcast features for control packets. We demonstrate that the broadcast capability of the PSC eliminates the cyclic control packet transmission delays of stand-alone AWG networks, thus achieving high bandwidth efficiency at lower delays.

We note that there is an extensive literature on general techniques for managing optical networks, including techniques for fault detection and recovery, see for instance [39]–[42]. Techniques for fault detection and management in star networks employing AWGs are also provided in [11] and [12]. Essentially the same techniques as in this literature can be employed to detect faults and signal the transition to the appropriate backup operating mode in our AWG||PSC network. More specifically, similar to [11], a protection controller, which is centrally located at the network hub, monitors the functionalities of the central AWG and PSC and initiates the transition to the PSC-only mode or AWG-only mode in case one of these central devices fails. Failures of the transmitters or receivers at a node, or a fiber cut can be detected using the techniques developed in [40]. These components can be integrated in a network management system as in [12]. While these fault detection and network management techniques apply also to our AWG||PSC network, note that the backup operating modes that our AWG||PSC network transitions to in case of a failure are fundamentally different from the backup modes studied in the existing literature.

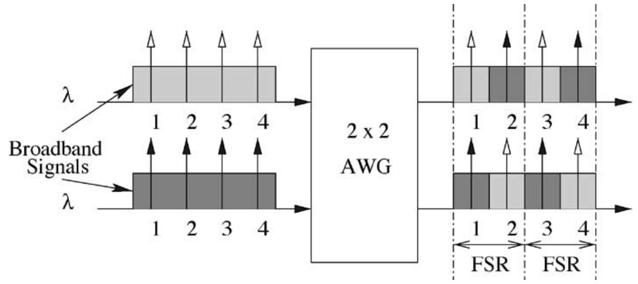

Fig. 1. Periodic wavelength routing of an AWG with $D = 2$ input and output ports when $R = 2$ FSRs are used. Each FSR provides one wavelength channel between an I/O port pair. A total of $D \cdot D \cdot R = D \cdot \Lambda = 2 \cdot 4$ wavelength channels connect the input ports to the output ports.

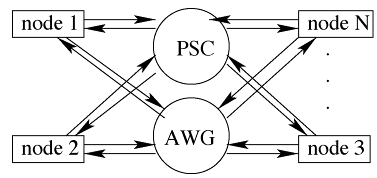

Fig. 2. Network architecture with parallel fiber connections from each of the N nodes to the AWG and to the PSC.

II. Properties of PSC and AWG

The passive star coupler (PSC) is a passive broadcasting device. In an $N \times N$ PSC, a signal coming from any input port is equally divided among the N output ports. The theory and construction of the PSC are detailed in [43], [44]. The broadcast property of the PSC makes it an ideal device for distributing information to all nodes in WDM networks. Star topology networks based on the PSC as the central broadcast device require a lower power budget compared to networks with a linear bus topology or a tree topology. These advantages have led to numerous proposals for PSC-based broadcast-and-select networks, see Section I-A. In these networks, the dynamic wavelength allocation is controlled by a MAC protocol. Chipalkatti *et al.* [15] and Mukherjee [1] provide surveys and network performance comparisons for different categories of MAC protocols.

The drawback of a PSC network is its lack of wavelength efficiency because each wavelength can only be used by one input port at a time. A collision occurs if a wavelength is used by more than one input port at the same time, resulting in a corrupted signal. Since each wavelength provides exactly one channel between a source-destination pair, expanding the transmission capacity of a PSC network requires more wavelengths. Also, broadcasting information to unintended nodes may lead to added processing burden for the nodes.

The arrayed-waveguide grating (AWG) is a passive wavelength-routing device. Dragone *et al.* [45], [46] discuss the construction and the properties of the AWG. Several works [31], [47]–[49] discuss the application of the AWG in multiplexing, demultiplexing, add–drop multiplexing, and routing. In the proposed AWG||PSC network, we use the AWG as a router. The crosstalk performance of AWG routers and the feasibility of AWG routers have been studied extensively (see, for instance, [50], [51]).

The wavelength reuse and periodic routing properties of the AWG are illustrated in Fig. 1. Four wavelengths are simultaneously applied at both input ports of a 2×2 AWG. The AWG routes every second wavelength to the same output port. This period of the wavelength response is referred to as free spectral range (FSR). Fig. 1 shows two FSRs, allowing two simultaneous transmissions between each AWG input/output (I/O) port pair. From Fig. 1, we also see that in order for a signal from one input port to reach all of the output ports at the same time, a multiwavelength or broadband light source is required.

In our network, we exploit two features of the AWG: 1) wavelength reuse and 2) periodic wavelength routing in conjunction with utilizing multiple FSRs. Wavelength reuse allows the same wavelengths to be used simultaneously at all of the AWG input ports. So, with a $D \times D$ AWG (D input ports and D output ports), each wavelength can be reused D times. Periodic wavelength routing and the utilization of multiple FSRs allow each I/O port pair to be connected by multiple wavelengths. We let R denote the number of utilized FSRs. Hence, $\Lambda = D \cdot R$ wavelengths are used at each AWG port.

Here we point out that the number of nodes N in a metropolitan or local area network is typically larger than D. Combiners are used to connect groups of transmitters to the input ports of the AWG and splitters are used to connect groups of receivers to the output ports of the AWG. With a given number of nodes, there is more than one way to construct a network by varying the parameters of the AWG and the combiners/splitters. For example, we can connect 16 nodes to a 4×4 AWG using four 4×1 combiners and four 1×4 splitters. Or, we can connect the 16 nodes using a 2×2 AWG and two 8×1 combiners and two 1×8 splitters. With, say, $\Lambda = 4$ wavelengths, the first case results in one wavelength channel per I/O port pair, i.e., $R = 1$. The second case results in two wavelength channels per I/O port pair, i.e., $R = 2$. In Section VI, we compare the throughput and delay performance of the network for different configurations of R and D.

III. Architecture

Fig. 2 shows the architecture of the proposed AWG||PSC network. In star networks without redundant fiber backup, each node is connected by one pair of fibers, one for the transmission of data, and one for the reception of data. In our network, we deploy one-to-one fiber backup for improved path protection and survivability, that is, each node is connected to the AWG||PSC network by two pairs of fibers.

The PSC and the AWG operate in parallel. The nodal architecture with transmitter and receiver connections to the AWG and to the PSC for the three different modes of operation (which are described in detail in Section IV) is depicted in Fig. 3. Each node is equipped with two fast tunable transmitters (TT), two fast tunable receivers (TR), each with a tuning range of $\Lambda = R \cdot D$ wavelengths, and one off-the-shelf broadband light emitting diode (LED). Due to the extensive spatial wavelength reuse, the tuning range (number of wavelengths) can be rather small. This allows for deploying electrooptic transceivers with negligible tuning times. More specifically, the transmission time of a packet is typically on the order of microseconds (e.g., 4.8 μs for a 1500 byte packet and 2.5 Gb/s link speed), while fast-tunable

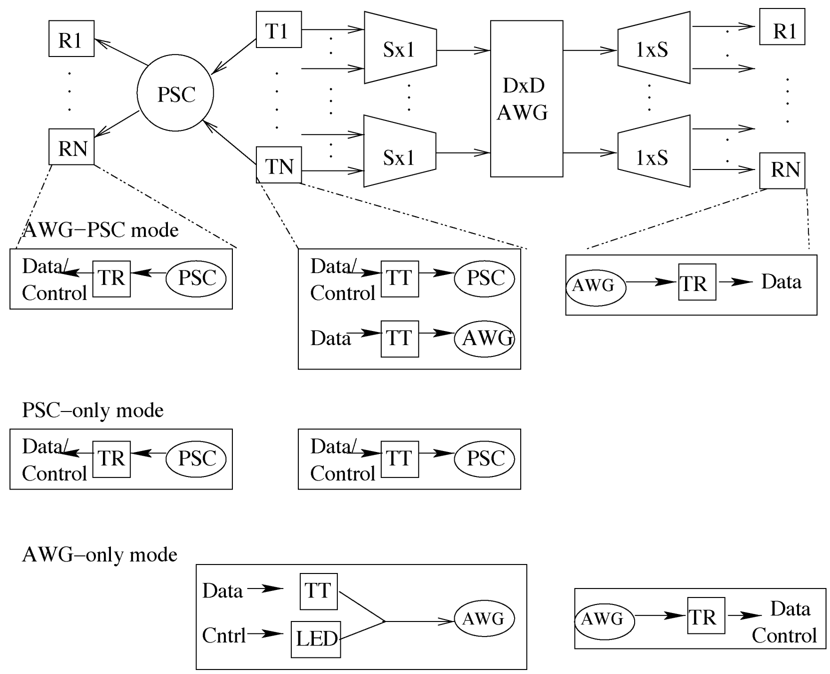

Fig. 3. Detailed node architecture with transmitter and receiver connections to the AWG and to the PSC for the three different modes of operation. A given node's transmitter and receiver are colocated, e.g., T1 and R1 are at the same location. The AWG and PSC are arranged in parallel; the network is "unfolded" here for illustration.

transmitters with a tuning time on the order of nanoseconds have been demonstrated to be feasible in a cost-effective manner [52], [53]. Similarly, electrooptic tunable filters (EOTFs) [54] are promising candidates for realizing fast tunable receivers which are expected not to be significantly more expensive than their fixed-tuned counterparts. We note that fast tunable filters are currently less mature than fast tunable lasers. However, we expect that the development of tunable receivers will attract more attention in the future since using tunable receivers improves the network efficiency and performance by means of load balancing over all wavelengths [55].

One TT and one TR are attached directly to one of the PSC's input ports and output ports, respectively. The TT and TR attached to the PSC are henceforth referred to as *PSC TT* and *PSC TR*, respectively. The second TT and TR are attached to one of the AWG's input ports and output ports via an $S \times 1$ combiner and a $1 \times S$ splitter, respectively. These are referred to as *AWG TT* and *AWG TR*. We note that an alternative architecture to the PSC TT-TR is to equip each node with a tunable PSC transmitter and two fixed-tuned PSC receivers, one tuned to the node's home channel and the other tuned to the control channel. The drawback of this architecture is the lack of data channel flexibility resulting in inefficient channel utilization. In addition, with our approach all wavelength channels can be used for data transmission, whereas with a fixed control channel one wavelength is reserved exclusively for control. Studies in [22], [56] have shown that, by allowing a node to receive data on any free channel, the TT-TR architecture has smaller delays and higher channel utilizations compared to the TT-FR architecture.

The LED is attached to the AWG's input port via the same $S \times 1$ combiner as the AWG TT. The LED is used for broadcast of control packets by means of spectral slicing over the AWG when the network is operating in AWG-only mode (discussed in more detail in Section IV). Two pairs of TTs and TRs allow the nodes to transmit and receive packets over the AWG and the PSC simultaneously. This architecture also enables transceiver backup for improved nodal survivability.

IV. MAC Protocols

We describe MAC protocols for the normal operating mode as well as the various backup modes. We define two levels of backup. The first level is the backup of the central network components, i.e., the PSC or the AWG. Because the AWG and the PSC operate in parallel, the two devices naturally backup each other. We have three different modes of operation: 1) *AWG-PSC mode*, with both AWG and PSC functional; 2) *PSC-only mode*, with AWG down; and 3) *AWG-only mode*, with PSC down. We present the MAC protocols for all three operating modes. The network's throughput and delay performance for each of the three operating modes is examined in Section VI. The second level of backup makes use of the two TT/TRs at each node and the two fiber pairs connecting each node to the central hub to enable transceiver and fiber backup.

Before we describe the MAC protocols for the various modes in detail, we outline the reasoning behind our MAC protocol designs. The nodes in our network employ tunable transmitters and tunable receivers. These tunable transmitters and receivers allow for high flexibility in the data transmissions and receptions and have the potential to achieve load balancing, as well as improved channel utilization and throughput-delay performance. However, with tunable transmitters and receivers, both channel collisions and receiver collisions can occur. A channel collision occurs: 1) when two nodes transmit on the same wavelength over the PSC at the same time or 2) when two nodes that are attached to the same AWG input port transmit

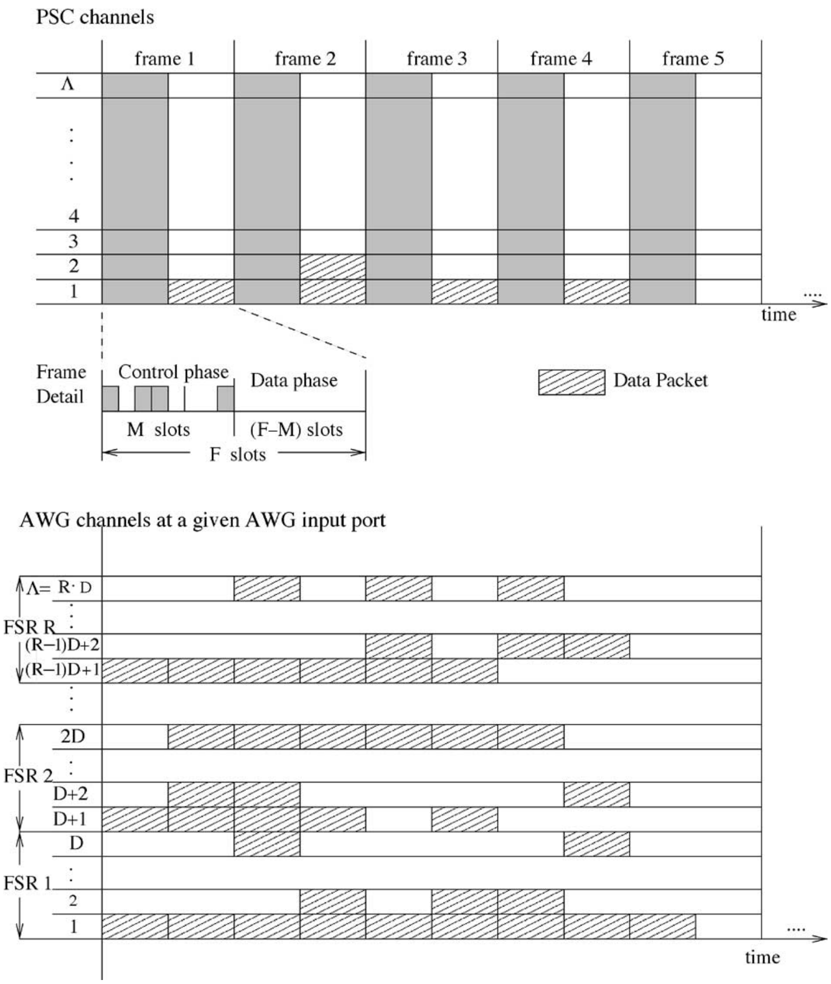

Fig. 4. AWG-PSC mode timing structure: The pretransmission coordination is conducted over a control wavelength channel (one of the Λ wavelength channels) on the PSC during the control phase. During the data phase, data packets are transmitted over the Λ PSC wavelength channels. The Λ wavelength channels at each of the D AWG ports are used for data packet transmission all the time.

