
METROPOLITAN AREA PACKET-SWITCHED WDM NETWORKS: A SURVEY ON RING SYSTEMS

MARTIN HERZOG, TECHNICAL UNIVERSITY BERLIN

MARTIN MAIER, CENTRE TECNOLÓGIC DE TELECOMUNICACIONS DE CATALUNYA (CTTC)

MARTIN REISSLEIN, ARIZONA STATE UNIVERSITY

ABSTRACT

We provide a comprehensive survey of packet-switched ring metro WDM networks. We first review current standardization and testbed activities, then we provide a categorization of ring WDM networks. We structure our survey according to a classification of the medium access control (MAC) protocols employed in the networks. Throughout the article we pay close attention to the key factors that govern the throughput-delay performance of the networks, such as the source vs. destination stripping of the data packets from the ring and the a priori or a posteriori access strategies. We also consider fairness aspects and QoS support in the networks.

From the optical networking perspective, the future Internet may be viewed as a three-level hierarchy consisting of backbone networks, metropolitan area networks, and local access networks, as illustrated in Fig. 1. The backbone networks will provide abundant bandwidth by employing wavelength division multiplexing (WDM) links that are interconnected with reconfigurable optical add-drop multiplexers (OADMs) and optical cross connects [1]. The metropolitan area networks, or metro networks for short, interconnect the backbone networks with the local access networks. The local access networks carry the data from and to the individual users. By employing advanced LAN technologies, such as Gigabit Ethernet (GbE), and broadband access, such as xDSL and cable modems, access networks provide increasing amounts of bandwidth. Most existing metro networks are based on synchronous optical network with synchronous digital hierarchy (SONET/SDH) technology, a circuit-switched networking technology. As will be discussed in more detail, circuit-switched SONET/SDH metro networks carry bursty packet traffic relatively inefficiently, resulting in a

bandwidth bottleneck at the metro level. This bandwidth bottleneck, which is widely referred to as *metro gap*, prevents the high-speed clients and service providers in local access networks from tapping into the vast amounts of bandwidth available in the backbone networks [2, 3].

This metro gap may become more severe as proxy cache servers are more widely deployed in the metro networks. These proxy caches, which are employed to reduce network latency, to balance server load, and to increase content availability [4], may result in an increase of local IP traffic and thus exacerbate the metro gap. This trend may be further intensified by the increased use of cellular phones and handheld devices employing next-generation wireless technologies, such as the Universal Mobile Telecommunication System (UMTS) and high-speed wireless local area networks (WLANs), for Internet services, which will increase the amount of locally maintained content, especially as home appliances, cars, and other electronic devices begin to utilize the metro network [5]. In addition, future peer-to-peer applications where each attached user will also operate as a server (Napster being simply a precursor of this controversial file sharing model) may dramatically increase the amount of intra-metro area traffic.

To address the metro gap, three main developments are underway. Two of these developments, which we will briefly discuss later, are:

This article was recommended for publication after undergoing the standard IEEE Communications Surveys and Tutorials review process, which was managed by John N. Daigle, Associate EiC.

- To enhance and adapt circuit-switched SONET/SDH technology to more efficiently support bursty packet traffic.
- To develop a standard for a packet-switched ring network, the *resilient packet ring (RPR)*.

The other main development, which is the focus of this survey, is the design of packet-switched optical WDM networks for the metropolitan area [3, 6]. The designed packet-switched WDM metro networks fall into two main categories:

- Networks with a physical star topology.
- Networks with a physical ring topology.

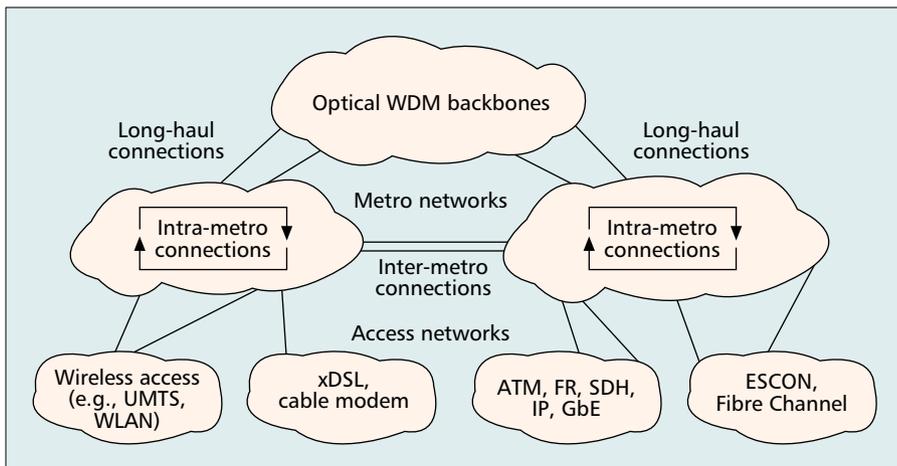
For completeness we mention that there are also bus networks, for example, AMTRAC [7], which have received relatively less interest. Star metro networks come in two flavors:

- Broadcast-and-select single-hop networks based on a wavelength-insensitive passive star coupler (PSC).
- Wavelength-routing networks based on a wavelength-selective arrayed-waveguide grating (AWG).

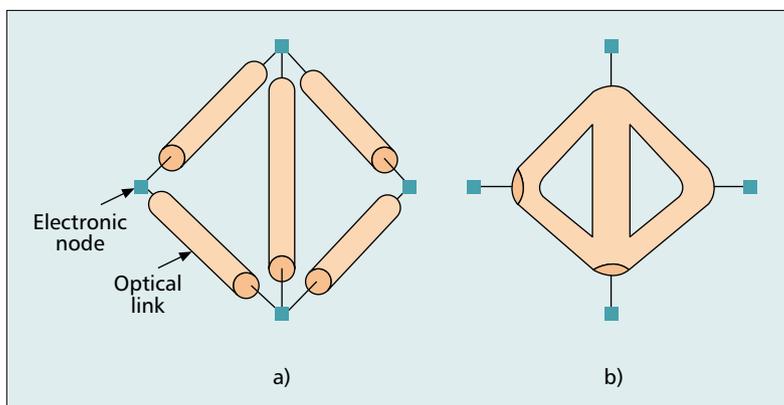
For extensive surveys on physical star networks, the interested reader is referred to [8–11]. In this article, we focus on optical ring metro WDM networks.

We note that packet-switched WDM ring networks have been reviewed to some extent in [12, 13]. However, to the best of our knowledge a comprehensive up-to-date survey of optical metro WDM network architectures and medium access control (MAC) protocols is missing. In our survey we concentrate on *physical* ring networks, which can also be used as building blocks for the design of WDM networks using a multiple ring approach [14]. For information on *logical* rings embedded on a physical AWG-based star network the interested reader is referred to [15, 16].

This survey is organized as follows. We first give a brief historical overview of optical networking in the metropolitan area. This historical overview concludes with a brief survey of current standardization activities on optical metro WDM networks. We provide a survey of a few selected experimental metro ring WDM testbed systems. The purpose of the discussion of these testbeds is twofold. First, it gives an illustration of the network architectures that are feasible with current optical equipment. Second, we introduce and explain several key photonic hardware components used in optical networks. In the main section of this survey, we provide a comprehensive survey of the packet-switched ring metro WDM networks that have been studied to date. In this section, we first introduce a categorization for ring networks. Our categorization is based on the medium access control (MAC) protocol employed in the network. We then comprehensively survey the ring WDM networks within the structure provided by our categorization. We comprehensively survey fairness control and QoS support for packet-switched metro WDM ring networks. We summarize the research and development work on packet-switched ring metro WDM networks to date and outline directions for future research and development.



■ **FIGURE 1.** Network hierarchy: Metro networks interconnect the local access networks, which employ a wide variety of networking technologies and provide increasingly high-speed service, with the high-speed optical backbone networks.



■ **FIGURE 2.** Optical metro networks: a) first generation, b) second generation.

HISTORICAL OVERVIEW AND STANDARDIZATION ACTIVITIES

Optical fiber is widely considered the medium of choice to provide enough bandwidth in the metro area to the ever increasing number of users and bandwidth-hungry applications, for example, video conferences, distributed games, visualization, supercomputer interconnection, or medical imaging applications that do not allow for image compression [17].

