Ring in the new for WDM resilient packet ring

MARTIN MAIER
AND MARTIN REISSLEIN

THE NEW HIGH-PERFORMANCE standard IEEE 802.17 for the resilient packet ring (RPR) aims to combine the appealing functionalities from synchronous optical network/synchronous digital hierarchy (SONET/SDH) networks with the advantages of Ethernet networks.

RPR is a ring-based architecture consisting of two counterdirectional optical fiber rings. Similar to SONET/SDH networks, the RPR is designed to provide fast recovery from a single link or node failure through protection mechanisms that ensure high availability and reliability. SONET/SDH efficiently supports legacy voice service through time-division multiplexing (TDM), in which time is divided into slots and a node sends its traffic in a periodically recurring fixed assigned slot. RPR is designed to carry this legacy TDM traffic with a high level of quality of service (QoS), e.g., low delay jitter. Similar to Ethernet, RPR provides the advantages of low equipment cost and simplicity and achieves an improved bandwidth utilization due to statistical multiplexing. With statistical multiplexing, a node is not assigned a prescribed slot for transmission; instead, the node can dynamically adjust its transmission rate according to the variations of the traffic, which is more efficient than TDM transmission for highly variable data traffic, for example, from Web applications.

The bandwidth utilization in RPR is further increased by means of spatial reuse. Spatial reuse is achieved in RPR through so-called destination stripping, which means that the destination node takes a transmitted packet off the fiber ring. Thus, a given transmission traverses only the ring segment from the source node to the destination node, allowing other nodes on the ring segment between the destination node and the source node to exchange transmissions at the same time on the same fiber ring. In other words, the destination stripping enables nodes in different ring segments to transmit simultaneously, resulting in spatial reuse and an increased bandwidth utilization. Furthermore, RPR provides fairness (unlike today's Ethernet) and allows the full ring bandwidth to be utilized under normal (that is, failure free) operation conditions (unlike today's SONET/SDH rings, where 50% of the available bandwidth is reserved for protection). Current RPR networks are single-channel systems; that is, each fiber carries a single wavelength channel and is expected to be primarily deployed in metropolitan areas.

©DIGITAL STOCK, 1996

Today's metropolitan-area (metro) networks present a significant bandwidth bottleneck between the increasingly higher-speed local and access networks and the huge bandwidth pipes of backbone networks. This bottleneck, often called a *metro gap*, prevents end users from tapping into the vast amount of backbone bandwidth. Next-generation metro networks must bridge the metro gap in order to tap into the vast amount of backbone bandwidth, enable new emerging services, and stimulate revenue growth. To this end, RPR is likely to be upgraded from a single-channel system to a multichannel system by means of wavelength division multiplexing (WDM), where multiple wavelength channels are used on each fiber.

This article describes two complementary approaches for WDM upgrading RPR. The first approach deploys WDM on the fiber ring; it does not require any additional fiber infrastructure but does require that *all* ring nodes are upgraded to support WDM. The second approach augments the RPR network by an additional single-hop star WDM subnetwork, which requires additional fibers. However, only a subset of the RPR nodes need to be upgraded to support WDM and attached to the star WDM subnetwork. This second approach enables cautious pay-as-you-grow nodal upgrades that can proceed in incremental fashion according to the growing traffic demands, but it requires the initial investment into the additional fibers to form the star subnetwork. Therefore, the first WDM upgrade of RPR saves on fiber requirements while the second one saves on nodal upgrade requirements.

Overview of resilient packet ring (RPR)

As depicted in Fig. 1, RPR is an optical dual-fiber bidirectional ring network where each fiber ring carries a single wavelength channel. Destination stripping in conjunction with shortest path routing is deployed to improve the spatial reuse of bandwidth. With shortest path routing, a packet is sent in the ring direction that provides the shortest path in terms of the number of hops (that is, the number of traversed intermediate