on the same wavelength at the same time. A receiver collision (destination conflict) occurs: 1) when two or more signals arrive simultaneously on different wavelength channels over the PSC at a node or 2) when two or more signals arrive simultaneously on different wavelength channels over the AWG at a node. In order to mitigate the collisions, a MAC protocol is typically employed, which arbitrates the access to the wavelengths. Generally, MAC protocols for single-hop WDM networks fall into the three main categories of: 1) preallocation protocols; 2) random access protocols; and 3) reservation (pretransmission coordination) protocols, comprehensively surveyed in [13]. (Since MAC protocols for single-hop networks employing an AWG have received relatively little attention so far, the survey [13] studies the large body of literature on MAC protocols for networks consisting of only a PSC. Some key learned lessons from this literature, however, are considered generally valid for single-hop WDM networks and guide the design of the MAC protocol for our single-hop network consisting of an AWG and a PSC.) Preallocation protocols statically assign a wavelength to a node during a periodically recurring time slot. This static periodic allocation typically gives high utilization for uniform nonbursty traffic, but is poorly suited for the bursty packet traffic. Random access protocols do not require any coordination among nodes. For medium to high traffic loads, however, collisions become very frequent resulting in small throughout and large delay. Reservation protocols employ pretransmission coordination (reservation signalling) to assign wavelengths and receivers on demand. With the so-called attempt-and-defer type of reservation protocol, data packets are only transmitted after a successful reservation. which completely avoids channel and receiver collisions. This approach is generally preferable in a tunable transceiver network carrying bursty packet traffic and we adopt it for our network. For the pretransmission coordination we transmit small control packets according to a random access protocol, namely a modified slotted Aloha protocol. This approach is adopted since (*i*) random access control packet transmission, as opposed to fixed assignments, makes the network scalable and (*ii*) for the typical large propagation delay to control packet transmission delay ratio, slotted Aloha is superior to carrier sensing based access.

A. AWG-PSC Mode

The wavelength assignment and timing structure are shown in Fig. 4. With a transceiver tuning range of Λ wavelengths, the PSC provides a total of Λ wavelength channels. The length of a PSC frame is F slots. The slot length is equal to the transmission time of a control packet (which is discussed shortly). Each PSC frame is divided into a control phase and a data phase. During the control phase, all of the nodes tune their PSC TR to a preassigned wavelength. (One of the wavelength channels on the PSC is used as control channel during the first M slots in a frame; in the remaining slots this channel carries data.)

Given N nodes in the network, if node i, $1 \leq i \leq N$, has to transmit a packet to node j, $i \neq j$, $1 \leq j \leq N$, node i randomly selects one of the M control slots and transmits a control packet in the slot. The slot is selected using a uniform distribution to ensure fairness. Random control slot selection, as opposed to fixed reservation slot assignment, also makes the network upgradable without service disruptions and scalable.

The nodes transmit their data packets only after knowing that the corresponding control packets have been successfully transmitted and the corresponding data packets successfully scheduled. All nodes learn of the result of the control channel transmission after the one-way end-to-end propagation delay (i.e., half the round-trip time). A control packet collision occurs when two or more nodes select the same control slot. A node with a collided control packet enters the backlog state and retransmits the control packet in the following frame with probability p.

The control packet contains three fields: destination address; length of the data packet; and the type of service. Defining the type of service enables circuit-switching. Once a control packet requesting a circuit is successfully scheduled, the node is automatically assigned a control slot in the following frame. This continues until the node releases the circuit and the control slot becomes available for contention.

A wide variety of algorithms can be employed to schedule the data packets (corresponding to successfully transmitted control packets) on the wavelength channels provided by the AWG and the PSC. To avoid a computational bottleneck in the distributed scheduling in the nodes in our very high-speed optical network, the scheduling algorithm must be simple. Therefore, we adopt a first-come-first-served and first-fit scheduling algorithm with a frame timing structure and a short scheduling window. (The scheduling window is the maximum number of frames that packets are allowed to be scheduled into the future.) The frames on the AWG are also F slots long, as the PSC frames. However, unlike the PSC frames, the AWG frames are not subdivided into control and data phase. Instead, the entire AWG frame is used for data. With this algorithm, data packets are assigned wavelength channels starting with the earliest available frame on the lowest FSR on the AWG. Once all the FSRs on the AWG are assigned for that frame, assignment starts on the PSC beginning with the lowest wavelength. Once all the AWG FSRs and PSC wavelengths are assigned in the earliest available frame, assignment starts for the next frame, again beginning with the lowest FSR on the AWG, and so forth. This continues until the scheduling window is full. The unassigned control packets are discarded and the nodes retransmit the control packets with probability p in the next frame. A node with a collided control packet or a data packet that did not get scheduled (even though the corresponding control packet was successfully transmitted) continues to retransmit the control packet, in each PSC frame with probability p, until the control packet is successfully transmitted and the corresponding data packet scheduled.

The nodes avoid receiver collision by tuning their PSC TR to the preassigned control wavelength during the control phase of each frame and executing the same wavelength assignment (scheduling) algorithm. Each node maintains the status of all the receivers in the network. Also, since both the PSC TR and the AWG TR may receive data simultaneously, in the case when two data packets are addressed to the same receiving node in the same frame, the receivers may be scheduled for simultaneous reception of data from both transmitting nodes. In case there are more than two data packets destined to the same receiving node, transmission for the additional packet(s) has to be scheduled for future frame(s).

We note that in our MAC protocol for the AWG-PSC mode, the PSC, which provides Λ wavelength channels, is completely allocated to control for a fraction of M/F of the time. This control makes it possible to use the AWG, which provides $D \cdot \Lambda$ wavelength channels, for data packet transmissions without any channel or receiver collisions all the time.

We note that we consider unicast traffic throughout this paper. However, we do point out that the AWG||PSC network provides a flexible infrastructure for efficient multicasting. A multicast with receivers at only one AWG output port can be efficiently conducted over the AWG, with the splitter distributing the traffic to all attached receivers. A multicast with receivers at several AWG output ports, on the other hand, might be more efficiently conducted over the PSC (to avoid repeated transmissions to the respective AWG output ports).

B. PSC-Only Mode

The network operates in the PSC-only mode when the AWG fails. The failure of the AWG can be detected with a central protection controller [11]. When the AWG fails, the controller signals to the nodes that the network continues operation in the PSC-only mode.

In the PSC-only mode, each frame has a control phase and a data phase as illustrated in Fig. 5. During the control phase, all of the nodes with data packets transmit their control packets in one of the M slots during the control phase. Nodes with collided packets retransmit their control packets following a backoff schedule similar to that of the AWG-PSC mode. The nodes that have successfully transmitted the control packet are assigned the earliest slot starting with the lowest available wavelength. Once the scheduling window is full, the control packets corresponding to unscheduled data packets are discarded and the corresponding nodes retransmit the control packets with probability p in the following frame.

C. AWG-Only Mode

The network operates in the AWG-only mode when the PSC fails. The PSC failure can be detected with a central protection controller [11], which signals to the nodes that the network continues operation in the AWG-only mode when the PSC fails.

Transmitting and receiving control packets over the AWG are more complicated compared to the PSC. First, recall that a multiwavelength or a broadband light source is required to transmit a signal from one input port to all output ports (see Fig. 1). Thus,

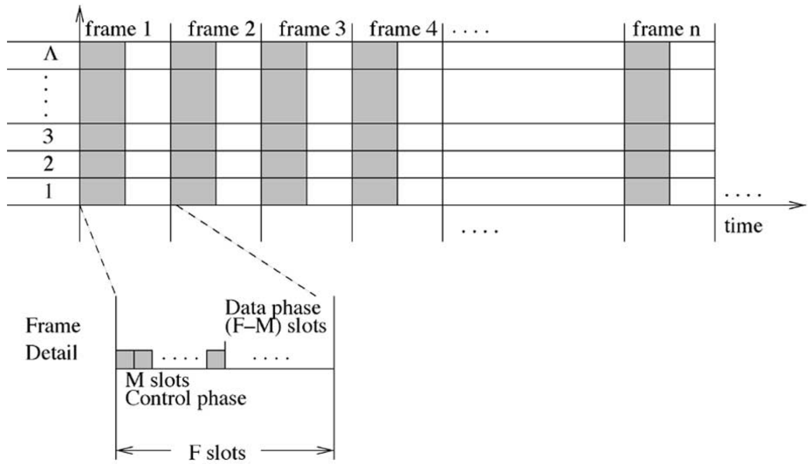

Fig. 5. *PSC-only mode* frame structure.

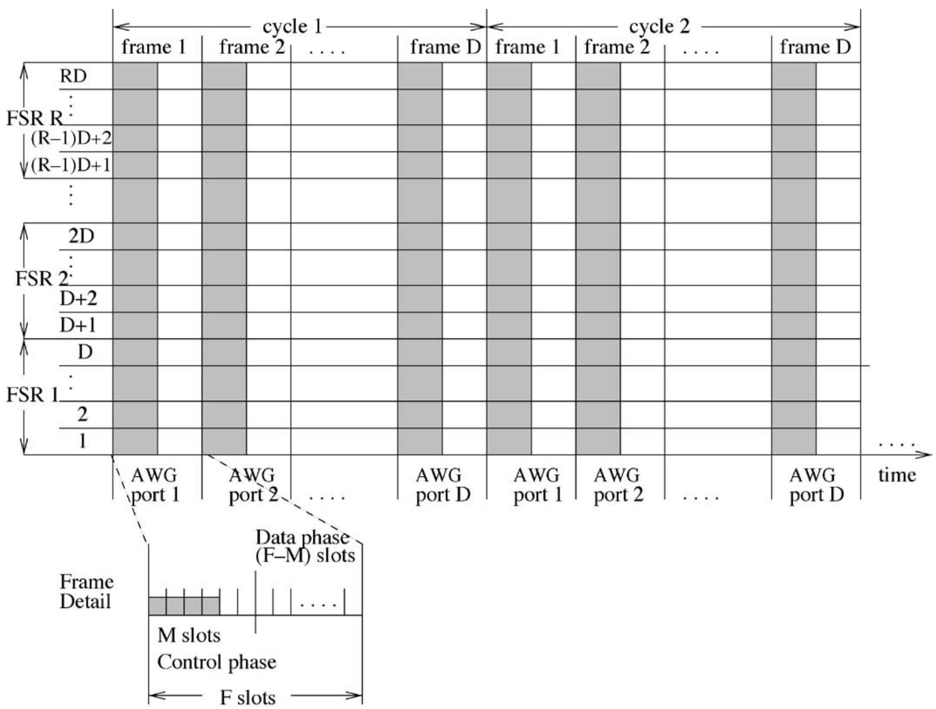

Fig. 6. *AWG-only mode* timing structure.

in the AWG-only mode the LED is used to broadcast the control packets by means of spectral slicing. Second, the transmission of control packets follows a timing structure consisting of cycles to prevent receiver collision of spectral slices. For example (see Fig. 1), if two nodes that are attached to different input ports broadcast control packets using their broadband light source, the wavelength routing property of the AWG slices the signals and sends a slice from each of the broadband signals to each output port. The TR at each node can only pick from one of the wavelengths at each output port to receive the control packet, resulting in receiver collision for the second control packet. Therefore, only the group of nodes attached to the same AWG input port via a common combiner is allowed to transmit control packets in a given frame. In the following frame, the next group of nodes attached to another combiner transmits control packets. This continues until all of the nodes have had a chance to transmit a control packet, and the cycle then starts over. Therefore, with a $D \times D$ AWG, a cycle consists of D frames. The control packet transmission cycle and the frame structure are depicted in Fig. 6. Methods for frame and cycle synchronization are beyond the scope of this paper (see, e.g., [57]–[59, Sec. 7.2.1] for techniques for distributed slot synchronization in WDM networks).

Control packets collide when two or more nodes attached to the same combiner select the same control slot. Nodes with collided control packets retransmit the control packets in the next transmission cycle with probability p.

In the AWG-only mode we distinguish data packet transmission without spatial wavelength reuse and data packet transmission with spatial wavelength reuse. If the scheduling window for data packets is one frame, then nodes can transmit data packets only in one frame out of the D frames in a cycle, which means that there is effectively no wavelength reuse. Full spatial wavelength reuse requires a scheduling window of at least D frames.

D. Transceiver and Fiber Backup

In this section, we describe the second level of backup, the backup of the nodal transceivers and fibers. We note that generally, nodal transceiver and fiber backup in single-hop networks are not as critical as in multihop networks. This is because a transceiver failure or fiber cut in a single-hop network affects only the traffic originating from or destined to the node with the failed transceiver or fiber cut. In a multihop network, on the other hand, a given node has to forward packets that originate from other nodes and are destined to other nodes. Thus, a transceiver failure or fiber cut at one node affects not only the traffic from/to the failed node, but also traffic that originates from other nodes and is destined to other nodes. Nevertheless, nodal transceiver and fiber backup may be important in certain networking scenarios even in single-hop networks and the proposed MAC protocol takes advantage of the node architecture to enable transceiver and fiber backup.

In our network architecture, we denote the fiber connecting the PSC TT of a node to the PSC as the *PSC uplink* and the fiber connecting the PSC TR of the node to the PSC as the *PSC downlink*. We denote the fiber connecting the AWG TT and the LED of a node to the AWG as the *AWG uplink* and the fiber connecting the AWG TR of the node to the AWG as the *AWG downlink*. Note that the failure of a transmitter or receiver at a node has the same effect as a cut of the corresponding fiber, e.g., a failure of the AWG TT has the same effect as a cut of the AWG uplink. We assume that at any time there is at most one failure in the network, i.e., either the AWG or the PSC fails, or one of the nodes experiences a failure, which is reasonable given the long mean time between failures of the network components.