There are two generations of optical metro networks. As depicted in Fig. 2a, in first-generation optical metro networks copper links are replaced with fiber links while the nodes at either end of the fiber remain electronic. In such *opaque* optical networks optical-electronic-optical (OEO) conversions of the signal take place at each node [18]. Initially, each fiber carried only one wavelength such as in Fibre Distributed Data Interface (FDDI) and IEEE 802.6 Distributed Queue Dual Bus (DQDB) networks. To cope with the increasing amount of data traffic and to fully exploit the gain bandwidth of the optical Erbium doped fiber amplifier (EDFA), *wavelength division multiplexing* (WDM) was introduced in the 1990s. With WDM, each fiber carries multiple wavelength channels, each operating at any arbitrary line rate, for example, electronic peak rate. After providing these huge pipes, attention turned from optical transmission to optical networking [19]. As shown in Fig. 2b, in second-generation optical networks, OEO conversions occur only at the source and destination nodes while all of the intermediate nodes are *optically bypassed*

by means of optical add-drop multiplexers (OADMs). The OADMs allow nodes to locally drop and add one or more wavelengths from or to an incoming or outgoing fiber link. By optically bypassing nodes, the electro-optic bottleneck is alleviated, and the number of electronic port cards can be reduced at each node, resulting in all-optical (OOO) node structures and significantly reduced network costs, which is one of the most important drivers for optics [20].

Optical bypasses can be used in ring WDM networks to build cost-effective node architectures [21] and to reduce the number of logical intermediate nodes between source-destination pairs, leading to a decreased logical mean hop distance [22]. The resulting all-optical lightpaths are able to provide *transparent* channels to users who are free to choose bit rate, modulation format, and protocol. This transparency enables the support of various legacy as well as future services, which may include Asynchronous Transfer Mode (ATM), Frame Relay (FR), Synchronous Digital Hierarchy (SDH), the Internet Protocol (IP), Enterprise System Connection (ESCON), and Fibre Channel, as illustrated in Fig. 1. We note that there are also hybrid forms of optical networks where not all intermediate nodes are optically bypassed and OEO conversion takes place not only at the source and destination nodes but also at a few selected intermediate nodes. This type of optical network is known as a *translucent* network.

Today's metropolitan area networks are mostly SONET/SDH ring networks. These networks are circuit-switched networks. The individual network nodes access the network bandwidth in a time-division multiplex fashion, that is, each node is periodically allocated a specific number of slots. The SONET/SDH technology may be combined with WDM to establish multiple SONET/SDH rings on one fiber. Also, SONET/SDH WDM rings may employ optical bypassing and traffic grooming to alleviate the computational burden and reduce the number of electronic port cards at bypassed nodes [23]. (Traffic grooming refers here to the routing of traffic destined to a node on the wavelengths that are not bypassed at the node. In general, traffic grooming in WDM networks aims at collecting lower-rate traffic and sending it on high-speed wavelength channels such that a smaller number of wavelengths is required and fewer wavelengths have to be dropped and electronically processed at each node.) The main drawback of SONET/SDH networks is that due to their time-division multiplex operation in conjunction with a circuit set-up time on the order of several weeks or months [24], they accommodate packet traffic only inefficiently [25], especially when the traffic is highly variable, giving rise to the so called metro gap described above [26, p. 14–16]. The metro gap is exacerbated by a number of additional drawbacks of SONET/SDH for which we refer the interested reader to [24] for details. We only note here one additional major drawback, namely that SONET/SDH is designed for symmetric traffic, which leads to inefficiencies when transporting asymmetric IP traffic.

The inefficiencies of SONET/SDH networks are addressed by three new technologies, the *data over SONET/SDH (DoS)* technologies: the Generic Framing Procedure (GFP) [27], Virtual Concatenation [28], and the Link Capacity Adjustment Scheme (LCAS) [29] currently being standardized by the International Telecommunication Union-Telecommunication Standardization Sector (ITU-T) and T1X1.5.

The GFP technology allows for the transport of data packets in SONET/SDH frames. Until now many network operators used proprietary technologies based on packet over SONET/SDH (PoS) for this purpose. With PoS the boundaries of the variable-size data packets are marked with con-

trol characters which requires the receivers to have lots of processing capacity since each incoming byte has to be monitored to recognize the boundaries. In addition, occurrences of the control character in the data packet have to be masked with byte stuffing, resulting in a fluctuating data rate depending on the content of the packet. In contrast, with Frame-Mapped GFP (GFP-F) [30] each data packet is preceded by a short header providing the length of the packet so that the receiver knows the beginning of the next packet in advance and no byte stuffing is required. The header is protected with a checksum that corrects single-bit errors. For storage networks, Transparent GFP (GFP-T) provides a method to transparently transport block-coded data, such as 8B/10B coded bytes, which is bandwidth efficient and introduces only small delays. 8B/10B coding, in which ten bits are transmitted for each byte, is common in storage networks and is also used in Gigabit Ethernet (GbE). The two additional bits are used to balance the numbers of ones and zeros and to transmit link control information.

The SONET/SDH technology offers data transmission only at specific rates from a prescribed set of rates. A GbE connection with a data rate of 1 Gb/s, for instance, would have to be transported via SONET/SDH at a data rate of 2.4 Gb/s, resulting in an overhead of 1.4 Gb/s. With Virtual Concatenation data rates of a much finer granularity are provided to reduce the overhead. This is achieved by virtually combining (concatenating) multiple SONET low-data-rate connections into an aggregate connection close to the desired data rate. The individual connections making up an aggregate connection can operate at different data rates and can travel on different paths through the network. Alignment is performed at the receiver. When LCAS is added, further flexibility is obtained in that the aggregate data rate can be adapted to the data rate currently required. For instance, the amount of data transported over the GbE connection might differ significantly at different times of day. To adapt the data rate, low-rate tributaries can be added or removed from the virtually concatenated connection. To add or remove connections, control packets are exchanged between the sender and the receiver. Note that both virtual concatenation and LCAS do not require any changes inside the SONET/SDH network, only the sender and the receiver are affected. In conjunction with control plane protocols such as Generalized Multiprotocol Label Switching (GMPLS) or Automatic Switched Transport Networks (ASTNs), DoS enables SONET/SDH-based networks to automatically adapt to the current traffic situation within seconds or minutes. This may be sufficient to achieve a high utilization in backbone networks where the traffic flows are aggregates of many individual flows and are thus relatively smooth. In metro networks, however, the traffic is more bursty and it is desired to efficiently share the available capacity between the nodes at the time scale of individual packets (packet switching) or bursts of packets (burst switching).

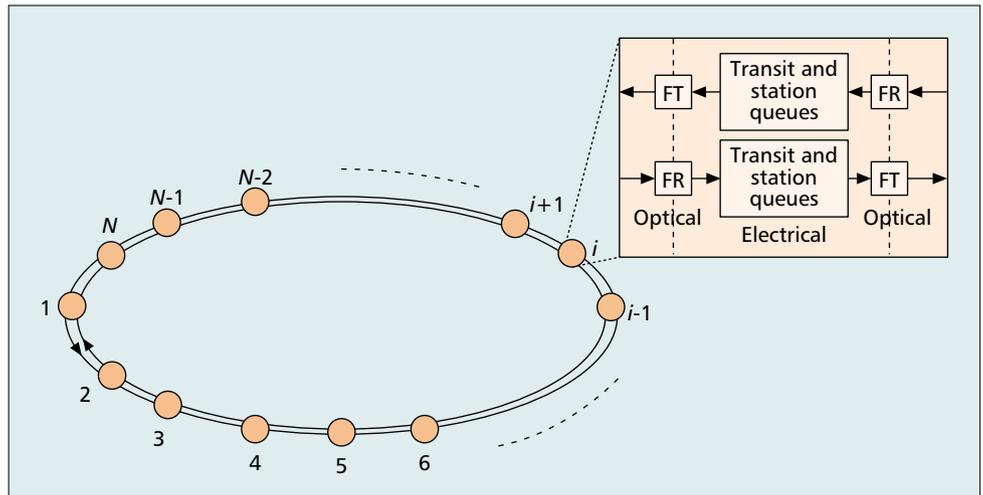
While the standardization efforts in the area of SONET/SDH are not specific to metro networks, the importance of the metro gap is reflected by the large number of recently initiated standardization activities and industry fora such as the Internet Engineering Task Force (IETF) working group (WG) for IP over Resilient Packet Ring (IPoRPR), the IEEE 802.17 resilient packet ring working group (RPRWG), the Metro Ethernet Forum (MEF), and the resilient packet ring (RPR) alliance, which comprises more than 70 companies.