0278-6648/06/$20.00 © 2006 IEEE

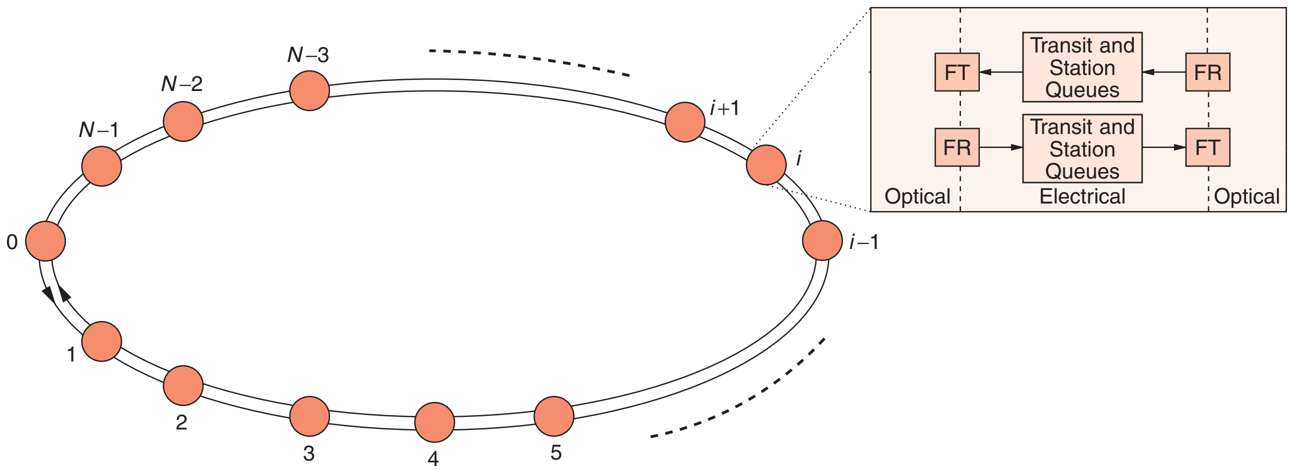

Fig. 1 Illustration of a resilient packet ring (RPR) network interconnecting *N* nodes and the node architecture

nodes) between the source node and the destination node. Each node is equipped with two fixed-tuned transmitters (FTs) and two fixed-tuned receivers (FRs), one for each fiber ring. The FTs transmit the traffic in the form of optical signals onto the fibers, while the FRs convert the incoming optical signal into an electrical signal; that is, each node performs optical-electrical-optical (OEO) conversion. Each node has separate (electrical) transit and station queues for either ring. Transit queues buffer the in-transit traffic traversing the node while station queues buffer the locally generated and received traffic.

Each RPR node employs a prioritized transmission policy. In particular, each node gives priority to forwarding in-transit traffic that traverses the node over the transmission of the node's own (locally generated) traffic. This ensures that the transit path is lossless; i.e., once a packet is put on the ring, it is not dropped at downstream nodes. On the downside, however, a node with a backlog of locally generated traffic must wait for the transit path to clear before it can send its own traffic. As a consequence, upstream nodes with a lot of traffic to send can keep the transit path congested, which in turn prevents downstream nodes from sending their own traffic. This phenomenon is referred to as *starving the downstream nodes* and gives rise to fairness problems. To achieve fairness among nodes, RPR nodes deploy a distributed fairness control algorithm. In this algorithm, a backlogged node based on local measurements sends fairness control packets to upstream nodes in order to throttle their data transmission rates and, thus, alleviate the congestion.

Finally, RPR provides resilience against any single link or node failure by means of wrapping and steering protection mechanisms. Wrapping occurs locally and requires both nodes adjacent to the failure to perform protection switching. With wrapping, upon detection of a link or node failure, the two ring nodes adjacent to the failed link or node switch all traffic arriving on the incoming fiber onto the outgoing fiber to reach the destination node going in the opposite direction. Steering is achieved by modifying the routing tables of each node after learning that a failure has occurred. With steering, after learning that a failure has occurred, a given source node injects the traffic in the direction opposite to the link or node failure, i.e., the source node steers the traffic away from the failure.

WDM ring networks

A straightforward approach to WDM upgrading RPR is the use of WDM on the bidirectional fiber ring, resulting in a WDM ring network. This approach leverages on the existing fiber infrastructure of RPR and does not require additional fibers. On the downside, however, all RPR nodes need to be WDM upgraded since WDM is applied on the entire ring. A plethora of WDM ring node architectures as well as appropriate multichannel medium access control (MAC) and fairness control protocols have been proposed and examined to date. Figure 2 depicts an example node architecture of a WDM ring network with N nodes for either ring direction. At each node, all N incoming wavelengths are demultiplexed. A single wavelength λ_{drop} is dropped at each node. The drop wavelength is routed to a burst-mode receiver while the status of the remaining $(N-1)$ wavelength channels is monitored by using 90/10 taps and an array of photodiodes. The burst-mode receiver recovers the clock of an arriving burst (packet) very quickly and does not need to receive a continuous signal. Each 90/10 tap splits off 10% of the optical power from the fiber. Subsequently, the wavelengths are multiplexed onto the outgoing ring fiber. With a 50/50 combiner and an external modulator, each node is able to send packets by activating one or more transmitters, each fixed tuned to a different wavelength channel. The 50/50 combiner collects signals from two input ports and equally combines them onto a common output port.