The failure of any of the transmitters or receivers at a node or a fiber cut can be detected with the techniques developed in [40] and is then signalled to the protection controller, which initiates the transition to the appropriate backup mode. More specifically, we define six states, illustrated in Fig. 7, where a node with a failed transmitter or receiver or fiber cut can still communicate. If a node has malfunctions that go beyond the six states, then the node is dropped from the network because the node cannot communicate with other nodes. For example, if a node has a failed PSC TR and a failed AWG TT, then the node cannot transmit control packets over the PSC with its functional PSC TT because the node cannot determine whether the control packets are successful in control packet contention and data packet scheduling (and thus maintain global knowledge in our distributed MAC protocol). Since the AWG TT is down, the node cannot transmit control packets over the AWG and keep track of them with its working AWG TR.

The backup operating modes of our MAC protocol for transceiver and fiber failures are as follows. If a node experiences a failure of its AWG transceiver and/or AWG fibers (i.e., node status 3, 4, or 6) then the network continues operating in the AWG-PSC mode, with some modifications of the scheduling of data packets originating from or destined to the node with the failure. More specifically, if a node has a failure of its AWG TT and/or cut of the AWG uplink fiber (i.e., the node status is 3),

Node State	PSC		AWG	
	TT	TR	TT	TR
0	u	u	u	u
1	d	u	u	u
2	u	d	u	u
3	u	u	d	u
4	u	u	u	d
5	d	d	u	u
6	u	u	d	d

u: up, functional
d: down, non–functional

Fig. 7. Node status based on transceiver and fiber functional status. A transmitter/receiver is considered up if both the transmitter/receiver and the corresponding uplink/downlink are up. A transmitter/receiver is considered down if either the transmitter/receiver or the corresponding uplink/downlink is down (or both are down).

then data packets from the node with the failure are only scheduled on the PSC. If the node experiences status 4, then all data packets to the node are scheduled on the PSC. If the node experiences node status 6, then all data packets to and from the node are scheduled on the PSC.

If a node experiences a failure of its PSC transceiver and/or PSC fibers (i.e., status 1, 2, or 5), then the network transitions to the *AWG-control mode*. In the AWG-control mode, control packets are transmitted over the AWG, similar to the AWG-only mode. Unlike in the AWG-only mode, however, the PSC continues to operate in the AWG-control mode and is used exclusively for data packet transmissions (to and from the nodes that can still transmit and receive over the PSC channels). The data packets from and to the node with the failure are scheduled on the AWG.

We briefly note that if either (*i*) there are two or more nodes that simultaneously experience AWG tranceiver/fiber failure (status 3, 4, or 6), or (*ii*) there are two or more nodes that simultaneously experience PSC tranceiver/fiber failure (status 1, 2, or 5), then our transceiver and fiber backup scheme still works. However if simultaneously one node experiences status 3, 4, or 6, and another node experiences status 1, 2, or 5, then one of the nodes needs to be dropped from the network because two such nodes cannot communicate with one another while maintaining global knowledge of the ongoing control and data packet transmissions in the network. [Only the combination of a node with status 1 and a node with status 3 could be accommodated at the expense of increased overhead by allowing for the simultaneous transmission of control packets over the PSC (from node with status 3) and the AWG (from node with status 1).] Since any malfunction within the network is usually fixed within a short period of time and the mean time between failures is typically large, the likelihood of dropping a node is fairly small.

V. Analysis

In this section, we develop a probabilistic model for the AWG||PSC network. In Table I we summarize the network parameters which have been introduced in the preceding two

TABLE I
NETWORK PARAMETERS AND THEIR DEFAULT VALUES

N	number of nodes in network	200
D	degree (number of ports) of AWG	4
R	number of utilized FSRs	2
Λ	$(= D \cdot R)$, number of wavelengths (transceiver tuning range)	8
p	packet re–transmission probability $(= M/N)$	0.85
F	number of slots per frame	340
M	number of control slots per frame	170
σ	packet generation probability (traffic load)	

sections and will now be used in our model. (The default values given for the network parameters in the table are used in the numerical work in Section VI and may be ignored for now.)

A. System Model

We make the following assumptions in the modeling of the proposed network and MAC protocols.

- *Fixed data packet size*: Data packets have a fixed size of $F/2$ slots. Both the control phase and the data phase on the PSC are $F/2$ slots long, i.e., $M = F - M = F/2$. On the AWG, each frame accommodates two data packets, as illustrated in Fig. 4. With a degree of D and R utilized FSRs (and a corresponding transceiver tuning range of $\Lambda = D \cdot R$), the AWG provides Λ wavelength channels at each of its D ports, for a total of $D^2 \cdot R$ wavelength channels. Thus, the AWG can accommodate at most $2 \cdot D^2 \cdot R$ data packets per frame.
- *Uniform unicast traffic*: A data packet is destined to any one of the N nodes, including the originating node, with equal probability $1/N$. (In our simulations, see Section VI, a node does not transmit to itself. We find that the assumption made in our analytical model that a node transmits to itself with probability $1/N$ gives very accurate results.)
- *Scheduling window*: The scheduling window is generally one frame. (For the AWG-only mode we consider a scheduling window of one frame as well as a scheduling window of one cycle.) In the AWG-PSC mode and the PSC-only mode, a node with collided control packet or with successfully transmitted control packet but no resources (for data packet scheduling) in the current frame retransmits its control packet in the following frame with probability p. In the case of the AWG-only mode, a node with collided control packet or with no transmission resources retransmits in the following cycle with probability p_A.
- *Nodal states and traffic generation*: There are two nodal states: idle and backlogged. A node with no data packet in its buffer is defined as idle and generates a new data packet with probability σ at the beginning of a frame. Let η denote the number of nodes in this idle state. A node is backlogged if it has: *i*) a control packet that has failed in the control packet contention or *ii*) a successful control packet but no transmission resources for scheduling the corresponding data packet. The number of backlogged nodes equals $N - \eta$. Backlogged nodes retransmit their control packets with probability p in a frame. If a node has successfully transmitted a control packet and the corresponding data packet has been successfully scheduled, then the node is considered idle and generates a new packet with probability σ in the following frame.
- *Receiver collision*: We ignore receiver collisions in our analysis to make the probabilistic model tractable. In our simulations in Section VI, on the other hand, we take receiver collisions into consideration. In particular, in the AWG-PSC mode we schedule a data packet on the AWG only if the AWG TR is available. If the AWG TR is busy (or the AWG channels are already occupied), we try to schedule the packet on the PSC. If the PSC TR is busy (or the PSC channels are already occupied), the data packet scheduling fails and the transmitting node retransmits another control packet in the following frame with probability p. In our simulations of the AWG-only mode (PSC-only mode), the data packet scheduling fails if the AWG TR (PSC TR) is busy. Our simulation results in Section VI indicate that the impact of receiver collision on throughput and delay is negligible. This is consistent with [10] which has shown that the effect of receiver collisions is negligible if the number of nodes N is moderately large, which is typical for metro networks.
- *Nonpersistence*: If a control packet fails (in control packet contention or data packet scheduling) we draw a new independent random destination for the corresponding data packet. Our simulations in Section VI do not assume nonpersistence and demonstrate that the nonpersistence assumed in the probabilistic model gives accurate results.

B. Control Packet Contention Analysis

A given control slot contains a successfully transmitted control packet if: *i*) it contains exactly one control packet corresponding to a newly arrived data packet (from one of the idle nodes) and no control packet from the backlogged nodes or *ii*) it contains exactly one control packet from a backlogged node and no control packet corresponding to newly arrived data packets. Let X_i, $i = 1 \ldots M$, denote the number of control packets in slot i. The probability of a given slot containing a successfully transmitted control packet is

$$P(X_i = 1) = \eta \frac{\sigma}{M}\left(1 - \frac{\sigma}{M}\right)^{\eta-1}\left(1 - \frac{p}{M}\right)^{N-\eta} + (N-\eta)\frac{p}{M}\left(1 - \frac{p}{M}\right)^{N-\eta-1}\left(1 - \frac{\sigma}{M}\right)^{\eta} := \kappa \quad (1)$$

where we assume for simplicity that the number of control packets corresponding to newly arrived data packets is independent of the number of control packets corresponding to backlogged data packets, which as our simulations indicate is reasonable.

The expected number of successfully transmitted control packets in each frame is $\sum_{i=1}^{M} P(X_i = 1)$, which has a binomial distribution $BIN(M, \kappa)$. Hence, the expected number of successful control packets per frame is $M \cdot \kappa$.

C. AWG-PSC Mode Data Packet Scheduling

We assume that a data packet from each of the nodes is destined to any other node with equal probability. There are an equal number of nodes attached to each of the combiners and the splitters of a $D \times D$ AWG. Thus, the probability that a control slot contains a successfully transmitted control packet for data transmission between a given I/O port pair is κ/D^2. For notational convenience, let $\rho := \kappa/D^2$.

In the AWG-PSC mode, the throughput of the network is the combined throughput of the AWG and the PSC. Nodes with successfully transmitted control packets are first scheduled using the wavelengths on the AWG. Let Z_A denote the expected throughput on the AWG (in packets per frame). With R FSRs serving each I/O port pair per half-frame, D input ports and D output ports, the expected number of packets transmitted per frame over the AWG is

$$Z_A = D^2 \cdot \sum_{i=1}^{2\cdot R} i \binom{M}{i} \rho^i (1-\rho)^{M-i} + 2 \cdot R \cdot D^2 \cdot \sum_{j=2R+1}^{M} \binom{M}{j} \rho^j (1-\rho)^{M-j}. \quad (2)$$

If all of the FSRs for a given I/O pair are used, then the next packet is scheduled on a PSC channel. Let Z_P denote the expected throughput over the PSC channels (in packets per frame). Let $q_{ij}[n]$ denote the probability that there are $n = 0, 1, \ldots, (M-2R)$, overflow packets from AWG input port i, $i = 1, \ldots, D$, to output port j, $j = 1, \ldots, D$. Recall that the control packets are uniformly distributed over the I/O port pairs. Thus, the overflows from all of the I/O port pairs have the same distribution. So we can drop the subscript ij. If the number of packets destined from an input port to an output port is R or less, then there is no overflow to the PSC. If the number of packets for the given I/O port pair is $R+n$ with $n \geq 1$, then there are n overflow packets. Hence

$$q[n] = \begin{cases} \sum_{i=0}^{2R} \binom{M}{i} \rho^i (1-\rho)^{M-i} & \text{for } n=0 \\ \binom{M}{n+2R} \rho^{n+2R} (1-\rho)^{M-n-2R} & \text{for } n=1,\ldots,M-2R. \end{cases} \quad (3)$$

Let $Q[m]$, $m = 1, \ldots, (M-2R) \cdot D^2$, denote the probability that there are a total of m overflow packets. To simplify the evaluation of $Q[m]$, we assume that the individual overflows are mutually independent. With this assumption, which as our verifying simulations (see Section VI) indicate gives accurate results, the distribution of the combined arrivals at the PSC $Q[m]$ is obtained by convolving the individual $q_{ij}[n]$'s, i.e.,

$$Q[m] = q_{11}[n] * q_{12}[n] * \cdots * q_{1D}[n] * \cdots * q_{DD}[n]. \quad (4)$$

With $Q[m]$, we obtain the expected PSC throughput as approximately

$$Z_P = \sum_{i=1}^{\Lambda} i \cdot Q[i] + \Lambda \cdot \sum_{j=\Lambda+1}^{(M-2R)\cdot D^2} Q[j]. \quad (5)$$

The combined throughput from both AWG and PSC channels is the sum of Z_A and Z_P. To complete the throughput analysis, we note that in equilibrium the throughput is equal to the expected number of newly generated packets, i.e.,

$$Z_A + Z_P = \sigma \cdot E[\eta]. \quad (6)$$

For solving this equilibrium equation we make the approximation that the number of idle nodes η has only small variations around its expected value $E[\eta]$, i.e., $\eta \approx E[\eta]$, which as our verifying simulations in Section VI indicate gives accurate results. By now substituting (2) and (5) into (6), we obtain

$$D^2 \cdot \sum_{i=1}^{2R} i \binom{M}{i} \left(\frac{\kappa}{D^2}\right)^i \left(1 - \frac{\kappa}{D^2}\right)^{M-i} + 2 \cdot R \cdot D^2 \cdot \sum_{j=2R+1}^{M} \binom{M}{j} \left(\frac{\kappa}{D^2}\right)^j \left(1 - \frac{\kappa}{D^2}\right)^{M-j} + \sum_{i=1}^{\Lambda} i \cdot Q[i] + \Lambda \cdot \sum_{j=\Lambda+1}^{(M-2R)\cdot D^2} Q[j] = \sigma \cdot \eta \quad (7)$$

where κ is given by (1) and $Q[\cdot]$ is given by (4). We solve (7) numerically for η, which can be done efficiently using for instance the bisection method. With the obtained η we calculate κ (and ρ) and then Z_A and Z_P.

D. Delay

The average delay in the AWG||PSC network is defined as the average time (in number of frames) from the generation of the control packet corresponding to a data packet until the transmission of the data packet commences. Since in the AWG-PSC mode the throughput of the network in terms of packets per frame is equal to $Z_A + Z_P$, the number of frames needed to transmit a packet is equal to $1/(Z_A + Z_P)$. Given that there are $N - \eta$ nodes in backlog and assuming that the propagation delay is smaller than the frame length, the average delay in number of frames is

$$\text{Delay} = \frac{N - \eta}{Z_P + Z_A}. \quad (8)$$

Propagation delays larger than one frame are considered in Appendix III.

E. PSC-Only Mode

In the PSC-only mode, the channels are shared by all of the nodes. We consider a scheduling window length of one frame. If a control packet is successfully transmitted, but the corresponding data packet cannot be transmitted due to lack of transmission resources, the node has to retransmit the control packet. The maximum number of packets transmitted per frame is equal to the number of channels Λ. The probability of a control slot containing a successfully transmitted control packet is given in (1). Hence, the expected number of successfully scheduled transmissions per frame Z_{PM} is

$$Z_{PM} = \sum_{i=1}^{\Lambda} i \binom{M}{i} \kappa^i (1-\kappa)^{M-i} + \Lambda \cdot \sum_{j=\Lambda+1}^{M} \binom{M}{j} \kappa^j (1-\kappa)^{M-j} \quad (9)$$

and in equilibrium the throughput is equal to the expected number of new packet arrivals, i.e.,

$$Z_{PM} = \sigma \cdot E[\eta]. \quad (10)$$

Z_{PM}, η, and κ are obtained by simultaneously solving (1), (9), and (10). Analogous to (8), the average delay is $(N - E[\eta])/Z_{PM}$ frames.