At present the IETF WG IPoRPR and the IEEE 802.17 RPRWG are working on the new resilient packet ring (RPR) standard for packet-switched metro ring networks, which was

anticipated to be completed by the time of this publication. The RPR network consists of a bidirectional dual-fiber ring using one wavelength for each direction, which is OEO converted at each node, that is, RPR operates without WDM. The counter-rotating rings provide protection against any single link or node failure. The RPR network and node architecture is shown in greater detail in Fig. 3. Every node is equipped with two fixed-tuned transmitters (FTs) and two fixed-tuned receivers (FRs), one for each fiber ring. RPR is an example of a *buffer insertion ring* where each node features three different types of electronic FIFO buffers or queues: reception, transmission, and insertion [31]. In general, the reception and transmission buffers store packets that are destined to or originate from the corresponding node. The insertion buffer temporarily stores the incoming ring traffic in the electrical domain in order to allow the local node to transmit a packet onto the ring. Specifically, in RPR each node has separate transit (insertion) and station (transmission, reception) queues for either ring, as depicted in Fig. 3. For each ring a node has one or two transit queues, one transmission queue called the stage queue, one reception queue, and one add_MAC queue which stores control packets generated by the local node. The nodes in the RPR network operate in one of two modes: single-queue mode or dual-queue mode. In the single-queue mode, the transit path consists of a single FIFO queue called the primary transit queue (PTQ). If the PTQ is not full, highest priority is given to add_MAC traffic. In the absence of local control traffic, priority is given to in-transit ring traffic over station traffic. In the dual-queue mode, the transit path comprises two queues, one for guaranteed class A traffic (PTQ) and one secondary transit queue (STQ) for class B (committed rate) and class C (best-effort) traffic. In the dual-queue mode, if both the PTQ and the STQ are not full, highest priority is given to add_MAC traffic, similar to the single-queue mode. If there is no local control traffic, PTQ traffic is always served first. If the PTQ is empty, the local transmission queue (stage queue) is served until the STQ reaches a prescribed queue threshold. If the STQ reaches that threshold, STQ in-transit ring traffic is given priority over station traffic such that in-transit packets are not lost due to buffer overflow. Thus, the transit path is lossless and a packet put on the ring is not dropped at downstream nodes. The dual-queue operation mode enables service differentiation.

The RPR network design makes use of the following four underlying principles:

- **Source stripping:** With source stripping the source node takes the transmitted packet from the ring.
- **Destination stripping:** With destination stripping, packets are removed from the ring by the receiving node rather than the transmitting node.
- **Spatial reuse:** As opposed to source stripping, destination stripping enables the destination stripping node and its downstream neighbor nodes to spatially reuse bandwidth on the ring, resulting in a higher degree of concurrency and an increased network capacity.
- **Shortest path routing:** With shortest path routing a given source node transmits packets to a destination node via



■ FIGURE 3. Resilient packet ring (RPR) network and node architecture connecting N nodes.

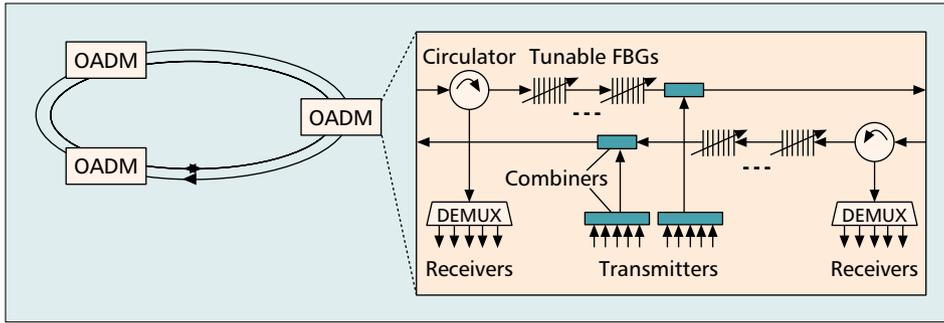
the shortest path (for example, given in terms of number of hops or distance) by using the appropriate ring. Spatial reuse and shortest-path steering are well understood, and it was shown within the MetaRing project that their use increases network capacity significantly [32, 33]. The RPR standard also defines fairness control algorithms that allow a congested downstream node to throttle the transmission rate of upstream nodes by sending fairness control packets upstream. For more details on RPR's fairness control the interested reader is referred to [34]. The two main limitations of RPR are:

- The use of only one wavelength in each fiber.
- The OEO conversion of all traffic at each node, that is, the fact that RPR belongs to the family of first-generation (opaque) networks.

WDM ring networks overcome these limitations by using multiple wavelengths in a fiber and optically bypassing transit traffic.

Recently, research and standardization work on Ethernet passive optical networks (EPONs), which are designed to carry Ethernet frames at standard Ethernet rates, have received considerable attention [35]. The EPON is a point-to-multipoint optical network with no active elements in the signal path from source to destination. The only interior elements used in an EPON are passive optical components, such as optical fiber, splices (which connect two fibers), and splitters, which fan out to multiple optical drop fibers connected to subscriber nodes. An EPON is an optical broadcast network, possibly augmented with a wavelength-routing WDM overlay. There are several EPON topologies suitable for the access network. Typically, EPONs have a tree topology, but other topologies such as ring, tree-and-branch, and bus are also possible. An EPON carries all data encapsulated in Ethernet frames. In addition to the standardization efforts, research on the design and evaluation of efficient multiple access schemes for EPONs have begun recently [36–38].

Newly adopted quality of service (QoS) techniques have made Ethernet networks capable of supporting voice, data, and video. These techniques include full-duplex support, prioritization (p802.1p), and virtual LAN (VLAN) tagging (P802.1Q). The standards work for Ethernet in the local subscriber access network is currently being done in the IEEE P802.3ah Ethernet in the First Mile (EFM) Task Force. Ultimately, the optical Ethernet has the potential to evolve from a pure LAN technology to a metropolitan area network technology that some predict will replace SONET, ATM, and Frame Relay [39].



■ FIGURE 4. KomNet metro WDM network.

EXPERIMENTAL SYSTEMS

In this section we survey three of the most recent experimental testbed systems for packet-switched ring metro WDM networks: KomNet, RINGO, and HORNET. The surveyed testbed systems illustrate the capabilities of currently readily available photonic hardware components. By way of explaining the functioning of these testbeds we also explain the functionalities of several key photonic networking components.

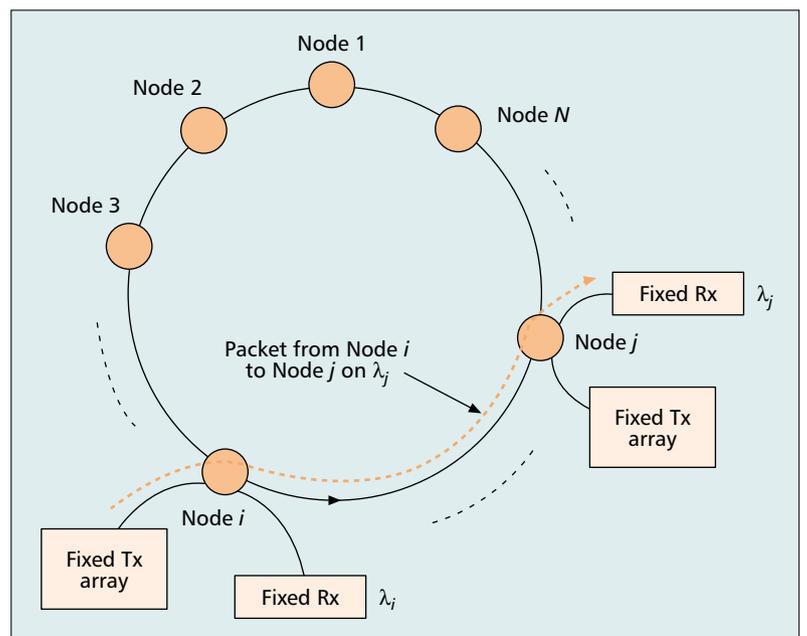
Most of the experimental ring metro WDM networks surveyed in this section operate at a line rate of 2.5 Gb/s. Depending on the used technology, the systems are suited either for circuit or packet switching. While transmitters have been demonstrated to be tunable across adjacent wavelengths in a few nanoseconds, fast tunable receivers are not yet mature. Therefore, most of the experimental packet-switched WDM ring networks use fixed-tuned receivers rather than tunable receivers.