WDM rings provide an increased capacity by deploying multiple wavelength channels on each fiber. Furthermore, WDM node architectures such as our example above allow the design of all-optical (OOO) node structures, in which a part of the optical signal (namely the signals on all wavelengths, except the drop wavelength λ_{drop}) remains in the optical domain and does not need to be converted into the electrical domain, electronically processed, and converted back to an optical signal. In doing so, a part of the incoming signal is able to optically bypass the node, avoiding OEO conversion. The resultant OOO node structures provide transparency in the sense that the optically bypassed nodes do not need to be equipped to electronically process and support the protocol, data rate, and modulation format used to generate the optical signals on the bypassed wavelengths. This transparency facilitates the network support of a wide variety of both legacy and future services. On the downside, WDM ring networks suffer from some limitations. First, similar to RPR, WDM rings are able to recover only from a single link or node failure. Also, ring networks provide only a small degree of connectivity; that is, each node is only connected to two other nodes. Consequently, due to missing alternate physical routes, traffic generally has to traverse multiple intermediate ring nodes—either optically bypassing them or being electronically stored and forwarded by them—on its way to the destination, resulting in a decreased bandwidth efficiency compared to other network topologies, for example, single-hop star networks, which are presented next.

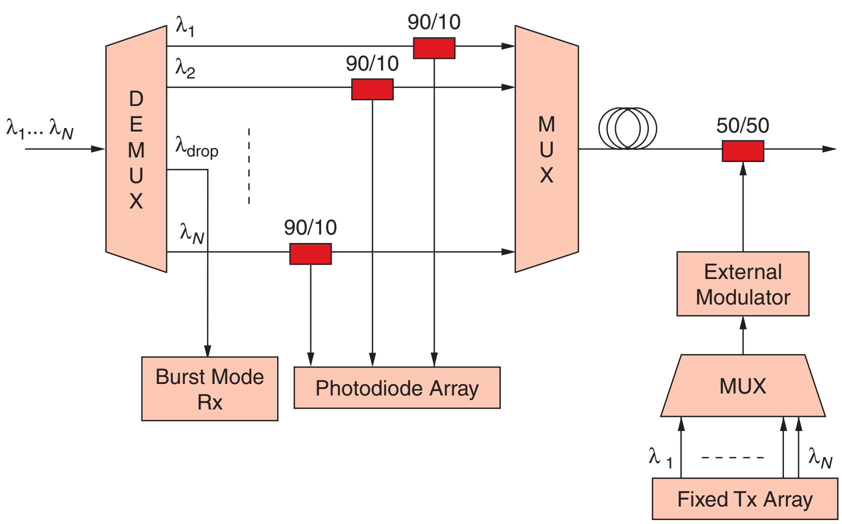

Fig. 2 Example node architecture of a WDM ring network: the signals on drop wavelength λ_{drop} are converted to the electrical domain and the signals on the other wavelengths stay in the optical domain

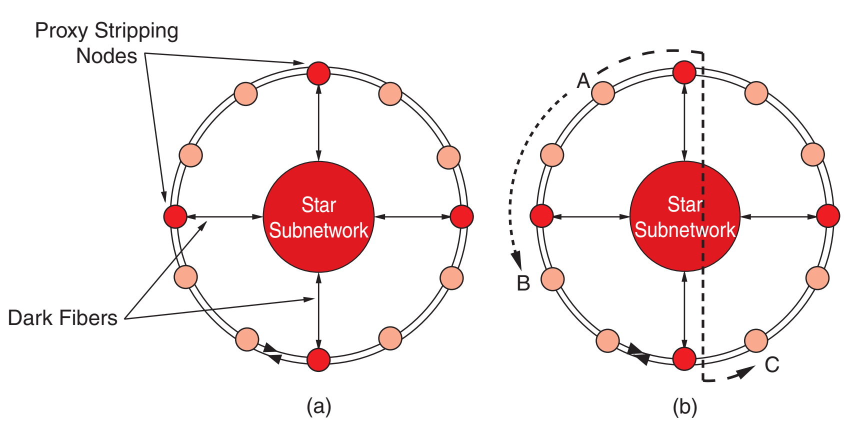

Fig. 3 Proxy stripping technique: (a) architecture: RPR with $N = 12$ nodes, where $P = 4$ nodes are interconnected by a dark-fiber single-hop star subnetwork; (b) operation: packets are sent on the shortest path and stripped by the destination