F. AWG-Only Mode

In the AWG-only mode, we consider two scenarios. In the first scenario, we set the length of the scheduling window to one frame. Recall that under this condition, there is no spatial wavelength reuse. In the second scenario, we set the length of the scheduling window to D frames, i.e., one cycle. In this scenario, there is full wavelength reuse.

Since transmissions in the AWG-only mode are organized into cycles, we define σ_A as the probability of an idle node having generated a new packet by the beginning of its transmission cycle. Given that an idle node generates a new packet with probability σ at the beginning of a frame, we have $\sigma_A = 1 - (1-\sigma)^D$. Similarly, we define p_A as the probability that a backlogged node retransmits a control packet at the beginning of a cycle, where $p_A = 1 - (1-p)^D$. For a $D \times D$ AWG, N/D nodes are allowed to transmit control packets in a given frame. Thus, the probability of a given control slot containing a successfully transmitted control packet is approximately

$$\kappa_A = \frac{\eta}{D}\left(\frac{\sigma_A}{M}\right)\left(1-\frac{\sigma_A}{M}\right)^{\frac{\eta}{D}-1}\left(1-\frac{p_A}{M}\right)^{\frac{(N-\eta)}{D}} + \frac{N-\eta}{D}\left(\frac{p_A}{M}\right)\left(1-\frac{p_A}{M}\right)^{\frac{(N-\eta)}{D}-1}\left(1-\frac{\sigma_A}{M}\right)^{\frac{\eta}{D}}. \quad (11)$$

The average throughput over the AWG in packets per *frame* is equal to the average number of packets transmitted from one given input port to the D output ports in one *cycle*. We assume that a control packet is destined to an output port with equal probability. The probability of a control slot containing a successfully transmitted control packet destined to a given output port is κ_A/D. The AWG accommodates up to R packets per I/O port pair per frame, since the R utilized FSRs provide R parallel wavelength channels between each I/O port pair. Without wavelength reuse, i.e., with a scheduling window of one frame, the nodes at a given input port can utilize the R wavelength channels that connect the considered input port to a given output port only during the latter half of one frame out of the D frames in a cycle. Hence, the expected number of successfully scheduled packets Z_{AM} per frame is

$$Z_{AM} = D \cdot \sum_{i=1}^{R} i\binom{M}{i}\left(\frac{\kappa_A}{D}\right)^i\left(1-\frac{\kappa_A}{D}\right)^{M-i} + R\cdot D\cdot \sum_{j=R+1}^{M}\binom{M}{j}\left(\frac{\kappa_A}{D}\right)^j\left(1-\frac{\kappa_A}{D}\right)^{M-j}. \quad (12)$$

We solve for η numerically using (11), (12) and the equilibrium condition $Z_{AM} = \sigma_A \cdot E[\eta]/D$. With the obtained η we calculate κ_A and then Z_{AM}.

In the second scenario, i.e., with full wavelength reuse, successful control packets destined for a given output port not scheduled in the current frame are scheduled in the following frame, up to D frames. So the AWG accommodates up to $R \cdot D$ $(= \Lambda)$ packets per I/O port pair per cycle. Hence, with wavelength reuse, the expected number of successfully scheduled packets Z_{RE} per frame is

$$Z_{RE} = D \cdot \sum_{i=1}^{R\cdot D} i\binom{M}{i}\left(\frac{\kappa_A}{D}\right)^i\left(1-\frac{\kappa_A}{D}\right)^{M-i} + R\cdot D^2\cdot \sum_{j=R\cdot D+1}^{M}\binom{M}{j}\left(\frac{\kappa_A}{D}\right)^j\left(1-\frac{\kappa_A}{D}\right)^{M-j}. \quad (13)$$

Z_{RE}, η, and κ_A are obtained by simultaneously solving (11), (13) and the equilibrium condition $Z_{RE} = \sigma_A \cdot E[\eta]/D$. With the obtained η we calculate κ_A and then Z_{RE}.

We note that the maximum number of packets that the AWG can accommodate in the AWG-only mode with full wavelength reuse per frame can be increased from $D \cdot \Lambda$ to $D\cdot\Lambda + \Lambda$ by employing spreading techniques for the control-packet transmissions. With spreading of the control-packet transmissions, the nodes at a given AWG input port can send data packets in parallel with their control packets during the first half of the frame as studied in [6]. We also remark that with an additional LED attached to the PSC, the nodes could send data packets in parallel with (spreaded) control packets over the PSC when the AWG||PSC network runs in the AWG-PSC mode. This would increase the number of packets that the AWG||PSC network can accommodate in the AWG-PSC mode per frame by Λ. In order not to obstruct the key ideas of the AWG||PSC network, we do not consider the spreading of control information in this paper.

In the scenario without wavelength reuse, there are two delay components. The first component is the delay resulting from the control packet contention and the scheduling process. This component equals the number of backlogged nodes divided by the throughput. The second component is the waiting period in the transmission cycle. All of the idle nodes generate a new packet with probability σ at the beginning a frame. But the nodes transmit control packets once every D frames. Hence, the expected waiting period from the generation of a new data packet to the transmission of the corresponding control packet is the mean of a truncated geometric distribution, i.e.,

$$I_{\text{del}} = \frac{\sum_{i=0}^{D}(D-i)\cdot\sigma\cdot(1-\sigma)^i}{1-(1-\sigma)^D}. \quad (14)$$

Combining the two components, the total mean delay (in number of frames) is

$$\text{Delay}_{AM} = \frac{N - E[\eta]}{Z_{AM}} + I_{\text{del}}. \quad (15)$$

In the scenario with wavelength reuse, there are three delay components. The first two components are the same as for the scenario without wavelength reuse. The third delay component occurs in the case when the number of scheduled packet is larger than $D \cdot R$. In this case, the packets scheduled in the future frames experience an average delay of $(Z_{RE} - D\cdot R)^+/(2\cdot D\cdot R)$ frames, where $(Z_{RE} - D\cdot R)^+ = \max(0, Z_{RE} - D\cdot R)$. To see this note that if $Z_{RE} > D \cdot R$, the packets not scheduled in the current frame have to wait an average $(Z_{RE} - D\cdot R)/(2\cdot D\cdot R)$

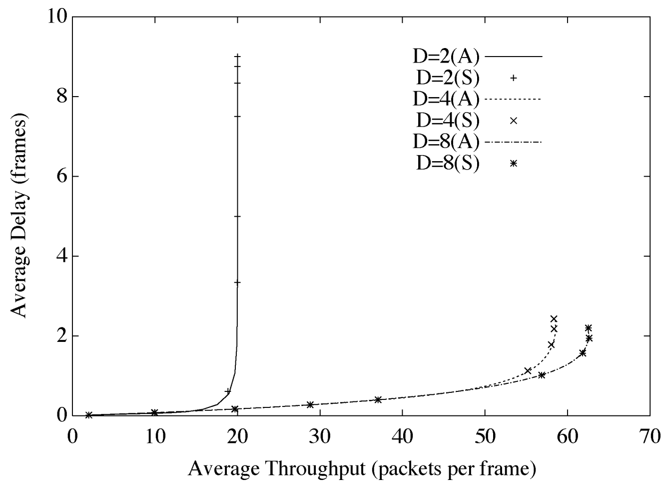

Fig. 8. Throughput-delay performance for AWG-PSC mode for AWG degree D =2, 4, and 8. (R = 2 fixed).

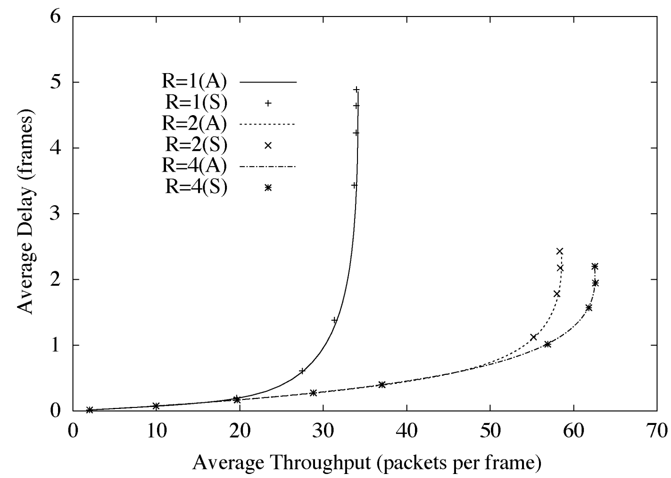

Fig. 9. Throughput-delay performance for AWG-PSC mode for R =1, 2, and 4 used FSRs. (D = 4, fixed).

frames for transmission. Combining the three components, the total mean delay (in frames) is

$$\text{Delay}_{RE} = \frac{N - E[\eta]}{Z_{RE}} + I_{\text{del}} + \frac{(Z_{RE} - D \cdot R)^+}{2 \cdot D \cdot R}. \quad (16)$$

VI. Numerical and Simulation Results

In this section, we examine the throughput-delay performance of the AWG||PSC network in the three operating modes: 1) AWG-PSC mode; 2) PSC-only mode; and 3) AWG-only mode, by varying system parameters around a set of default values, which are summarized in Table I. (We set $p = M/N$ as this setting gives typically a large probability κ of success in the control packet contention. Note from (1) that κ is maximized for $p = (M - \eta\sigma)/(N - \eta - 1)$.) We provide numerical results obtained from our probabilistic analysis (marked (A) in the plots) as well as from simulations of the network (marked with (S) in the plots). Each simulation was run for 10^6 frames including a warm-up phase of 10^5 frames; the 99% confidence intervals thus obtained were always less than 1% of the corresponding sample mean. Throughout the simulations, we used the σ values 0.01, 0.05, 0.10, 0.15, 0.2, 0.4, 0.6, 0.8, and 1.0. We note that in contrast to our probabilistic analysis, our simulations do take receiver collisions into consideration. Also, in the simulations a given node does not transmit to itself. In addition, in the simulations, we do not assume nonpersistence, i.e., the destination of a data packet is not renewed when the corresponding control packet is unsuccessful.

Fig. 8 compares the throughput-delay performance of the network for different AWG degrees D =2, 4, and 8 (with the number of used FSRs fixed at $R = 2$, thus the corresponding Λ values are 4, 8, and 16). For small σ, the throughput-delay performance for the three D values are about the same. For large σ, the throughput for $D = 2$ peaks at 20 packets per frame and the delay shoots up to very large values. A network constructed using $D = 8$ achieves higher throughput at lower delays compared to the $D = 4$ network at high traffic levels. Recall that the wavelength reuse property of the AWG allows each wavelength to be simultaneously used at all of the input ports, thus providing $D \cdot \Lambda$ channels. Furthermore, each AWG FSR at each port accommodates two data-packet transmissions per frame. Thus, the maximum combined throughput of AWG and PSC is $2 \cdot D \cdot \Lambda + \Lambda$ data packets per frame. For $D = 2$, the average throughput reaches a maximum of 20 packets per frame as indicated in the graph. The maximum throughput for $D = 4$ and $D = 8$ are 72 and 272 packets per frame, respectively. For these two cases, the throughput is primarily limited by the number of successful control packets (per frame); whereas the data packet scheduling is the primary bottleneck for $D = 2$.

Fig. 9 compares the throughput-delay performance of the network for different numbers of used FSRs R =1, 2, and 4 (with the AWG degree fixed at $D = 4$, thus, the corresponding Λ values are 4, 8, and 16). The throughput for $R = 1$ peaks at 32 packets per frame and the delay grows to large values, while the throughput and delay for $R = 2$ and $R = 4$ are approximately the same. Increasing R increases the number of channels for each I/O port pair on the AWG, thus, increasing the number of channels in the network. For $R = 1$, the maximum throughput is $2 \cdot D \cdot \Lambda + \Lambda = 36$ packets per frame. The throughput is primarily limited by the scheduling capacity of the network. For $R = 2$ and $R = 4$ the maximum throughputs are 72 and 144 packets per frame, respectively. For these two cases, the throughput is primarily limited by the number of control packets that are successful in the control packet contention. The conclusion is that increasing the number of channels for each I/O port pair does not yield measurable improvements in throughput or delay when there are not enough successful control packets.

In Fig. 10, we fix the number of wavelengths in the network $(\Lambda = 8)$ and examine the throughput-delay performance for different combinations of D and R with $D \cdot R = 8$. We examine the cases: ($D = 2$, $R = 4$), ($D = 4$, $R = 2$), and ($D = 8$, $R = 1$). We observe that ($D = 2$, $R = 4$) has the shortest delay up to a throughput of about 34 packets per frame, and a maximum throughput of 40 packets per frame. The delays for ($D = 4$, $R = 2$) and ($D = 8$, $R = 1$) are approximately the same up to a throughput of approximately 48 data packets per frame. At higher traffic levels, the ($D = 8$, $R = 1$) network achieves higher throughput at lower delays compared to the ($D = 4$, $R = 2$) network due to the larger number of channels in the

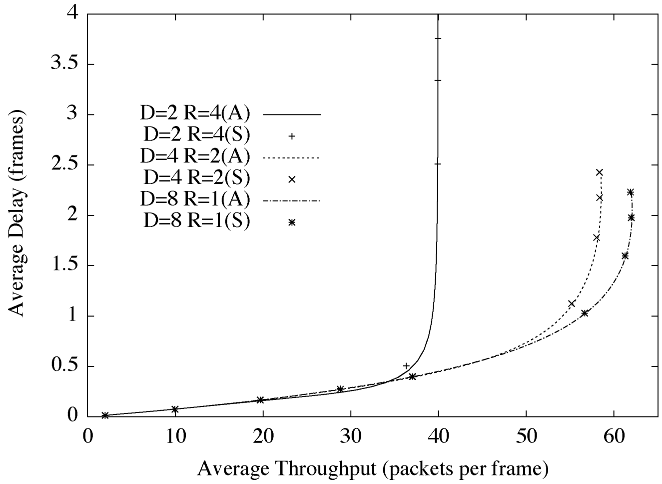

Fig. 10. Throughput-delay performance for AWG-PSC mode for fixed tuning range $\Lambda = R \cdot D = 8$ wavelengths.