KOMNET

The KomNet metro WDM field trial network consists of three optical add-drop multiplexers (OADMs) interconnected in a bidirectional fiber ring topology [40, 41]. The structure of an OADM is shown in detail in Fig. 4. On each fiber, multiple wavelengths can be dropped by deploying tunable fiber Bragg gratings (FBGs). The FBGs reflect the desired wavelengths back to the circulator, which takes them off the ring and forwards them to the demultiplexer. By using wavelength-insensitive combiners, multiple wavelengths can be added to each fiber. Each FBG has a relatively small insertion loss of 0.1 dB. The FBGs can be mechanically tuned within the millisecond range. Therefore, KomNet is well suited for (λ) circuit switching, but is inefficient for packet switching due to the relatively large tuning time of each FBG.

RINGO

The packet-switched RING Optical network (RINGO) has a unidirectional fiber ring network architecture [42, 43]. It features N nodes, where N also equals the number of wavelengths. Each node is equipped with an array of fixed-tuned transmitters and one fixed-tuned receiver operating on a given wavelength that identifies the node. That is, node j drops wavelength λ_j from the ring. Thus, in order to communicate with node j , a given node i has to transmit data by using the laser operating on wavelength λ_j , as illustrated in Fig. 5. All wavelengths are slotted with the slot length equal to the transmission time of a fixed-size data packet plus guard time. Each node performs λ -monitoring, that is, checks the state of the wavelength occupation, on a slot-



■ FIGURE 5. RINGO metro WDM network.

by-slot basis to avoid channel collisions. This approach is a multi-channel generalization of the *empty-slot approach*. In the empty-slot approach one bit at the beginning of each slot indicates the state of the corresponding slot, that is, whether the slot is free (empty) or occupied. A monitoring node is only allowed to use empty slots for its transmissions.

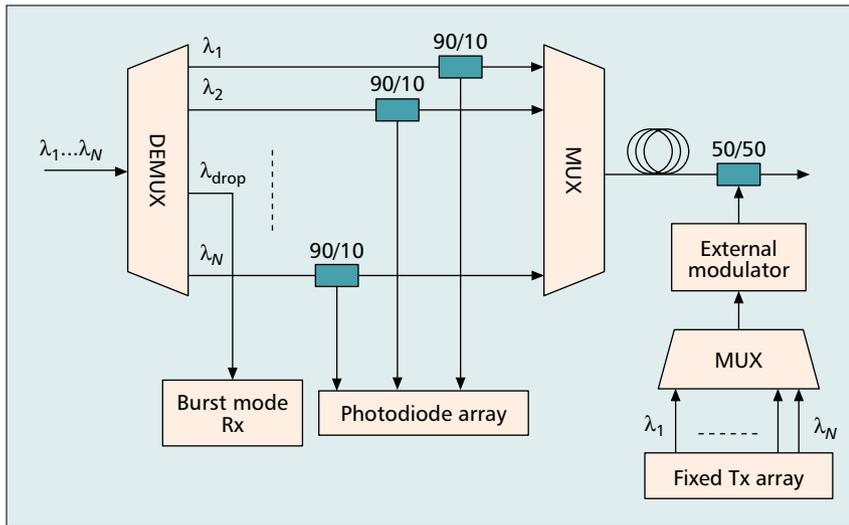
Figure 6 depicts the node structure in greater detail. At each node all wavelengths are demultiplexed.

The drop wavelength is routed to a burst mode receiver while the status of the remaining wavelengths is monitored by using 90/10 taps and an array of photodiodes. A burst mode receiver recovers the clock for each optical burst (packet) very quickly and does not need to receive a continuous signal. A 90/10 tap splits off 10 percent of the optical power from the fiber. Subsequently, the wavelengths are multiplexed on the outgoing ring fiber. With a 50/50 combiner and an external modulator the node is able to send data packets by activating one or more fixed-tuned transmitters. A 50/50 combiner collects signals from two input ports and equally combines them onto one output port. Both input signals thereby experience a combining loss of 3 dB.

HORNET

The Hybrid Optoelectronic Ring NETWORK (HORNET) is a unidirectional WDM ring network [44, 45]. All wavelengths are slotted with the slot length equal to the transmission time of a fixed-size packet (plus guard time). Each wavelength is shared by several nodes for data reception. Every node is equipped with one fast tunable transmitter [46, 47] and one fixed-tuned burst mode receiver [48]. As shown in Fig. 7, the node structure consists of a slot manager, a smart drop, and a smart add module [49].

Access to all wavelengths is governed by means of a *Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA)*



■ FIGURE 6. RINGO node structure.

MAC protocol [50, 51]. When a node transmits a packet it multiplexes a subcarrier tone onto the packet at a sub-carrier frequency that corresponds to the wavelength on which the packet is sent. The destination address of the packet is modulated onto the sub-carrier multiplexed (SCM) tone using a combination of amplitude shift keying (ASK) and frequency shift keying (FSK). For carrier sensing, the slot manager taps off a small amount of optical power and detects it with one photodiode, as illustrated in Figs. 7 and 8. The payload data from all wavelengths collide at baseband while the SCM tones remain intact. The composite SCM signal is demultiplexed into the individual SCM tones using a collection of bandpass filters. The SCM tone corresponding to the drop wavelength of the node is FSK-demodulated while the other SCM tones are ASK-demodulated. The outcome of the ASK demodulation indicates the absence or presence of a packet on the corresponding wavelength. This allows the node to determine whether a wavelength is free for a packet transmission, which is conducted with the smart add module. The outcome of the FSK demodulation indicates whether there is a packet on the node's drop wavelength. If there is a packet, it is taken off the ring with the node's burst mode receiver. The outcome of the FSK demodulation also gives the destination address of the packet. If the destination address does not match the node's address, then the node forwards the packet using its smart add module.

WDM RINGS AND ACCESS PROTOCOLS

In this section, we provide a comprehensive survey of packet-switched ring metro WDM networks. We first discuss a generic WDM ring network architecture from which essentially all studied architectures can be derived with a few modifications. We also introduce a classification of the networks based on the employed MAC protocol. We then survey the networks in the individual categories of our classification.

Most packet-switched ring WDM networks are based on a unidirectional all-optical fiber ring, as shown in Fig. 9. At each node an optical add-drop-multiplexer (OADM) drops a prescribed wavelength from the ring and allows the addition of data at any arbitrary wavelength. A node transmits data on

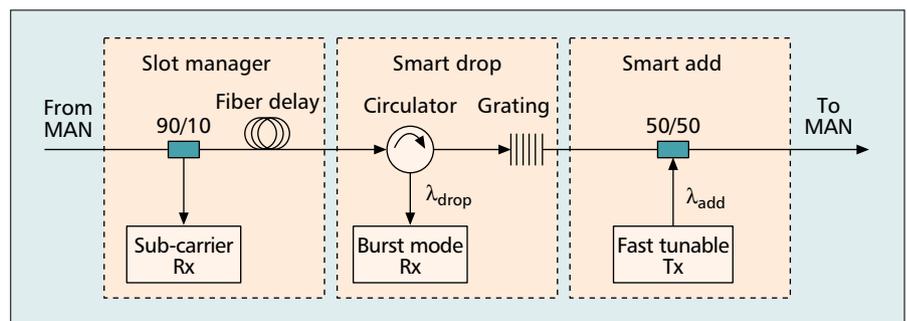
the added wavelength while it receives data on the dropped wavelength. Data on the dropped wavelength are removed from the ring and optically-electronically converted. If the number of nodes N is equal to the number of wavelengths W , as depicted in Fig. 9 for $N = W = 4$, each node has a dedicated *home channel* for reception. However, in general $N \geq W$ since the number of available wavelengths is limited, for example, for cost reasons or finite transceiver tuning ranges. With $N \geq W$ the system is referred to as *scalable* since the number of nodes is independent of the number of available wavelengths.