WDM on star subnetwork

To avoid the above mentioned shortcomings of WDM ring networks, a subset $P \leq N$ of the RPR ring nodes may be interconnected by a single-hop star subnetwork, as shown in Fig. 3(a) for $N = 12$ and $P = 4$. The star subnetwork is best built by using dark fibers, which are abundantly available in the metropolitan area. Dark fibers have recently been installed by most conventional carriers, a growing number of public utility companies, and new network operators who make use of their right-of-ways, especially in metropolitan areas, to install a fiber infrastructure that exceeds their current needs. The fibers that are not needed remain unlit, i.e., dark. These dark fibers provide a cost-effective way to build very high capacity networks or upgrade the capacity of existing (ring) networks. Buying one's own dark fibers is a promising solution to reduce network costs as opposed to leasing bandwidth, which is an ongoing expense. Nodes can be attached to the single-hop star subnetwork one at a time, according to their traffic demands, in a pay-as-you-grow manner.

Nodes attached to the star subnetwork perform proxy stripping and are referred to as *proxy stripping nodes* as illustrated in Fig. 3(b). The basic idea of proxy stripping is to take a packet off the ring and send it over the star subnetwork whenever the star subnetwork provides a shorter path in terms of the number of hops (whereby the transmission from one proxy stripping node through the star subnetwork to another proxy stripping node counts as one hop). If the transmission over the star subnetwork does not provide a smaller hop count, then the packet stays on the ring.

To illustrate the proxy stripping, consider the transmissions from Node A to Node B and from Node A to Node C in Fig. 3(b). First, recall that in RPR spatial reuse is achieved by means of shortest path routing and destination stripping. For the transmission from Node A to Node B, the counterclockwise transmission over the ring has a hop count of three. The transmission over the star subnetwork would also have a hop count of three (one hop in the clockwise direction from Node A to the closest proxy stripping node, plus one hop over the star subnetwork to the proxy stripping node closest to Node B, plus one hop in the counter clockwise direction to Node B). Since proxy stripping does not provide a smaller hop count, the packet stays on the ring. More specifically, the source Node A sends the packets in the counterclockwise direction on the ring (shortest path routing) and the destination Node B takes the packets off the ring (destination stripping). The traversed proxy stripping node performs simple packet forwarding on the ring.

If the shortcuts of the star subnetwork provide a shorter path than either peripheral fiber ring, then the intermediate proxy stripping nodes perform proxy stripping instead of simple packet forwarding. Proxy stripping makes use of RPR's built-in shortest path routing and destination stripping. As shown in Fig. 3(b) for the transmission from Node A and Node C, Node A sends its packets destined for Node C to its closest proxy stripping node (shortest path routing). Now, instead of simply forwarding the packets on the clockwise peripheral ring, the proxy stripping node pulls the packets from the ring and sends them across the single-hop star subnetwork to the proxy stripping node closest to the destination Node C. The receiving proxy-stripping node forwards the packets on the shortest path along the counterclockwise ring towards Node C, which finally takes the packets from the ring (destination stripping). Practically, proxy stripping can be implemented with the help of the topology database, which is built and continuously updated in each node by RPR's built-in topology discovery protocol.

The hub of the single-hop star subnetwork may be a passive star coupler (PSC), an arrayed waveguide grating (AWG), or a combination of both. The PSC is an optical broadcast device where at each input port all incoming wavelengths are broadcast to all output ports. In contrast, the AWG is a wavelength router where at a given input port each incoming wavelength is routed to a different output port. The wavelength routing characteristics of an AWG are such that no wavelength channel collisions occur at the AWG output ports if a given wavelength is used at multiple AWG input ports simultaneously. As a consequence, an AWG allows that all wavelength channels can be spatially reused at all AWG input ports without resulting channel collisions at the AWG output ports. Thus, by using an AWG rather than a PSC, single-hop star WDM networks with a high degree of spatial wavelength reuse can be designed.