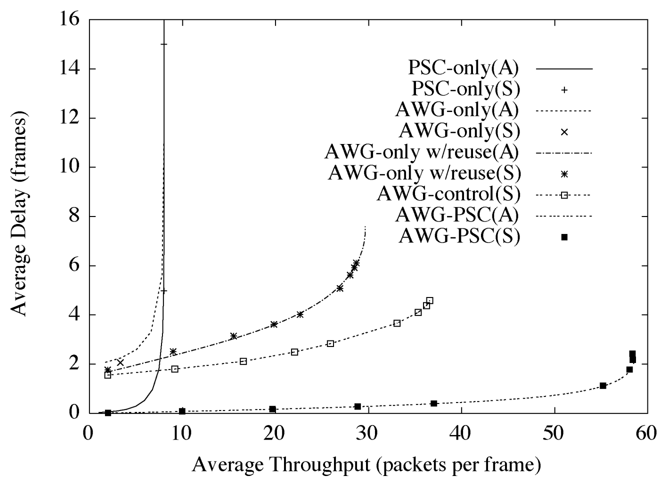

Fig. 11. Throughput-delay performance of AWG||PSC network for different modes of operation.

($D = 8$, $R = 1$) network. The combination ($D = 2$, $R = 4$) achieves the shortest delay at small σ due to higher channel utilization from the larger number of FSRs. The throughput for ($D = 2$, $R = 4$) is bounded by the scheduling capacity of $2 \cdot D \cdot \Lambda + \Lambda = 40$ data packets per frame.

Fig. 11 compares the throughput-delay performance of the network in the following modes: PSC-only mode, AWG-only mode without wavelength reuse (i.e., a scheduling window of one frame), AWG-only mode with wavelength reuse (i.e., a scheduling window of one cycle), AWG-control mode (with wavelength reuse), and AWG-PSC mode. The PSC-only mode has a maximum throughput of 8 data packets per frame. This is expected because the maximum number of channels in a PSC-network is equal to the number of available wavelengths, $\Lambda = 8$. The AWG-only mode with wavelength reuse achieves throughputs up to roughly 30 packets per frame. This is primarily due to the larger number of $D \cdot \Lambda = 32$ available wavelength channels with spatial wavelength reuse. The delay for the AWG-only mode is larger than for both the PSC-only mode and the AWG-PSC mode at low traffic. This is due to the cyclic control packet transmission in the AWG-only mode.

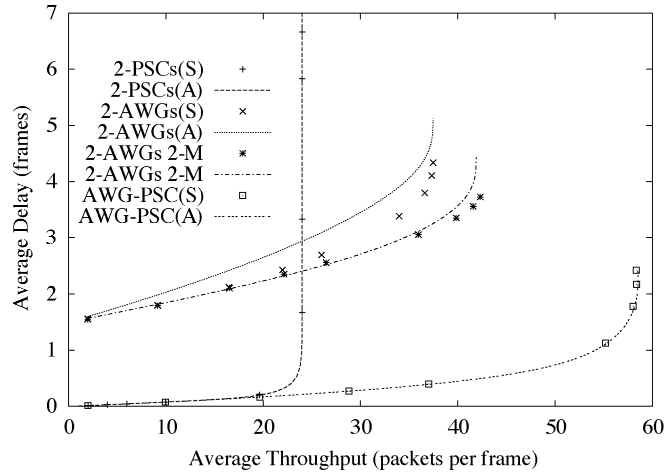

Fig. 12. Throughput-delay performance comparison for three networks: PSC||PSC, AWG||AWG, and AWG||PSC.

The AWG-PSC mode achieves the largest throughput and the smallest delays for all levels of traffic.

We note that the results in Fig. 11 also give an indication of the performance of the AWG||PSC network for node transceiver and fiber backup. If a node's AWG transceiver and/or fiber has failed (status 3, 4, or 6 in Fig. 7), the network continues operating in the AWG-PSC mode with the failed node's traffic transmitted over the PSC, which typically has a minor impact on the overall network performance (especially when the contribution of the failed node to the overall traffic is small). If a node's PSC transceiver and/or fiber has failed (status 1, 2, or 5 in Fig. 7), the network operates in the AWG-control mode, in which control packets are transmitted over the AWG, similar to the AWG-only mode, but which provides Λ additional wavelength channels for data transmission over the PSC. Thus, the AWG-control mode achieves a throughput that is up to Λ packets per frame larger than the throughput in the AWG-only mode, and smaller delays than the AWG-only mode, as we observe from Fig. 11.

Overall, we observe from Fig. 11 that for a given level of delay, the throughput for the AWG||PSC network is significantly larger than the total throughput obtained by combining the throughput of a stand-alone AWG network with the throughput of a stand-alone PSC network. The AWG||PSC network in the AWG-PSC mode has a maximum throughput of 59 packets per frame and a delay of no more than three frames. For the same level of delay, the throughput of a stand-alone PSC network and a stand-alone AWG network are 8 and 12 packets per frame, respectively. So by combining the AWG and the PSC in the AWG||PSC network, we effectively tripled the total combined throughput of two stand-alone networks.

A. Comparison of AWG||PSC Network With AWG||AWG Network and PSC||PSC Network

In this section, we compare the AWG||PSC network to its peers of homogeneous two-device networks. Fig. 12 compares the throughput-delay performance of the AWG||PSC network with a PSC||PSC network (consisting of two PSCs in parallel) and an AWG||AWG network (consisting of two AWGs in parallel). The throughput-delay performance of these homogeneous two-device networks is analyzed in detail in Appendix I.

In brief, in the PSC||PSC network an idle node generates a new packet with probability σ at the beginning of a frame. In the AWG||AWG network an idle node generates a new packet with probability $\sigma_A = 1 - (1 - \sigma)^D$ at the beginning of a cycle and data packets are scheduled with full wavelength reuse, i.e., a scheduling window of one cycle.

We observe that the average throughput of the AWG||PSC network is significantly larger and the delay significantly smaller than for the other two two-device networks. In the PSC||PSC network, we observe a maximum average throughput of 24 packets per frame. We imposed the control packet contention only on one of the devices. This allows for the scheduling of up to two data packets per frame on the second PSC, which effectively allows for the scheduling of up to three data packets per wavelength on the PSC||PSC network in each frame. With $\Lambda = 8$ wavelengths available, the PSC||PSC network has a maximum throughput of 24 data packets per frame. An alternative framing structure is to have control packet contention on both PSCs. This would double the number of contention slots per frame, but would reduce the scheduling capacity to 16 data packets per frame. Since the number of wavelength channels is the obvious bottleneck for the PSC||PSC network, we chose the former framing method to alleviate the bottleneck for data transmission.

For the AWG||AWG network, we present numerical and simulation results for two framing structures. The first framing structure has control contention only on one of the AWGs. The second framing structure (marked 2-M in the plots) has control packet contention slots and data slots imposed on both devices. We observe that the framing structure with control contention on both AWGs achieves larger throughput and smaller delays compared to the framing structure with contention over one AWG. The maximum average throughput for one control slot contention and two control contentions are 37 packets and 42 packets per frame, respectively. Using one control contention per frame, the maximum throughput is $3 \cdot D \cdot \Lambda = 96$ data packets per frame. Using two control contentions per frame, the maximum throughput is $2 \cdot D \cdot \Lambda = 64$ data packets per frame. (Note that a preallocation MAC protocol in conjunction with perfectly smooth periodic packet traffic could achieve a throughput of $4 \cdot D \cdot \Lambda$ data packets per frame over the AWG||AWG network, whereas it could only achieve a maximum throughput of $2 \cdot \Lambda + 2 \cdot D \cdot \Lambda$ over the AWG||PSC network. Packet data traffic is however typically not smooth and periodic, but rather bursty and aperiodic and thus more efficiently accommodated with the pretransmission coordination MAC protocol with dynamic wavelength allocation studied here.) Although the structure with two control contentions per frame has fewer data slots, it has a larger probability of success for control packet contention, thus resulting in larger throughput and smaller delay. The primary reason that the throughput levels in both of these framing structures are significantly smaller than their data scheduling capacity is the lower traffic as a result of the cyclic control packet transmission structure. For $\sigma = 1$ an idle node in the PSC||PSC or AWG||PSC network generates a new packet with probability one at the beginning of a frame, whereas an idle node in the AWG||AWG network generates a new packet with the corresponding probability $\sigma_A = 1$

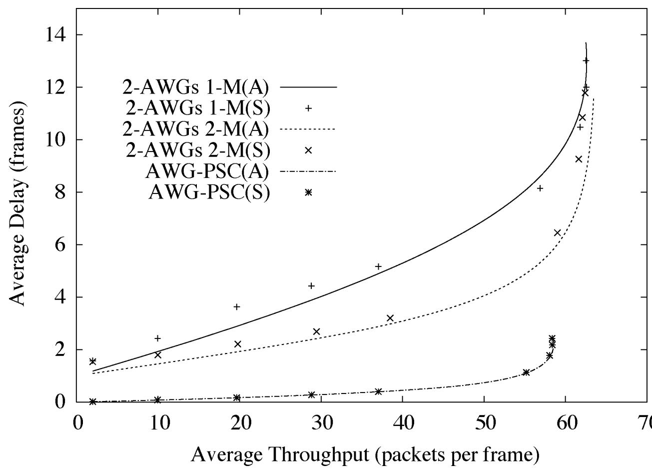

Fig. 13. Throughput-delay performance comparison for three networks: D-buffered AWG||AWG with one control, D-buffered AWG||AWG with two controls, and AWG||PSC.

at the beginning of a cycle (consisting of D frames). In other words, the AWG||AWG network is "fed" with a smaller input traffic rate since each node generates at most one new packet in a cycle. Thus, the maximum number of control packets corresponding to new data packet in a 200-node network with a 4×4 AWG is 50 control packets per frame.

To get a better understanding of the relative performance of the AWG||PSC network with respect to the AWG||AWG network, we consider an alternative operation of the AWG||AWG network, which ensures that both networks are "fed" with the same traffic rate. Specifically, we equip each node in the AWG||AWG network with D packet buffers; one for each of the frames in a cycle. (Each node in the AWG||PSC continues to have only one packet buffer.) Each node in the AWG||AWG network generates a new packet with probability σ at the beginning of a frame if the buffer corresponding to that frame is idle. As explained in Section IV-C the nodes in the AWG||AWG network can only send control packets in the one frame (out of the D frames in the cycle) that is assigned to the node's combiner. Whereas in the single-buffer operation considered in Section IV-C and Section V-F, a node sends at most one control packet in that assigned frame, in the D-buffer operation considered here a node sends up to D control packets—one for each of the packets in its D buffers—in the assigned frame. The control packet contention and data packet scheduling for this D-buffer operation of the AWG||AWG network and the resulting throughput-delay performance are analyzed in detail in Appendix II.

Fig. 13 compares the throughput-delay performance for the AWG||PSC network with the throughput-delay performance of the AWG||AWG network with D-buffer operation, both with control packet contention on one AWG and on two AWGs. We observe that the AWG||AWG network with D-buffer operation achieves somewhat larger throughput than the AWG||PSC network. However, the AWG||PSC network achieves significantly smaller delay throughout. While the comparison in Fig. 13 is fair in that both networks are "fed" with the same traffic rate, the AWG||AWG network is given the advantage of D packet buffers and a scheduling window of D frames

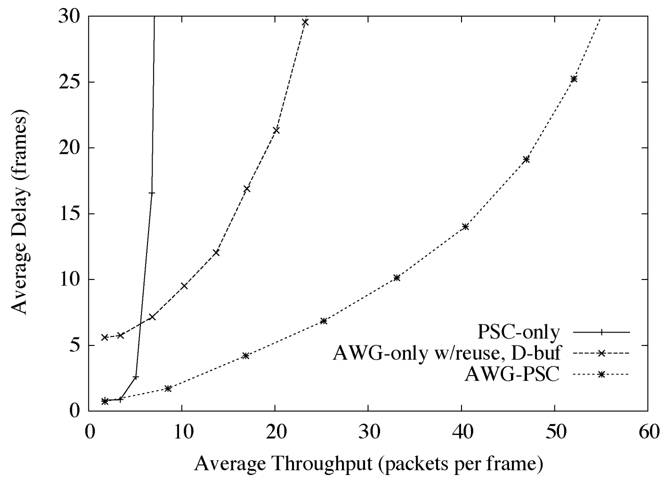

Fig. 14. Throughput-delay performance of AWG||PSC network for three modes of operation for self-similar traffic and large node buffers.

(both resulting in higher complexity), whereas the AWG||PSC network has a single packet buffer and a scheduling window of one frame. The comparisons in both Fig. 12 and Fig. 13 indicate that the AWG||PSC network achieves good throughput-delay performance at low complexity.

B. AWG||PSC Network Performance for Self-Similar Traffic

In this section, we examine the throughput-delay performance of the AWG||PSC network in its three operating modes for a more general traffic model. In particular, we consider self-similar packet traffic with a Hurst parameter of 0.75, which we generate from ON/OFF processes with Pareto distributed on-duration and geometrically distributed off-duration [60]. We equip each node with a large buffer such that essentially no packet is dropped and run simulations to obtain the packet throughput and delay.

In Fig. 14, we compare the throughput-delay performance for the PSC-only mode, the AWG-only mode with wavelength reuse, and the AWG-PSC mode. For the AWG-only mode with wavelength reuse we consider the D-buffer operation in which a node can send up to D control packets in its assigned frame, as introduced in the preceding section. The single-buffer operation of the AWG in which a node can send at most one control packet in its assigned frame (out of the D frames in a cycle) would give significantly larger delays for the persistent packet bursts of the self-similar traffic. We observe from Fig. 14 that similar to the results for Bernoulli traffic in Fig. 11, the AWG-PSC mode achieves significantly larger throughput for a given level of delay than the combined throughput of PSC-only mode and AWG-only mode. For a wide range of delay levels the AWG||PSC network achieves close to twice the throughput of a stand-alone PSC network and a stand-alone AWG network. We also observe by comparing Fig. 14 with Fig. 11 that for a given level of throughput, the delay for each of the operating modes for the self-similar traffic scenario is larger than the delay for the Bernoulli traffic scenario. This is due to the arrival of persistent packet bursts and the node buffering in the self-similar traffic scenario.