Each node is equipped with one or more fixed-tuned and/or tunable transmitters and receivers. We adopt the $FT^i\text{-}TT^j\text{-}FR^m\text{-}TR^n$ notation to describe the node architecture, where $i, j, m, n \geq 0$ [8]. That

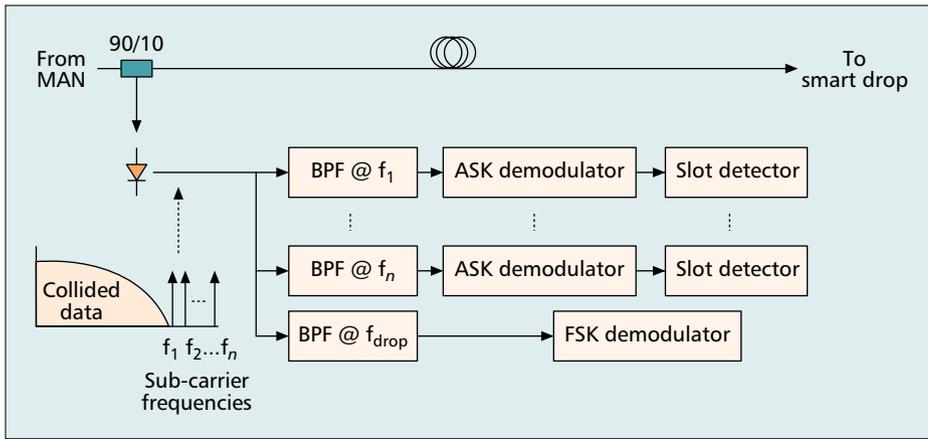
is, each node is equipped with i fixed-tuned transmitters, j tunable transmitters, m fixed-tuned receivers, and n tunable receivers. For example, a $TT\text{-}FR$ node structure means that each node has one tunable transmitter and one fixed-tuned receiver.

When a node inserts a packet on a given wavelength while another packet is currently passing the ring on the same wavelength, a *channel collision* occurs and both packets are disrupted. *Receiver collisions* can also occur with tunable receivers. Receiver collisions are also known as destination conflicts, and they can occur when a node's receiver is not tuned to the wavelength of an incoming packet. This can happen if the destination node does not know about the transmission or another packet is currently received on a different wavelength. Clearly, both channel and receiver collisions have a detrimental impact on the throughput-delay performance of the network. The degradation of network performance due to channel or receiver collisions can be mitigated or completely avoided at the architecture and/or protocol level. For example, equipping each node with a receiver fixed-tuned to a home channel (either dedicated to a single node or shared by multiple nodes) prevents receiver collisions. Similarly, allocating to each node a separate home channel for transmission avoids channel collision at the expense of scalability. However, in scalable systems, that is, systems with $N \geq W$, each wavelength channel is typically shared by multiple nodes, giving rise to channel collisions. Clearly, medium access control (MAC) protocols are needed to govern access to the wavelength channels and to mitigate or prevent channel (and receiver) collisions.

Packet-switched ring WDM networks can be classified according to a number of different criteria, for example, uni-



■ FIGURE 7. HORNET node structure.



■ FIGURE 8. Structure of the HORNET slot manager.

directional vs. bidirectional rings or dedicated vs. shared protection [52]. We introduce a classification of the networks according to the MAC protocols that they employ. As illustrated in Fig. 10, we introduce the main categories of slotted rings, multitoken rings, and meshed rings. Slotted-ring MAC protocols, in which the time is divided into fixed-length slots, can be further classified into protocols with channel inspection and without channel inspection, and those making use of a separate control channel. Protocols with channel inspection determine the status (empty or occupied) of a slot before sending a packet, whereas protocols without channel inspection do not perform such a check of the slot status. A control channel is an additional wavelength channel that is used exclusively for the transmission of control information and does not carry any payload data. MAC protocols with channel inspection use one of two different access strategies: *a priori* or *a posteriori*. With a *a priori* access the packet to transmit in the upcoming slot is selected before the channel inspection of the slot is completed. This has the advantage that the packet selection can be performed without strict timing constraints. The drawback is that the drop wavelength of the destination of the selected packet may turn out to be occupied in the upcoming slot, in which case the packet cannot be transmitted. Also, if some other wavelength is free in this slot, it is not possible to select a different packet for any such free wavelength, resulting in a potential waste of bandwidth. With a *a posteriori* access, on the other hand, the packet to transmit in an upcoming slot is selected after the inspection of the slot is completed. This has the advantage that only packets whose destination drop wavelength is empty in the slot are considered. The drawback is that the *a posteriori* packet selection needs to be performed under tight timing constraints since there is only a small fiber delay between the slot inspection and the packet transmission into the slot, as illustrated in Figs. 6 and 7.

In multitoken rings the time is not slotted. Instead, on each wavelength channel there is a special control packet, the *token*, that travels around the ring. A given node can hold the token for some time duration governed by the MAC protocol and transmit data packet(s) on the corresponding wavelength while it holds the token.

Finally, a meshed-ring network is a ring network that is augmented by additional fibers that create short-cuts between prescribed nodes on the ring. Although the meshed ring is, strictly speaking, not a “pure” ring network, we include it in our survey for completeness, and because meshed-ring networks are closely related to ring networks.

We now comprehensively survey the packet-switched ring metro WDM networks. Our discussion proceeds from left to right in the classification illustrated in Fig. 10, that is, we

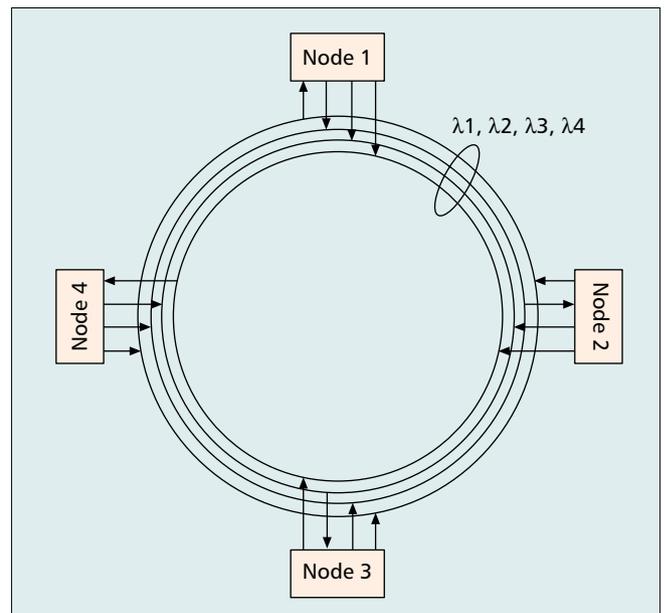
begin with slotted rings without channel inspection and end with meshed rings.

SLOTTED RINGS WITHOUT CHANNEL INSPECTION

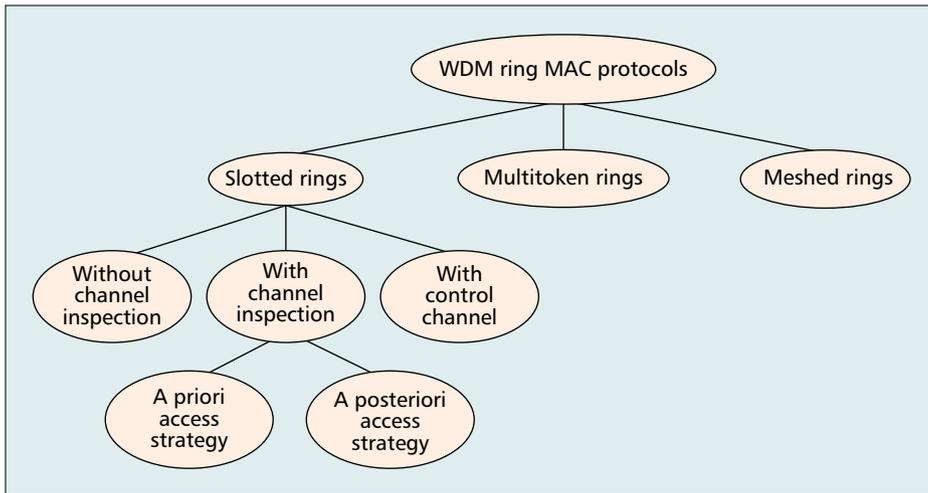
A simple way to avoid channel and receiver collisions is the deployment of time division multiple access (TDMA). Time is divided into slots equal to the packet transmission time. Typically, these time slots are of a fixed size with multiple slots circulating at each wavelength on the ring, as illustrated in Fig. 11. The slots at different wave-

lengths are typically aligned. With TDMA, channel and receiver collisions are avoided by statically assigning each slot to a prescribed source-destination pair. Thus, a fixed amount of capacity is allocated to each pair of nodes which is well suited for uniform regular traffic at medium to high loads, but leads to wasted bandwidth and low channel utilization in the case of bursty traffic.