Figure 4 depicts a hybrid ring-star network where $P = 4$ of $N = 12$ RPR nodes are interconnected by an AWG-based single-hop star WDM network. The P nodes connected to both the ring and the star subnetwork perform the proxy stripping and are also referred to as *ring-and-star homed nodes*. The AWG has D input ports and D output ports, where $D \geq 1$ determines the spatial reuse factor of the set of wavelengths used in the star subnetwork. To each AWG port S RPR nodes are attached by means of an $S \times 1$ combiner and a $1 \times S$ splitter, where $S \geq 1$ and $D \cdot S = P$. More precisely, each node attached to the star subnetwork is equipped with an additional pair of one tunable transmitter (TT) and one tunable receiver (TR) for transmission and reception on the star subnetwork, respectively. To tap into the vast amount of capacity of the star subnetwork, both the TT and TR are assumed to be tunable over the entire set of wavelengths used in the star WDM subnetwork. (Alternatively, the tunable transceivers may be replaced with arrays of multiple fixed-tuned transceivers according to technological feasibility issues, cost considerations, and traffic demands.) Each TT and TR is connected to a separate combiner input port and splitter output port, respectively. The wavelength channel access on the star subnetwork may be governed by using an appropriate multichannel medium access control (MAC) protocol. If necessary, optical amplifiers are deployed to compensate for insertion and propagation losses in the star subnetwork.

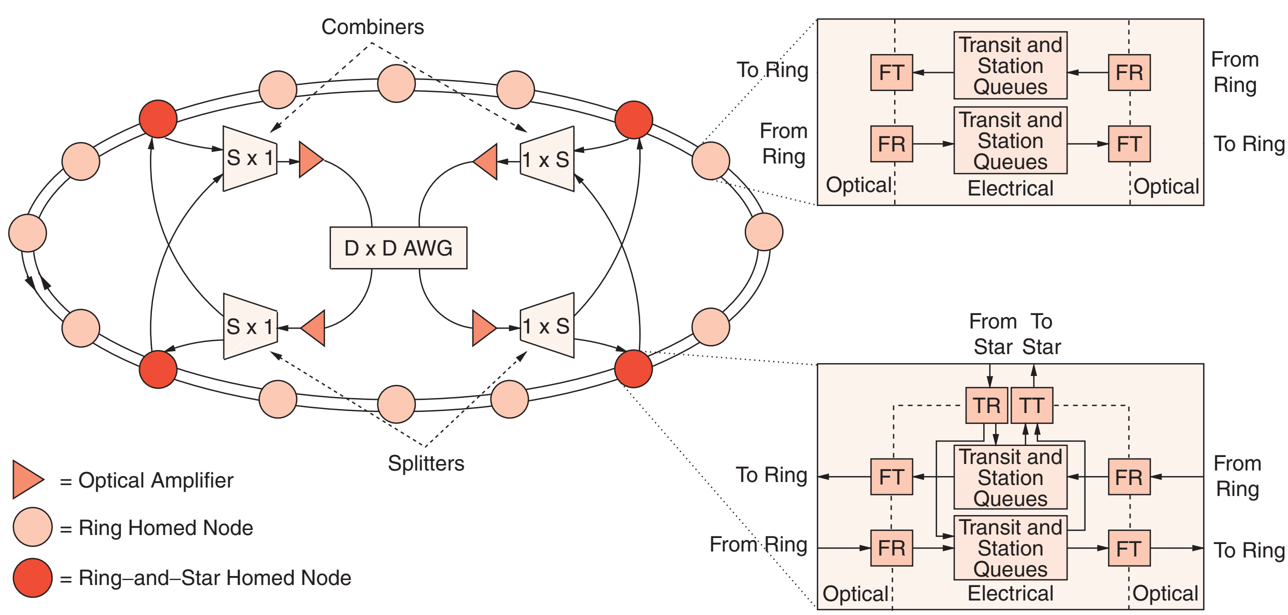

Fig. 4 WDM on star subnetwork: a subset $P = 4$ of $N = 12$ RPR nodes are interconnected by an AWG-based single-hop star WDM subnetwork