VII. CONCLUSION

To address the problem of the single point of failure in single-hop WDM networks, we have proposed and evaluated the AWG||PSC network, a novel single-hop WDM network, consisting of an AWG in parallel with a PSC. The AWG||PSC network achieves high survivability through *heterogeneous protection* (i.e., the AWG and the PSC protect each other); the network remains functional when either the AWG or the PSC fails. The AWG||PSC network provides enhanced throughput-delay performance by exploiting the respective strengths of the AWG (periodic wavelength routing, spatial wavelength reuse) and the PSC (efficient broadcast) during normal operation. We note that the heterogeneous protection proposed and studied in this paper is a general approach, i.e., it can be applied to the PSC based networks reported in the literature in analogous fashion.

Several aspects of the network remain to be explored in detail in future work. One avenue for future work is to examine the throughput-delay performance of the network for a wider variety of general traffic patterns. We also note that the network provides a flexible infrastructure for efficient optical multicasting, which is another topic for future research. A multicast destined to the receivers at one AWG output port could be conducted over the AWG, while a multicast destined to receivers at several AWG output ports may be conducted more efficiently over the PSC.

APPENDIX I
THROUGHPUT-DELAY ANALYSIS FOR PSC||PSC NETWORK AND AWG||AWG NETWORK

In this appendix, we analyze the throughput-delay performance of the PSC||PSC network and the AWG||AWG network. We make the following traffic assumptions for these two homogeneous networks.

- A node selects one of the two devices with equal probability for transmission.
- Each node can have at most one data packet in the buffer to ensure a fair comparison with the AWG||PSC network.

A. PSC||PSC Network

For the PSC||PSC network with control packet contention over one PSC, the control packet contention analysis is the same as in Section V-B. Because we can schedule up to three data packets per frame on each wavelength; one data packet per frame on the PSC with contention phase, two data packets per frame on the PSC dedicated to data, the throughput for the PSC||PSC network is

$$Z_{2PM} = \sum_{i=1}^{3\cdot\Lambda} i\binom{M}{i}\kappa^i(1-\kappa)^{M-i} + 3\cdot\Lambda \cdot \sum_{j=3\Lambda+1}^{M}\binom{M}{j}\kappa^j(1-\kappa)^{M-j}. \quad (17)$$

The equilibrium condition for the PSC||PSC network is $Z_{2PM} = \sigma \cdot E[\eta]$, which is used to solve numerically for η. The average delay (in frames) is $(N - E[\eta])/Z_{2PM}$.

B. AWG||AWG Network

For the AWG||AWG network, we consider two scenarios: 1) control contention over one AWG and 2) control contention over both AWGs. In the case of control contention over one AWG, the contention analysis is the same as in Section V-F. The throughput is modified to reflect the additional two data packets that can be scheduled per FSR per frame on the AWG dedicated to data transmission

$$Z_{1M} = D \cdot \sum_{i=1}^{3\Lambda} i \binom{M}{i} \left(\frac{\kappa_A}{D}\right)^i \left(1 - \frac{\kappa_A}{D}\right)^{M-i} + 3 \cdot R \cdot D^2 \cdot \sum_{j=3\Lambda+1}^{M} \binom{M}{j} \left(\frac{\kappa_A}{D}\right)^j \left(1 - \frac{\kappa_A}{D}\right)^{M-j}. \quad (18)$$

The equilibrium condition is $Z_{1M} = \sigma_A \cdot E[\eta]/D$, which is again used to solve numerically for η.

In the scenario of control contention over both AWGs, we assume that a node selects one of the two devices with equal probability for transmission. We define σ_{2A} as the probability that a given idle node generates a new packet by the beginning of its transmission cycle and sends this control packet to a given AWG. Clearly, $\sigma_{2A} = 1 - (1 - \sigma/2)^D$. Similarly, we define p_{2A} as the probability that a given backlogged node retransmits a control packet over a given AWG at the beginning of a given cycle. Clearly, $p_{2A} = 1 - (1 - p/2)^D$. The probability that a given control slot on a given AWG contains a successfully transmitted control packet is approximately

$$\kappa_{2A} = \frac{\eta\sigma_{2A}}{DM}\left(1 - \frac{\sigma_{2A}}{M}\right)^{\frac{\eta}{D}-1}\left(1 - \frac{p_{2A}}{M}\right)^{\frac{(N-\eta)}{D}} + \frac{(N-\eta)p_{2A}}{DM}\left(1 - \frac{p_{2A}}{M}\right)^{\frac{(N-\eta)}{D}-1}\left(1 - \frac{\sigma_{2A}}{M}\right)^{\frac{\eta}{D}}. \quad (19)$$

This κ_{2A} is used to evaluate the average throughput over a given AWG, which—for a scheduling window of one cycle—is given by

$$Z_{2M} = D \cdot \sum_{i=1}^{\Lambda} i \binom{M}{i} \left(\frac{\kappa_{2A}}{D}\right)^i \left(1 - \frac{\kappa_{2A}}{D}\right)^{M-i} + R \cdot D^2 \cdot \sum_{j=\Lambda+1}^{M} \binom{M}{j} \left(\frac{\kappa_{2A}}{D}\right)^j \left(1 - \frac{\kappa_{2A}}{D}\right)^{M-j}. \quad (20)$$

The equilibrium condition is $Z_{2M} = \sigma_{2A} \cdot E[\eta]/D$, which is again used to solve numerically for η. The average throughput of the AWG||AWG network (in packets per frame) is then given as $2 \cdot Z_{2M}$ and the average delay in the network (in frames) is $(N - E[\eta])/(2 \cdot Z_{2M}) + I_{\text{del}} + (Z_{2M} - D \cdot R)^+/(2 \cdot D \cdot R)$.

Appendix II
Throughput-Delay Analysis for the AWG||AWG Network With D-Buffer Operation

In this appendix, we analyze the throughput-delay performance of the AWG||AWG network with D-buffer operation and full wavelength reuse (i.e., a scheduling window of one cycle). In the D-buffer operation, an idle buffer corresponding to a given frame (out of the D frames in the cycle) generates a new packet with probability σ at the beginning of that frame. In the frame assigned to the node for control packet transmission, control packets are sent for all packets that have been newly generated in the past D frames. In addition, control packets are sent for each backlogged (packet) buffer with probability p. Let η_D denote the total number of idle buffers in the network. Note that the are $D \cdot N - \eta_D$ backlogged buffers in the network. Also note that each frame is assigned N/D nodes for control packet transmission. Thus, in equilibrium, there are $\eta_D/D = \eta$ newly generated packets contenting in a given frame. In addition, there are $(D \cdot N - \eta_D)/D = N - \eta$ backlogged buffers contending in a given frame. Thus, the probability of a control slot containing a successfully (without collision) transmitted control packet is κ given in (1). The throughput of the AWG||AWG network in D-buffer operation with control packet contention on one AWG is thus obtained by replacing κ_A by κ in (18) and σ_A by σ in the corresponding equilibrium condition.

The throughput of the AWG||AWG network in D-buffer operation with control packet contention on two AWGs is obtained by replacing κ_{2A} by

$$\eta\left(\frac{\sigma}{2M}\right)\left(1 - \frac{\sigma}{2M}\right)^{\eta-1}\left(1 - \frac{p}{2M}\right)^{N-\eta} + (N-\eta)\left(\frac{p}{2M}\right)\left(1 - \frac{p}{2M}\right)^{N-\eta-1}\left(1 - \frac{\sigma}{2M}\right)^{\eta} \quad (21)$$

in (20) and σ_{2A} by σ in the corresponding equilibrium condition.

Appendix III
Analysis of Impact of Propagation Delay

Recall that the analysis in Section V assumed that the one-way end-to-end propagation delay in the network is less than one frame. In this appendix, we develop a more general analytical model which accommodates larger propagation delays. This more general model allows us to accurately characterize the performance of the AWG||PSC network for the larger propagation delays in realistic networking scenarios.

For our analysis, we assume that all nodes are equidistant from the central AWG||PSC. (This can be achieved in a straightforward manner by employing standard low-loss fiber delay lines.) Let τ denote the one-way end-to-end (from a given node to the central AWG||PSC and on to an arbitrary node) propagation delay in integer multiples of frames (as defined in Section IV). We furthermore assume that each node has a buffer that holds $\tau + 1$ packets.

In a typical scenario with a distance of 50 km from each node to the central AWG||PSC and a propagation speed of $2 \cdot 10^8$ m/s, the one-way end-to-end propagation delay is 0.5 ms. With an OC48 transmission rate of 2.4 Gb/s and a frame size of 1596 bytes (corresponding to a maximum size Ethernet frame) the propagation delay is $\tau = 94$ frames. (Buffering the corresponding 94 packets requires at most 150 kbytes of buffer in the electronic domain.) Note that if we had considered a frame size corresponding to the maximum size of a SONET frame of 1600 kbytes, the propagation delay would only be a fraction of one frame, which is accommodated by the analysis in Section V.

We now proceed with the analysis for a propagation delay of multiple frames. The basic time unit in our analysis is the slot, i.e., the transmission time of a control packet, as defined in Section IV. Note that a propagation delay of τ frames is equivalent to a delay of $\tau \cdot F$ slots. For our analysis, we introduce the concept of time-sequenced buffering.

A. Time-Sequenced Buffering at Nodes

We view a given node's buffer capable of holding $\tau + 1$ packets as consisting of $\tau + 1$ *buffer slots*, as illustrated in Fig. 15. Each buffer slot can hold one packet. In each frame, one of the buffer slots is the *active* buffer slot. The active buffer slot behaves exactly in the same way as the single-packet buffer considered in Section V, i.e., if idle, it generates a new packet with probability σ and sends a control packet. If backlogged it sends a control packet with probability p.

The other τ buffer slots are *inactive*. The inactive buffer slots do not generate any new packets nor do they send any packets into the network. The purpose of the inactive buffer slots is to hold the data packets that correspond to the control packets that are propagating in the network.

A given buffer slot that is active in a given frame is inactive in the following τ frames (allowing each of the τ other buffer slots to be active for one frame), and then becomes again active $\tau + 1$ frames later.

Suppose a buffer slot is active in a given frame and in one of the M control slots in this frame sends out a control packet. This control packet arrives back at the node by the time the buffer slot becomes again active at the start of the $(\tau+1)$th frame (i.e., after "sitting out" for τ frames). If the control packet is successful in control packet contention and data packet scheduling, the corresponding data packet is sent out in this $(\tau + 1)$th frame.

Also if the control packet is successful, a new data packet is generated with probability σ at the beginning of this $(\tau + 1)$th frame. If a new data packet is generated, the corresponding control packet is sent in one of the M control slots of the $(\tau + 1)$th frame. Note that we have tacitly assumed here that the nodal processing takes no more than $F - M$ slots. If the processing delay is larger, it can be accommodated in a straightforward manner by adding more buffer slots.

For an illustration of the concept of time-sequenced buffering, consider the buffer slots of a given node depicted in Fig. 15. Suppose buffer slot 1 is empty prior to time $t = 0$, and generates a new packet, designated by D(1), at $t = 0$. The control packet corresponding to D(1), designated by C(1), is sent in one of the M control slots of the frame that is sent between $t = 0$ and $t = F$ (slots). By the time $t = F$, this frame is completely "on the fiber," as illustrated in the second snapshot in Fig. 15. (Note that this frame contains no data packets, as we assumed that buffer slot 1 was empty before $t = 0$.) At $t = F$, buffer slot 1 becomes inactive, while buffer slot 2 becomes active. Suppose the node generates a new data packet D(2) at $t = F$. At $t = 2F$ the frame with the control packet C(2) is completely on the fiber and buffer slot 3 becomes active, and so on.

At time $t = \tau F$ the frame containing C(1) starts to arrive back at the node. By time $t = \tau F + M$, the control packet is completely received and its processing commences. With an assumed processing delay of less than $F - M$ slots, the processing is completed by $t = (\tau + 1)F$, which is exactly when buffer slot 1 becomes again active. Suppose C(1) was successful and the corresponding D(1) is scheduled on the AWG. Also suppose a new data packet $D(\tau + 2)$ is generated at $t = (\tau + 1)F$. By $t = (\tau + 2)F$, the frame containing $D(1)$ and $C(\tau + 2)$ is completely on the fiber, and buffer slot 2 becomes active, and so on.

B. Network Analysis

The key insight to the analysis of the network with time-sequenced buffering at the nodes is that in steady state it suffices to consider only the active buffer slot at each of the N network nodes. Specifically, at each instance in time, each node has exactly one active buffer slot. This active buffer slot is either idle or backlogged (similar to the way a node is either idle or backlogged in the analysis of Section V). A buffer slot is considered idle if: 1) it contains no data packet or 2) it successfully transmitted a control packet the last time it was active and the corresponding data packet has been successfully scheduled (although this data packet may still be in the buffer slot.)

An active buffer slot is considered backlogged if it contains a data packet whose corresponding control packet failed in the control packet contention or data packet scheduling. Let η denote the number of idle nodes (active buffer slots). Clearly, the number of backlogged nodes (active buffer slots) is $N - \eta$.

Now note that the control packet contention with time-sequenced buffer in a given frame is analogous to the control packet contention with the single-packet buffer considered in Section V. In a given frame, each of the η idle active buffer slots generates a new data packet and sends a control packet with probability σ. Each of the $N - \eta$ backlogged active buffer slots retransmits a control packet with probability p. Thus, the expected number of successful control packets per frame is $M \cdot \kappa$, as given in Section V-B.

Next note that the time-sequenced buffering does not interfere with the data packet scheduling as described in Section IV and analyzed in Section V. Thus, the throughput results derived for the different operating modes in Section V apply without any modification to the time-sequenced buffer scenario.