The only packet-switched network that falls into this category of slotted rings without channel inspection is the Metropolitan Area Wavelength Switched Optical Network (MAWSON) [53–55]. MAWSON is based on a FT^W-FR or alternatively a TT-FR node architecture. N nodes are connected to the ring via OADMs that use FBGs, as discussed earlier, for dropping a different wavelength for reception at each node. In MAWSON, the number of nodes N is equal to the number of wavelengths W and each node has a dedicated home channel, which avoids receiver collisions. In other words, a given wavelength channel interconnects $(N - 1)$ source nodes and one destination node. With the FT^W-FR node structure, broadcasting and multicasting can be achieved by simultaneously turning on multiple lasers, but only unicasting is considered in the evaluation of the MAC protocol.



■ FIGURE 9. Unidirectional ring WDM network with $N = 4$ nodes and $W = 4$ wavelengths.



■ FIGURE 10. Classification of ring WDM network MAC protocols.

Time is divided into fixed-size slots, which are assumed to be aligned across all W wavelengths. Each slot is further subdivided into header and data fields, as shown in Fig. 12. The slots on a given wavelength channel are assigned dynamically on demand. To this end, the header of each slot consists of $(N - 1)$ Request/Allocation (R/A) minislots which are statically preassigned in a TDMA fashion to $(N - 1)$ source nodes. Each R/A minislot essentially consists of two fields, one for requests and one for allocations. More precisely, node i , when ready to send variable-size data packets to node j , uses the request field of its assigned R/A minislot on j 's home wavelength channel to make a request. Upon receipt of node i 's request, node j allocates one or more data minislots to node i by using the allocation field of its assigned R/A minislot on i 's home wavelength. After receiving its allocation, node i transmits the data packet using the allocated data minislots.

To save costs the node architecture and protocol of MAWSON are kept simple, for example, no carrier sensing capabilities are required. Due to in-band signalling no additional control channel and control transceivers are needed. The protocol completely avoids both channel and receiver collisions, achieves good throughput performance, and provides fairness by allocating slots in a round-robin manner. However, the R/A procedure introduces some overhead and additional delay since the request and allocation takes at least one round-trip time around the ring.

SLOTTED RINGS WITH CHANNEL INSPECTION

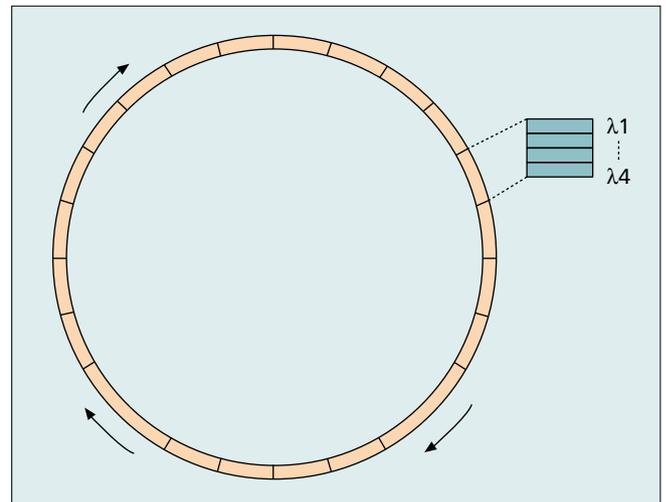
In most slotted WDM rings channel collisions are avoided by enabling the nodes to check the status (used/unused) of each slot. Generally, this is done by tapping off some power from the fiber and delaying the slot while the status of each wavelength is inspected in the tapped-off signal and electronically processed. A packet can then be inserted in a slot at an unused wavelength. Packets waiting for transmission are stored in *virtual output queues* (VOQs). Typically, a node maintains separate VOQs either for each destination or for each wavelength. In the latter case packets arriving at a node from the higher layer are put in the VOQ associated with the drop wavelength of the packet's destination. In WDM ring networks it is typically the responsibility of the MAC protocol to select the appropriate VOQ from which to send a packet in a time slot according to a given access strategy. This can be done a priori, that is, without taking the status of the slots into account, or

a posteriori, that is, with taking the status of the slots into account. In the a priori access strategy each node selects a VOQ prior to inspecting the slot status, whereas in the *a posteriori* strategy each node first checks the status of a slot and then selects an appropriate (non-empty) VOQ.

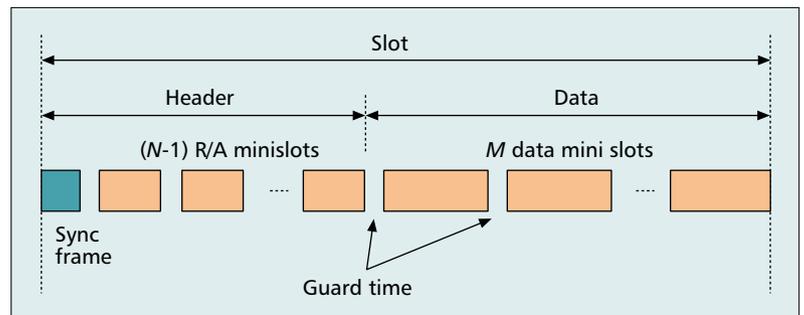
We now describe the networks that fall into the category of slotted rings with channel inspection (Fig. 10). We cover both the networks with a priori access and the networks with a *a posteriori* access in this section, as many networks can be operated with either access strategy.

RINGO — We have already presented the network architecture of the RING Optical network (RINGO), and now we discuss the MAC protocol for RINGO. First, recall from the earlier discussion that RINGO uses a FT^W -FR node architecture. Each node has channel inspection capability built with commercially available components. Nodes execute a multi-channel empty-slot MAC protocol that can also be applied to a TT-FR node architecture.

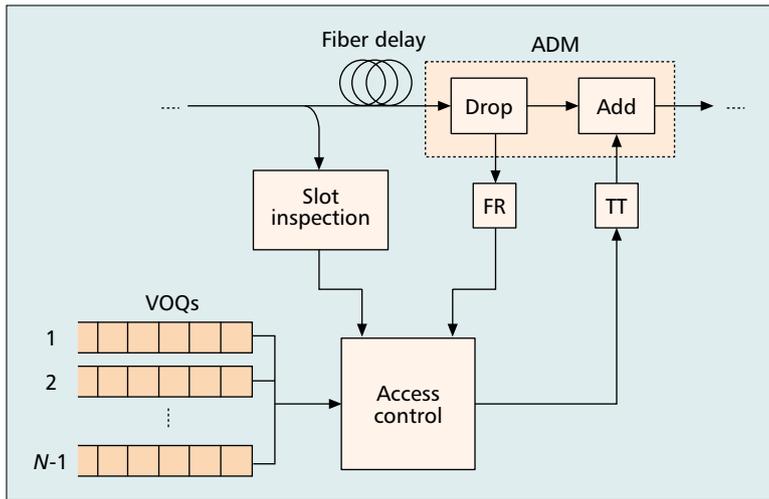
A MAC protocol with a *a posteriori* queue selection has been implemented in the RINGO testbed [56]. The number of wavelengths is assumed to be equal to the number of nodes,



■ FIGURE 11. Slotted unidirectional WDM ring with $W = 4$ wavelengths.



■ FIGURE 12. Slot structure of Request/Allocation Protocol in MAWSON.



■ FIGURE 13. SRR node architecture with VOQs and channel inspection capability.

and each node has one VOQ with first-come-first-served (FCFS) queuing discipline for each wavelength, meaning the oldest packet is sent first. Only the VOQs where the corresponding wavelengths have been found to be empty (unused) are allowed to send data packets in the free time slot. If a TT-FR node architecture is used, only one packet can be sent per time slot and the longest among those queues is chosen.