Deploying WDM on an AWG-based star subnetwork has several advantages compared to WDM ring networks. Apart from the high spatial wavelength reuse factor of the underlying AWG, the spatial reuse factor on the ring network is also increased. To see this, note that by attaching an increasing number of P RPR nodes to the star subnetwork, more nodes are able to perform proxy stripping and benefit from the shortcuts and large capacity of the star WDM subnetwork. As a result, packet transmissions are restricted to smaller segments on the ring, and fewer ring bandwidth resources are utilized, leading to an increased spatial reuse factor. Furthermore, the hybrid ring-star network provides resilience against multiple link and/or node failures. Specifically, each ring segment between two adjacent proxy stripping nodes is able to fully recover from a single link or node failure by using the star subnetwork to bypass the failed ring segment. Thus, the number P of nodes attached to the star subnetwork determines the number of fully recoverable link and/or node failures, provided that in each ring segment no more than one failure occurs. And finally, note that the hybrid ring-star network provides a higher degree of connectivity than bidirectional ring networks. Specifically, in the hybrid ring-and-star network, each ring-and-star homed node is connected to $P-1$ other ring-and-star homed nodes through the star subnetwork, in addition to being connected to two ring homed nodes through the ring network. Consequently, packets can be sent along multiple alternative short-cut routes, which help decrease path lengths and improve load balancing. In particular, hot-spot nodes, i.e., nodes that generate or receive a lot of traffic, should be made ring-and-star homed nodes so that they can send/receive their traffic directly across the single-hop short-cuts of the star subnetwork rather than traversing multiple intermediate nodes on the peripheral ring.

Hybrid WDM upgrades of RPR

A future direction for WDM upgrading single-channel RPR networks are hybrid WDM upgrades that deploy WDM on both the peripheral ring and the star subnetwork. A promising approach is to interconnect RPR nodes with the highest traffic demands by means of a single-hop star WDM network, each node possibly with a different node structure. For instance, it appears reasonable to equip hot-spot nodes with an array of multiple transceivers while the remaining nodes attached to the star subnetwork use a single (tunable) transceiver. Moreover, operating a few additional wavelength channels on the peripheral fiber ring enables the efficient transport of multicast and broadcast traffic on the ring. Note that this approach would require to upgrade all ring nodes to WDM. However, the costs are reduced because only a few wavelengths are used on the ring, and, therefore, each node is required to operate only on a few wavelength channels.

Read more about it

• F. Davik, M. Yilmaz, S. Gjessing, and N. Uzun, “IEEE 802.17 resilient packet ring tutorial,” *IEEE Commun. Mag.*, vol. 42, no. 3, pp. 112–118, Mar. 2004.

• P. Yuan, V. Gambiroza, and E. Knightly, “The IEEE 802.17 media access protocol for high-speed metropolitan-area resilient packet rings,” *IEEE Network*, vol. 18, no. 3, pp. 8–15, May/ June 2004.

• A. Carena, V. De Feo, J.M. Finochietto, R. Gaudino, F. Neri, C. Piglione, and P. Poggiolini, “RingO: An experimental WDM optical packet network for metro applications,” *IEEE J. Select. Areas Commun.*, vol. 22, no. 8, pp. 1561–1571, Oct. 2004.

• M. Herzog, M. Maier, and M. Reisslein, “Metropolitan area packet-switched WDM networks: A survey on ring systems,” *IEEE Commun. Surveys Tutorials*, vol. 6, no. 2, pp. 2–20, 2004.

• M. Maier and M. Reisslein, “AWG-based metro WDM networking,” *IEEE Commun. Mag.*, vol. 42, no. 11, pp. S19–S26, Nov. 2004.

• M. Herzog, M. Maier, and A. Wolisz, “RINGOSTAR: An evolutionary AWG based WDM upgrade of optical ring networks,” *IEEE/OSA J. Lightwave Technol.*, vol. 23, no. 4, Apr. 2005.

About the authors

Martin Maier is an associate professor at the Institut national de la Recherche Scientifique (INRS), Montreal, Canada. He was educated at the Technical University Berlin, Germany, and received the Dipl.-Ing. and Dr.-Ing. degrees (both with distinctions) in 1998 and 2003, respectively. He is author of the book *Metropolitan Area WDM Networks—An AWG Based Approach* (Kluwer, 2003) and can be reached at maier@ieee.org.

Martin Reisslein is an associate professor in the Department of Electrical Engineering at Arizona State University, Tempe. He received his Ph.D. in systems engineering from the University of Pennsylvania, Philadelphia, in 1998. He teaches and conducts research in the areas of wireless multimedia networking, video traffic characteristics, optical networking, and engineering education. He is editor-in-chief of *IEEE Communications Surveys and Tutorials* and can be reached at reisslein@asu.edu.