Finally, note that the delays for the different operating modes as derived in Section V are scaled by the propagation delay of τ frames when considering the time-sequenced buffer scenario. Specifically, for the AWG-PSC mode, there is a delay component of τ frames for the initial control packet. In addition, there is a delay component due to control packet retransmissions (if control packet contention or data packet scheduling failed). This second delay component is the expected number of backlogged nodes $N - E[\eta]$ divided by the expected throughput $Z_A + Z_P$ (similar to the case analyzed in Section V-D), but is now scaled by the propagation delay τ. Thus, the average delay is

$$\text{Delay} = \tau \cdot \left(1 + \frac{N - \eta}{Z_P + Z_A}\right)$$

in frames, where we make again the reasonable approximation $E[\eta] \approx \eta$.

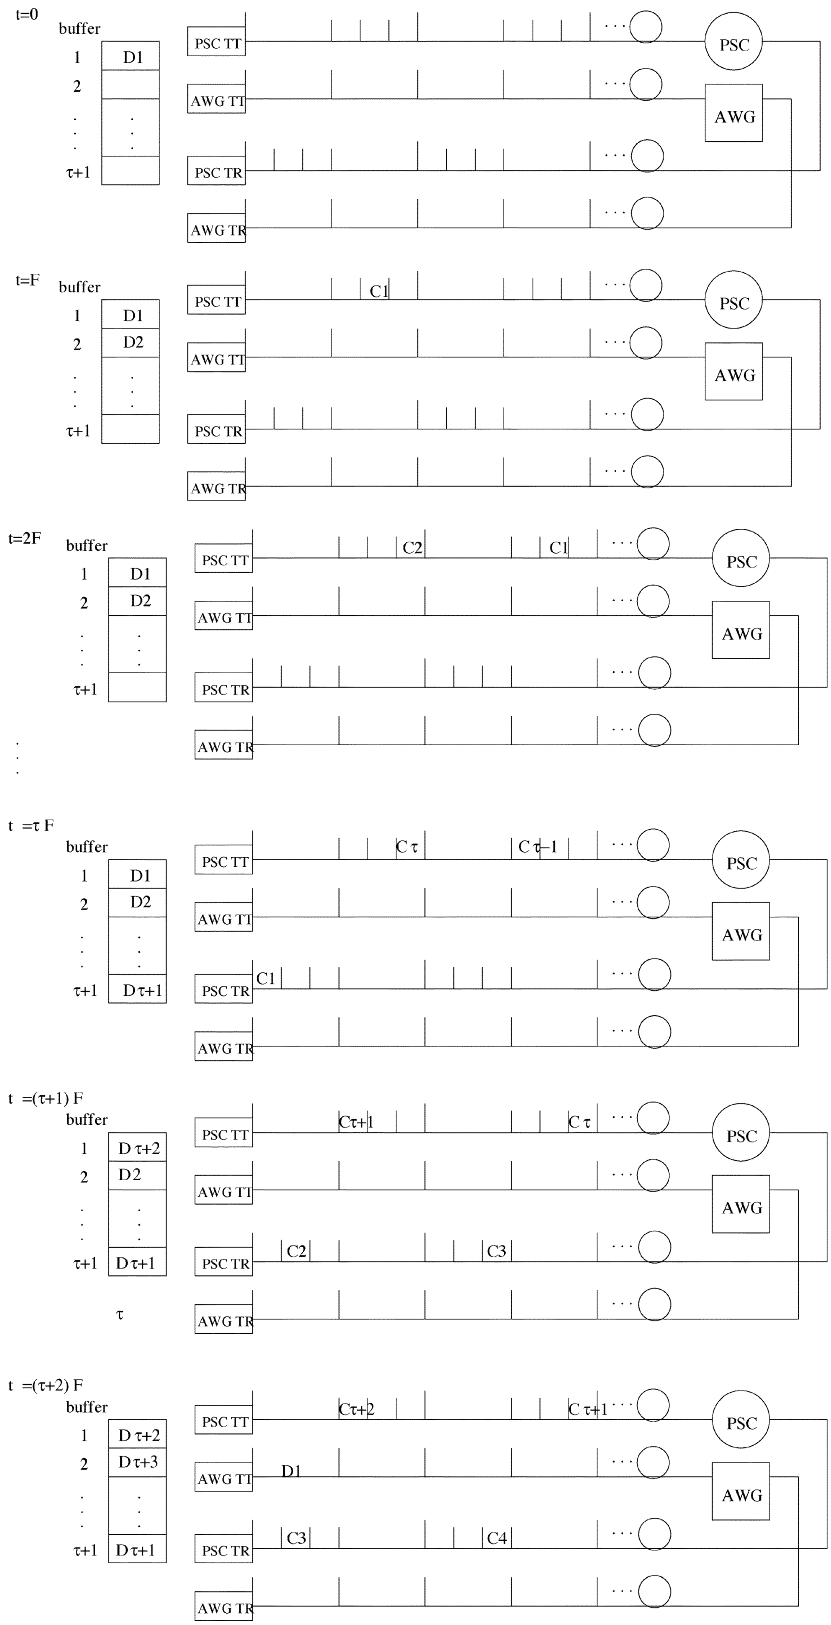

Fig. 15. Illustration of time-sequenced buffering.

In analogous fashion, the average delay for the PSC-only mode is

$$\text{Delay} = \tau \cdot \left(1 + \frac{N - \eta}{Z_P}\right) \text{ frames.}$$

As discussed in Section V-F, in the AWG-only mode with wavelength reuse, there are two additional delay components, cyclic control transmission delay I_{del} and scheduling delay if the data packet is not immediately transmitted. These two delay components are not affected by the propagation delay. Thus, the

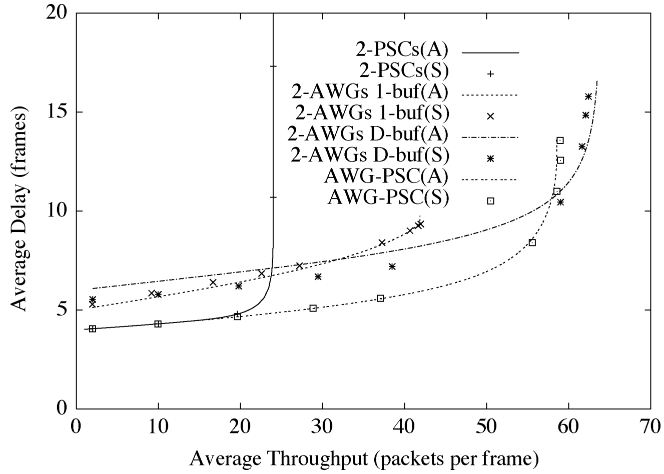

Fig. 16. Throughput-delay performance comparison for two-device networks for a propagation delay of $\tau = 4$ frames ($N = 200$ fixed).

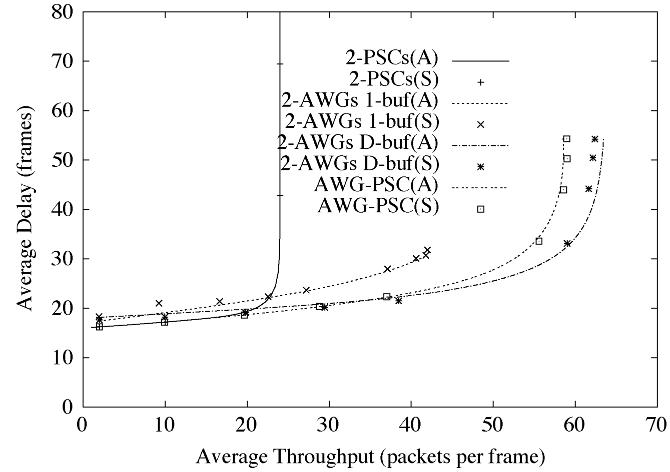

Fig. 17. Throughput-delay performance comparison for two-device networks for a propagation delay of $\tau = 16$ frames ($N = 200$ fixed).

average delay (in frames) for the AWG-only mode with spatial wavelength reuse is

$$\text{Delay}_{RE} = \tau \cdot \left(1 + \frac{N - \eta}{Z_{AM}}\right) + I_{\text{del}} + \frac{(Z_{RE} - D \cdot R)^{+}}{2 \cdot D \cdot R}.$$

C. Numerical and Simulation Results

In this section, we examine the throughput-delay performance of the two-device networks, AWG||PSC, AWG||AWG, and PSC||PSC with time sequenced buffering. For the AWG||AWG network we consider both single buffer and D-buffer operation. For the D-buffer operation we combine the time-sequenced buffering introduced in this appendix with the D packet buffers analyzed in Appendix II, for a total of $D \cdot (\tau + 1)$ packet buffers at each node of the AWG||AWG network with D-buffer operation. (Each node has only $\tau + 1$ packet buffers in the other considered networks.) Throughput we consider the AWG||AWG network with control packet contention on both AWGs and a scheduling window of D frames (the PSC||PSC and AWG||AWG networks have a scheduling window of one frame.) The numerical and simulation results are presented for one-way end-to-end propagation delays of $\tau = 4$ frames, $\tau = 16$ frames, and $\tau = 96$ frames in Fig. 16, Fig. 17, and Fig. 18, respectively. We observe that the throughputs for all of the networks are independent of the τ values and are the same. The throughput for the three networks are also the same as the throughput for a propagation delay of less than one frame, see Fig. 12. Thus, the time-sequenced buffering allows us to effectively utilize the full transmission capacity of the networks even for large propagation delays. Also it allows us to apply the probabilistic analytical model developed in Section V.

We observe that the AWG||PSC network has smaller delay compared to the AWG||AWG network for small τ. As the propagation delay τ increases the gap in delay between the AWG||PSC network and the AWG||AWG network becomes smaller. For small τ, the relatively larger delay for the AWG||AWG network is due to the cyclic control packet transmission. As τ increases the delay due to the cyclic control

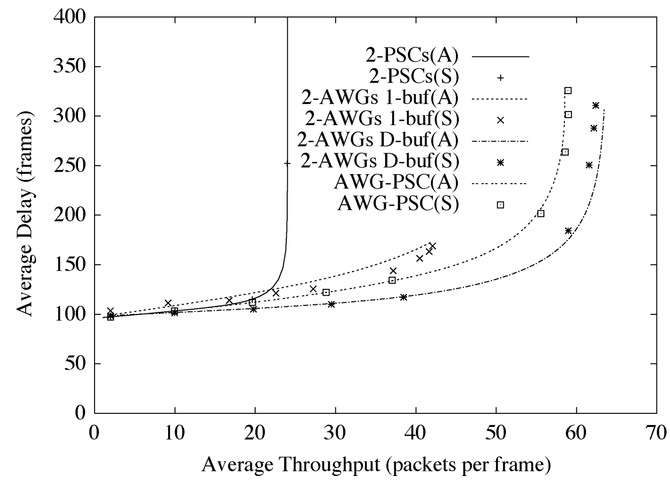

Fig. 18. Throughput-delay performance comparison for two-device networks for a propagation delay of $\tau = 96$ frames ($N = 200$ fixed).

packet transmission becomes less and less dominant. We also observe that the single-buffer AWG||PSC network gives larger throughput than the single-buffer AWG||AWG network. The throughput of the D-buffer AWG||AWG network is somewhat larger (at the expense of more complexity) than the throughput of the single-buffer AWG||PSC network. Overall, the results indicate that the low-complexity AWG||PSC network gives favorable throughput-delay performance for realistic propagation delays.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their insightful comments on an earlier version of the paper, which have greatly improved this work.

REFERENCES

[1] B. Mukherjee, "WDM-based local lightwave networks part I: single-hop systems," *IEEE Network*, vol. 6, pp. 12–27, May 1992.
[2] A. Bianco, E. Leonardi, M. Mellia, and F. Neri, "Network controller design for SONATA—a large-scale all-optical passive network," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 2017–2028, Oct. 2000.