The overhead of the RINGO empty-slot MAC protocol is very small. To identify the status of a given slot, a single bit is sufficient. All wavelengths are used for data transmission and no separate control channel or control transceivers are required. It was demonstrated that all-optical packet-switched ring WDM networks are feasible with currently available technology. However, owing to the fixed slot size, the transmitted packets have to be of fixed size. Note that variable-size packets can be transmitted in slotted rings without segmentation and reassembly by means of buffer insertion techniques exploiting optical delay lines [57]. More precisely, a transmitting node inserts a sufficiently long optical delay line to delay the in-transit ring traffic until the node has completed the transmission of its (variable-size) data packet. In doing so, the collision of the in-transit ring traffic and the locally injected traffic is avoided.

Synchronous Round Robin (SRR) — SRR is another empty-slot MAC protocol for a unidirectional WDM ring network with fixed-size time slots and destination stripping [58–60]. Each node is equipped with one tunable transmitter and one fixed-tuned receiver (TT-FR), where the transmitter is assumed to be tunable across all W wavelengths on a per-slot basis. If $N = W$ each node has its own home wavelength channel for reception. In the more general case $N > W$ each wavelength is shared by multiple destination nodes [58].

In SRR, each node has $(N - 1)$ separate first-in-first-out (FIFO) VOQs, one for each destination, as shown in Fig. 13. SRR uses an a priori access strategy. Specifically, each node scans the VOQs in a round robin manner on a per-slot basis, looking for a packet to transmit. If such a deterministically selected VOQ is nonempty, the first (oldest) packet is transmitted, provided the current slot was sensed empty. If the selected VOQ is empty the first packet from the longest queue among the remaining VOQs is transmitted, again provided the current slot is unused. If the current slot is occupied, that is, a transmission is not possible as it would result in a channel collision, then no packet is transmitted from the selected VOQ. For the transmission attempt in the next slot, the next VOQ is selected according to the round-robin scan-

ning of SRR. In doing so, under heavy uniform load conditions, when all VOQs are non-empty, the SRR scheduling algorithm converges to round-robin TDMA.

For uniform traffic, SRR asymptotically achieves a bandwidth utilization of 100 percent. However, the presence of unbalanced traffic leads to wasted bandwidth due to the nonzero probability that the a priori access strategy selects a wavelength channel whose slot is occupied while leaving free slots unused. It was shown in [61] that a *posteriori* access strategies avoid this drawback, resulting in an improved throughput-delay performance, albeit at the expense of increased complexity.

SRR achieves good performance requiring only local information on the backlog of the VOQs, which also avoids the well-known head-of-line (HOL) blocking problem. Owing to destination stripping, slots can be spatially reused several times as they propagate along the ring. On the

other hand, slot reuse raises fairness control problems, particularly for nonuniform traffic. A node to which a large amount of slots is directed generates a large amount of free slots, and nodes immediately downstream are in a favorable position with respect to other nodes. We will address this fairness problem later. Note that in order to provide quality of service (QoS), SRR requires additional modifications, which are discussed in greater detail later.

HORNET — The architecture of the Hybrid Optoelectronic Ring Network (HORNET) was described earlier. Recall that HORNET is a unidirectional destination-stripping ring WDM network with a TT-FR network structure. Similar to SRR, to avoid HOL blocking, each node uses VOQs, one for each wavelength, and both a priori and a *posteriori* access strategies can be used. Nodes sense the availability of each slot by monitoring *subcarrier multiplexed* (SCM) tones. The SCM-based carrier sensing scheme requires fewer hardware components than demultiplexing, separately monitoring, and subsequently multiplexing, all wavelengths of the WDM comb, as done in RINGO. More precisely, instead of the demultiplexer, photodiode array, and multiplexer used in RINGO, the HORNET channel inspection scheme requires only a single photodiode.

HORNET's Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) MAC protocol initially assumed fixed-size slots that are well suited for the transport of fixed-size packets, for example, ATM cells [44]. The CSMA/CA MAC protocol can be extended to support variable-size IP packets. Two CSMA/CA MAC protocols, both supporting variable-size packets, are proposed and investigated in [62]. In the first protocol, slots of different sizes circulate along the ring. The slot sizes are chosen according to the predominant IP packet lengths as typically found in traffic measurements. For example, three slots sizes can be chosen such that 40, 552, and 1500 byte long IP packets are accommodated. A dedicated node controls the size and number of slots such that they match the packet size distribution. A variant of this protocol for a TT-FR^W architecture has been proposed in [63].

The second protocol is unslotted and operates similarly to CSMA/CD with collision detection and backoff. More precisely, when a wavelength is sensed idle, a given node begins to transmit a packet. When another packet arrives on the same wavelength before the transmission is complete, the packet transmission is aborted. In this case, the incomplete packet is marked by adding a jamming signal to the end of the packet. Aborted transmissions are retried after some backoff time interval.

A more bandwidth-efficient modification of the second unslotted CSMA/CA protocol was examined in [64]. In the examined Carrier Sense Multiple Access with Collision Preemption (CSMA/CP) protocol, variable-size IP packets do not necessarily have to be transmitted in a single attempt. Instead, packets are allowed to be transmitted and received as fragments that are reassembled at the receiver. Thus, successfully transmitted parts of the original IP packet are not retransmitted, resulting in a higher channel utilization.

Besides demonstrating the feasibility of the SCM-based channel inspection approach, the HORNET project also proved the feasibility of fast tunable transmitters. These allow for replacing arrays of multiple fixed-tuned transmitters with a single tunable transmitter. In HORNET the number of nodes is independent from the number of wavelengths, and it is thus considered scalable. Generally, each wavelength is allowed to be shared by multiple destination nodes with packet forwarding at intermediate nodes, resulting in translucent multihop networks. Note that intermediate nodes not only forward packets toward the destination but also provide signal regeneration in the electrical domain. On the other hand, the CSMA/CA random access protocol does not provide QoS, and the destination stripping gives rise to fairness problems.

Several *a posteriori* buffer selection schemes for the HORNET architecture are studied by Bengi and van As [65, 66]. Recall that in an empty-slot protocol, each unused slot on any wavelength channel can be used for packet transmission by a source node. However, when more than one wavelength channel carries an empty slot in the current slot period, one packet (or equivalently, one VOQ) corresponding to one of the empty channels has to be chosen according to a prescribed selection rule. Due to the short time between channel inspection and packet transmission, the *a posteriori* packet selection process has to be performed at a high speed in the electronic domain, which increases the processing complexity compared to a priori packet selection scheme. Five different *a posteriori* VOQ selection strategies are described and examined in [65]:

- Random Selection: The VOQ from which a packet is to be transmitted is selected randomly according to a uniform distribution.
- Longest Queue Selection: The longest VOQ is chosen upon buffer contention.
- Round-Robin Selection: The VOQ is chosen in a round-robin fashion.
- Maximum Hop Selection: The packet (VOQ) associated with the maximum hop distance between source and destination node is selected when buffer contention arises.
- C-TDMA Selection: The channel-oriented TDMA (C-TDMA) scheme first attempts to select the packet according to a round-robin policy. If that selection would prevent a transmission, either due to an empty VOQ or an occupied slot, then the longest VOQ that allows for a packet transmission is chosen. This scheme is largely equivalent to the SRR scheme with *a posteriori* access

It was found that the random and round-robin buffer selection schemes provide a satisfactory compromise between performance and implementational complexity.

FT-TR Rings — Jelger and Elmirghani [67] proposed a unidirectional empty-slot WDM ring network that uses source stripping. Each node is equipped with one fixed-tuned transmitter and one tunable receiver (FT-TR). Packets are buffered in a single FIFO transmit queue at each node. In the applied source-stripping scheme, a sender must not reuse the slot it just marked empty. It was shown that for source-stripping rings this simple mechanism ensures fairness in that a node

can not starve the entire network. However, the mechanism does not ensure fairness for destination-stripping rings.

The performance of the network was compared for both source and destination stripping in [68]. By means of simulation it was shown that destination stripping clearly outperforms source stripping in terms of throughput, delay, and packet dropping probability.