[3] N. P. Caponio, A. M. Hill, F. Neri, and R. Sabella, "Single-layer optical platform based on WDM/TDM multiple access for large-scale 'switchless' networks," *Eur. Trans. Telecommun.*, vol. 11, no. 1, pp. 73–82, Jan./Feb. 2000.
[4] K. Kato, A. Okada, Y. Sakai, K. Noguchi, T. Sakamoto, A. Takahara, A. Kaneko, S. Suzuki, and M. Matsuoka, "10-Tbps full-mesh WDM network based on a cyclic-frequency arrayed-waveguide grating router," in *Proc. Eur. Conf. Optical Communications 2000 (ECOC'00)* , vol. 1, Munich, Germany, Sept. 2000, pp. 105–107.
[5] M. Maier, *Metropolitan Area WDM Networks—An AWG Based Approach*. Norwell, MA: Kluwer, 2003.
[6] 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," *Comput. Networks*, vol. 41, no. 4, pp. 407–433, Mar. 2003.
[7] A. Okada, T. Sakamoto, Y. Sakai, K. Noguchi, and M. Matsuoka, "All-optical packet routing by an out-of-band optical label and wavelength conversion in a full-mesh network based on a cyclic-frequency AWG," in *Optical Fiber Communication Conf. and Exhibit 2001 (OFC'01) Tech. Dig.* , Anaheim, CA, Mar. 2001, Paper ThG5, pp. ThG5-1–ThG5-3.
[8] F. Ruehl and T. Anderson, "Cost-effective metro WDM network architectures," in *Optical Fiber Communication Conf. and Exhibit 2001 (OFC'01) Tech. Dig.* , Anaheim, CA, Mar. 2001, Paper WL1, pp. WL1-1–WL1-3.
[9] M. Maier and A. Wolisz, "Demonstrating the potential of arrayed-waveguide grating based single-hop WDM networks," *Opt. Networks Mag.*, vol. 2, no. 5, pp. 75–85, Sept. 2001.
[10] M. Scheutzow, M. Maier, M. Reisslein, and A. Wolisz, "Wavelength reuse for efficient packet-switched transport in an AWG-based metro WDM network," *J. Lightwave Technol.*, vol. 21, pp. 1435–1455, June 2003.
[11] A. Hill, M. Brierley, R. Percival, R. Wyatt, D. Pitcher, K. Pati, I. Hall, and J.-P. Laude, "Multistar wavelength-router network and its protection strategy," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 1134–1145, Sept. 1998.
[12] Y. Sakai, K. Noguchi, Y. Yoshimura, T. Sakamoto, A. Okada, and M. Matsuoka, "Management system for full-mesh WDM AWG-star network," in *Proc. 27th Eur. Conf. Optical Communications*, vol. 3, Amsterdam, Netherlands, Sept./Oct. 2001, pp. 264–265.
[13] M. Maier, M. Reisslein, and A. Wolisz, "Toward efficient packet switching metro WDM networks," *Opt. Networks Mag.*, vol. 3, no. 6, pp. 44–62, Nov./Dec. 2002.
[14] M.-S. Chen, N. Dono, and R. Ramaswami, "A media-access protocol for packet-switched wavelength division multiaccess metroplitan area networks," *IEEE J. Select. Areas Commun.*, vol. 8, pp. 1048–1057, Aug. 1990.
[15] R. Chipalkatti, Z. Zhang, and A. S. Acampora, "Protocols for optical star-coupler network using WDM: performance and complexity study," *IEEE J. Select. Areas Commun.*, vol. 11, pp. 579–589, May 1993.
[16] P. W. Dowd, "Random access protocols for high-speed interprocessor communication based on an optical passive star-topology," *J. Lightwave Technol.*, vol. 9, pp. 799–808, June 1991.
[17] A. Ganz and Y. Gao, "Time-wavelength assignment algorithms for high performance WDM star based systems," *IEEE Trans. Commun.*, vol. 42, pp. 1827–1836, Feb./March/April 1994.
[18] I. M. I. Habbab, M. Kavehrad, and C.-E. W. Sundberg, "Protocols for very high-speed optical fiber local area networks using a passive star topology," *J. Lightwave Technol.*, vol. LT-5, pp. 1782–1794, Apr. 1987.
[19] M. J. Karol and B. Glance, "A collision-avoidance WDM optical star network," *Comput. Networks ISDN Syst.*, vol. 26, pp. 931–943, Mar. 1994.
[20] K. Li, "Probabilistic analysis of cyclic packet transmission scheduling in WDM optical networks," in *Proc. IEEE Int. Conf. Parallel Processing* , Toronto, ON, Canada, Mar. 2000, pp. 531–538.
[21] H. C. Lin and C. H. Wan, "A hybrid multicast scheduling algorithm for single-hop WDM networks," *J. Lightwave Technol.*, vol. 19, pp. 1654–1664, Nov. 2001.
[22] J. Lu and L. Kleinrock, "A wavelength division multiple access protocol for high-speed local area networks with a passive star topology," *Perform. Evaluat.*, vol. 16, no. 1–3, pp. 223–239, Nov. 1992.
[23] N. Mehravari, "Performance and protocol improvements for very high speed optical fiber local area networks using a passive star topology," *J. Lightwave Technol.*, vol. 8, pp. 520–530, Apr. 1990.
[24] Y. Ofek and M. Sidi, "Design and analysis of hybrid access control to an optical star using WDM," in *Proc. IEEE Infocom'91*, Bal Harbour, FL, May 1991, pp. 20–31.
[25] G. N. M. Sudhakar, N. D. Georganas, and M. Kavehrad, "Slotted Aloha and reservation Aloha protocols for very high-speed optical fiber local area networks using passive star topology," *J. Lightwave Technol.*, vol. 9, pp. 1411–1422, Oct. 1991.
[26] B. Kannan, S. Fotedar, and M. Gerla, "A two level optical star WDM metropolitan area network," in *Proc. IEEE Globecom'94 Communications* , San Francisco, CA, Nov. 1994, pp. 563–566.
[27] M. Janoska and T. Todd, "Coupled reservation protocols for hierarchical single-hop photonic networks," *Inst. Elect. Eng. Proc. Commun.*, vol. 144, no. 4, pp. 247–255, Aug. 1997.
[28] C. J. Chae, H. Park, and Y. Park, "Hybrid optical star coupler suitable for wavelength reuse," *IEEE Photon. Technol. Lett.*, vol. 10, pp. 279–281, Feb. 1998.
[29] S. B. Alexander, R. Bondurant, D. Byrne, V. Chan, S. Finn, R. Gallager, B. Glance, H. Haus, P. Humblet, R. Jain, I. Kaminow, M. Karol, R. Kennedy, A. Kirby, H. Le, A. Saleh, B. A. Schofield, J. Shapiro, N. Shankaranarayanan, R. Thomas, R. Williamson, and R. Wilson, "A precompetitive consortium on wide-band all-optical networks," *IEEE J. Lightwave Technol.*, vol. 11, pp. 714–735, May/June 1993.
[30] D. Banerjee, J. Frank, and B. Mukherjee, "Passive optical network architecture based on waveguide grating routers," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 1040–1050, Sept. 1998.
[31] B. Glance, I. P. Kaminow, and R. W. Wilson, "Applications of the integrated waveguide grating router," *J. Lightwave Technol.*, vol. 12, pp. 957–962, June 1994.
[32] B. Glance and M. Karol, "Large capacity multiaccess optical packet network," *IEEE Photon. Technol. Lett.*, vol. 6, pp. 872–875, July 1994.
[33] A. Hill, S. Carter, J. Armitage, M. Shabeer, R. Harmon, and P. Rose, "A scalable and switchless optical network structure, employing a single 32 × 32 free-space grating multiplexer," *IEEE Photon. Technol. Lett.*, vol. 8, pp. 569–571, Apr. 1996.
[34] D. K. Hunter, M. H. M. Nizam, M. C. Chia, I. Andonovic, K. Guild, A. Tzanakaki, M. O'Mahony, L. Bainbridge, M. Stephens, R. Penty, and I. White, "WASPNET: A wavelength switched packet network," *IEEE Commun. Mag.*, vol. 37, pp. 120–129, Mar. 1999.
[35] D. Jung, S. Shin, C. Lee, and Y. Chung, "Wavelength-division-multiplexed passive optical network based on spectrum-slicing techniques," *IEEE Photon. Technol. Lett.*, vol. 10, pp. 1334–1336, Sept. 1998.
[36] M. J. Spencer and M. Summerfield, "WRAP: a medium access control protocol for wavelength-routed passive optical networks," *IEEE J. Lightwave Technol.*, vol. 18, pp. 1657–1676, Dec. 2000.
[37] K. Bengi, "Performance of single-hop WDM LAN's supporting real-time services," *Photon. Network Commun.*, vol. 1, no. 4, pp. 287–301, Dec. 1999.
[38] E. Arthurs, J. Cooper, M. Goodman, H. Kobrinski, M. Tur, and M. Vecchi, "Multiwavelength optical crossconnect for parallel-processing computers," *Inst. Elect. Eng. Electron. Lett.*, vol. 24, no. 2, pp. 119–120, Jan. 1988.
[39] *Fiber Channel Standard, Annex I*, FC-PH4.3, 1993.
[40] C.-S. Li and R. Ramaswami, "Automatic faul dectection, isolation, and recovery in transparent all-optical networks," *IEEE J. Lightwave Technol.*, vol. 15, pp. 1784–1793, Oct. 1997.
[41] M. H. M. Nizam, D. K. Hunter, and D. G. Smith, "A dynamic reconfiguring tool for improving multiwavelength transport network robustness," in *Proc. IEEE Int. Conf. Communications 1997 (ICC'97)* , vol. 1, Montreal, QC, Canada, June 1997, pp. 246–250.
[42] H. Takeshita, S. Takahashi, T. Kato, S. Kitamura, H. Harano, and N. Henmi, "A demonstration of an optical cross-connect system for a self-healing optical network," in *Proc. 23rd Eur. Conf. Optical Communications* , vol. 3, Edinborough, U.K., Sept. 1997, pp. 335–338.
[43] A. Saleh and H. Kogelnik, "Reflective single-mode fiber-optic passive star couplers," *J. Lightwave Technol.*, vol. 6, pp. 392–398, Mar. 1988.
[44] M. Tabiani and M. Kavehrad, "Theory of an efficient N × N passive optical star-coupler," *J. Lightwave Technol.*, vol. 9, pp. 445–448, Apr. 1991.
[45] C. Dragone, "Optimum design of a planar array of tapered waveguides," *J. Opt. Soc. Amer.*, vol. 7, pp. 2081–2093, Nov. 1990.
[46] C. Dragone, C. Edwards, and R. Kistler, "Integrated optics $N \times N$ multiplexer on silicon," *IEEE Photon. Technol. Lett.*, vol. 3, pp. 896–899, Oct. 1991.
[47] Y. Hibino, "An array of photonic filtering advantages," *Circuits Devices*, pp. 23–27, Nov. 2000.
[48] K. McGreer, "Arrayed waveguide gratings for wavelength routing," *IEEE Commun. Mag.*, vol. 36, pp. 62–68, Dec. 1998.
[49] Y. Tachikawa, Y. Inoue, M. Ishii, and T. Nozawa, "Arrayed waveguide grating multiplexer with loop-back optical paths and its applications," *J. Lightwave Technol.*, vol. 14, pp. 977–984, June 1996.

[50] P. Bernasconi, C. Doerr, C. Dragone, M. Cappuzzo, E. Laskowski, and A. Paunescu, "Large $n \times n$ waveguide grating routers," *J. Lightwave Technol.*, vol. 18, pp. 985–991, July 2000.

[51] A. Yu, M. J. O'Mahony, and A. M. Hill, "Transmission limitation of all-optical network based on $N \times N$ multi/demultiplexer," *Electron. Lett.*, vol. 33, pp. 1068–1069, June 1997.

[52] O. A. Lavrova, G. Rossi, and D. J. Blumenthal, "Rapid tunable transmitter with large number of ITU channels accessible in less than 5 ns," in *Proc. ECOC'00*, vol. 2, Munich, Germany, Sept. 2000, pp. 169–170.

[53] D.-R. Wonglumsom, I. M. White, K. V. Shrikhande, M. S. Rogge, S. M. Gemelos, F.-T. An, Y. Fukashiro, M. Avenarius, and L. G. Kazovsky, "Experimental demonstration of an access point for HORNET—a packet-over-WDM multiple-access MAN," *J. Lightwave Technol.*, vol. 18, pp. 1709–1717, Dec. 2000.

[54] B. Mukherjee, *Optical Communication Networks*. New York: McGraw-Hill, 1997.

[55] A. Similjanic, "An efficient channel access protocol for an optical star network," in *Proc. IEEE Int. Conf. Communications (ICC'98)*, Atlanta, GA, 1998, pp. 514–519.

[56] K. M. Sivalingam, "Design and analysis of a media access protocol for star coupled WDM networks with TT-TR architecture," in *Optical WDM Networks—Principles and Practice*, K. M. Sivalingam and S. Subramaniam, Eds. Norwell, MA: Kluwer, 2000, ch. 9.

[57] M. Cerisola, T. Fong, R. Hofmeister, L. Kazovsky, C. Lu, A. Mekkittikul, and P. Poggiolini, "Novel distributed slot synchronization technique for optical WDM packet networks," in *Optical Fiber Communication Conf. and Exhibit (OFC'96) Tech. Dig.* , 1996, Paper FD1, pp. 314–315.

[58] R. Hofmeister, C. Lu, M. Ho, P. Poggiolini, and L. Kazovsky, "Distributed slot synchronization (DSS): a network-wide slot synchronization technique for packet-switched optical networks," *J. Lightwave Technol.*, vol. 16, pp. 2109–2116, Dec. 1998.

[59] R. Ramaswami and K. Sivarajan, *Optical Networks: A Practical Perspective*. New York: Morgan Kaufmann, 2002.

[60] K. Park and W. Willinger, Eds., *Self-Similar Network Traffic and Performance Evaluation*. New York: Wiley, 2000.

Chun Fan received the B.S. degree in industrial engineering from the University of Arizona, Tucson, in 1990, the Masters' degree in international management from the American Graduate School of International Management, Thunderbird, AZ, in 1992, and the M.S. degree in electrical engineering from Arizona State University, Tempe, in 2002. He is currently working toward the Ph.D. degree in the Department of Electrical Engineering, Arizona State University.

He has held a variety of management positions with Golden Optic Disc Company, Beijing, China (1990–1991), Emerson Motor Company (1992–1997), Beijing Feijie Hi-Tech Development Company, Beijing, China (1997–1998), and Advanced Fiber Communications, Petaluma, CA (1998–2000).

Martin Maier received the Dipl.-Ing. and the Dr.-Ing. degrees (both with distinctions) in electrical engineering from the Technical University of Berlin, Berlin, Germany, in 1998 and 2003, respectively.

He was a Visiting Researcher at the University of Southern California, Los Angeles, Arizona State University, Tempe, and the Massachusetts Institute of Technology, Cambridge. He also participates in the National Research Project TransiNet. He is author of *Metropolitan Area WDM Networks—An AWG Based Approach* (Norwell, MA: Kluwer, 2003). His research interests include network and node architectures, routing and switching paradigms, protection, multicasting, and the design, performance evaluation, and optimization of MAC protocols for optical WDM networks, with a particular focus on metro networks.

Dr. Maier was the recipient of the two-year Deutsche Telekom Doctoral Scholarship from June 1999 through May 2001. He is also a corecipient of the Best Paper Award, presented at the SPIE Photonics East 2000—Terabit Optical Networking Conference.

Martin Reisslein (A'96–S'97–M'98–SM'03) received the Dipl.-Ing. (FH) degree from the Fachhochschule Dieburg, Dieburg, Germany, in 1994 and the M.S.E. degree from the University of Pennsylvania, Philadelphia, in 1996, both in electrical engineering, and the Ph.D.degree in systems engineering from the University of Pennsylvania, in 1998.

He is an Assistant Professor with the Department of Electrical Engineering, Arizona State University, Tempe. During the academic year 1994–1995, he visited the University of Pennsylvania as a Fulbright Scholar. From July 1998 through October 2000, he was a Scientist with the German National Research Center for Information Technology (GMD FOKUS), Berlin, Germany. While there, he taught courses on performance evaluation and computer networking at the Technical University, Berlin. His research interests are in the areas of Internet quality of service, video traffic characterization, wireless networking, and optical networking.

Dr. Reisslein is Editor-in-Chief of the *IEEE Communications Surveys and Tutorials* and has served on the Technical Program Committees of IEEE Infocom, IEEE Globecom, and the IEEE International Symposium on Computer and Communications. He has organized sessions at the IEEE Computer Communications Workshop (CCW). He maintains an extensive library of video traces for network performance evaluation, including frame size traces of MPEG-4 and H.263 encoded video at http://trace.eas.asu.edu. He is corecipient of the Best Paper Award of the SPIE Photonics East 2000—Terabit Optical Networking Conference.