Clearly, with only one tunable receiver at each node, receiver collisions can occur. Receiver collisions can be avoided in a number of ways. In one approach, arriving packets that find the destination's receiver busy re-circulate on the ring until the receiver of the destination is free, that is, is tuned to the corresponding wavelength [68]. Alternatively, receiver collisions can be completely avoided at the architecture level by replacing each node's tunable receiver with an array of W fixed-tuned receivers, each operating at a different wavelength (FT-FR^W) [13]. Another proposal to resolve receiver contention is based on optical switched delay lines (SDLs) [69]. A destination node puts all simultaneously arriving packets except one into optical delay lines such that packets can be received sequentially.

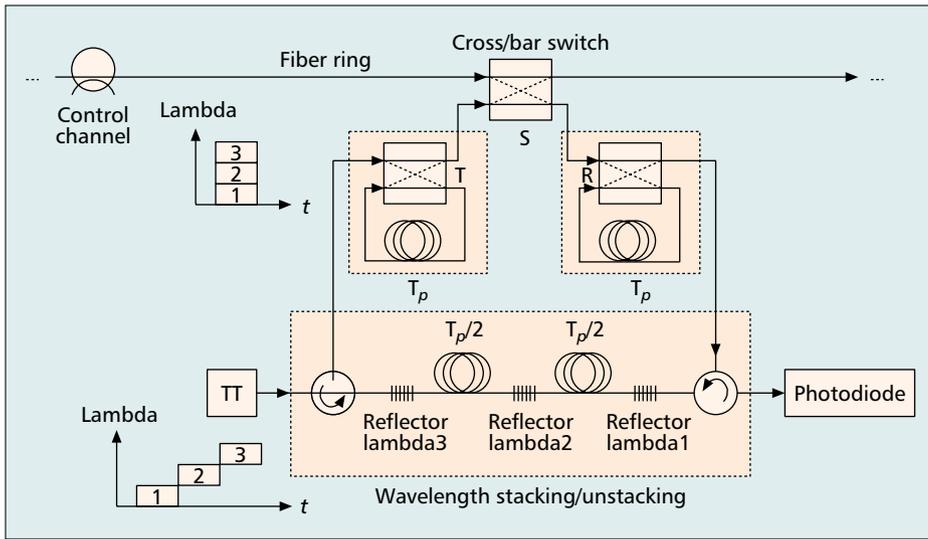
SLOTTED RINGS WITH CONTROL CHANNEL

In slotted ring networks with control channel, the status of the slots is transmitted on a separate control channel (CC) wavelength. Each node is typically equipped with an additional transmitter and receiver, both fixed-tuned to the control wavelength. A separate control channel wavelength enables nodes to exchange control information at high line rates and eases the implementation of enhanced access protocols with fairness control and QoS support, as we will see shortly.

Bidirectional HORNET with SAR-OD — An extended version of the original unidirectional TT-FR HORNET ring architecture in which SCM is replaced with a separate control channel wavelength is investigated in [70]. Transmission on the control channel (and data wavelengths) is divided into fixed-size slots. The control channel carries the wavelength availability information such that nodes are able to “see” one slot into the future. Two counter-directional fiber rings, each carrying W data wavelengths, and an additional control channel wavelength operate in parallel. On each ring every node deploys one fast-tunable transmitter and one fixed-tuned receiver for data, and one transmitter fixed-tuned to the control channel wavelength. Thus, the control channel-based HORNET network is a CC-FT²-TT²-FR⁴ system.

A modified MAC protocol able to efficiently support variable-size packets over the bidirectional ring network was examined. This Segmentation and Reassembly on Demand (SAR-OD) access protocol aims at reducing the number of segmentation and reassembly operations of variable-size packets. Specifically, the transmission of a packet from a given VOQ starts in an empty slot. If the packet is larger than a single slot, the transmission continues until it is complete or the following slot is occupied, that is, the packet is segmented only if required to avoid channel collisions. If a packet has to be segmented, it is marked incomplete and the transmission of the remaining packet segment(s) continues in the next empty slot(s) on the corresponding wavelength. By means of simulation it was shown that SAR-OD reduces segmentation/reassembly overhead by approximately 15 percent compared to a less intelligent approach where all packets larger than one slot are segmented irrespective of the state of successive slots.

The control channel-based bidirectional HORNET ring network preserves the advantages of the original unidirectional



■ FIGURE 14. Node architecture for wavelength stacking.

al HORNET ring, for example, scalability and a small number of hardware components. Bidirectional dual-fiber rings provide an improved fault tolerance against node/fiber failures and survivability compared to unidirectional single-fiber rings [71, 72]. Furthermore, the control channel can also be used to achieve efficient fairness control, as described in greater detail in a later section.

Variable-Size Packets without Segmentation/Reassembly — An access protocol for a control channel-based slotted-ring WDM network that completely avoids segmentation and reassembly of variable-size packets was studied by Bengi in [73, 74]. The access protocol is an extended version of Bengi's original protocol described above. The architecture differs from the control channel-based HORNET in that a unidirectional ring is deployed, and each node uses an additional transmitter fixed-tuned to the node's drop wavelength, resulting in a CC-FT²-TT-FR² system. The additional transmitter is used to forward dropped packets that are destined to downstream nodes that share the same drop wavelength.

The extended MAC protocol relies on a frame-based slot reservation strategy including reservation of successive slots for data packets longer than the given slot size and immediate access for packets shorter than the slot length. Each node is equipped with two VOQs for each wavelength, one for short packets and one for long packets. The ring is subdivided into multiple *reservation frames* with the frame size equal to the largest possible packet length. In these frames, multiple consecutive slots are reserved to transmit long packets without segmentation. A single reservation control packet containing all reservations circulates on the control channel. Each node maintains a table in which the reservations of all nodes are stored. When the control packet passes, a node updates its table and is allowed to make a reservation. The additional fixed-tuned transmitter is used to forward packets concurrently with transmitting long packets within multiple contiguous slots. Besides the support of long packets via reservation, short packets fitting into one slot are accommodated by means of immediate access of empty and unreserved slots.

The proposed protocol provides immediate medium access for packets shorter than one slot and completely avoids the segmentation and reassembly of longer variable-size packets, resulting in reduced complexity. The reservation protocol also enables QoS support, as discussed in greater detail later. On the other hand, the reservation protocol introduces some delay overhead, and reserved slots on their way back to the

source node cannot be spatially reused after destination stripping.

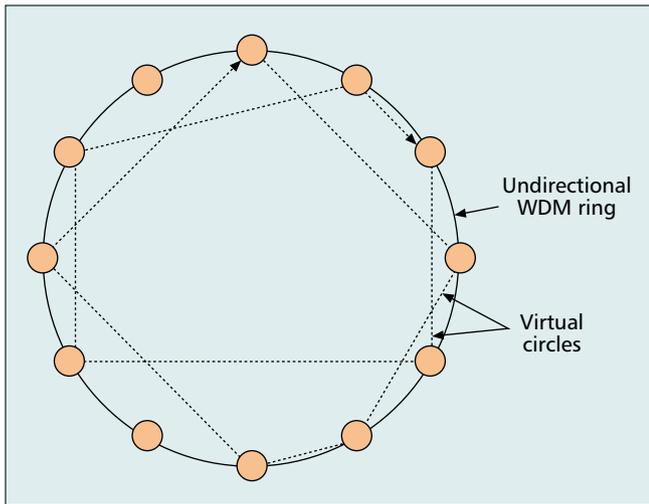
Wavelength Stacking — The *wavelength stacking* technique studied by Smiljanic *et al.* transmits a packet using all wavelengths of a time slot of a control channel-based slotted unidirectional WDM ring [75–77]. Each node is equipped with one fast-tunable transmitter and one photodiode. Time is divided into slots of duration T_p . The length of a data packet is W time slots. A fast-tunable laser at a given node starts transmission W time slots before its scheduled time slot. As illustrated in Fig. 14 for $W = 3$, in each following time slot it transmits data on a different wavelength.

The signal passes through the array of fiber gratings separated by delay lines so that the W segments of the data packet transmitted at different wavelengths are aligned in time. The packet is then transmitted to the network on all wavelengths in parallel by setting switch S to the cross state. On the receiver side, the reverse procedure is performed. A packet is received when switch S is in the cross state and is then unstacked by passing through the same array of fiber gratings and delay lines. Note that a single broadband wavelength-insensitive photodiode without optical filter is sufficient for packet reception since at most one wavelength needs to be converted from the optical to the electronic domain at any given time. The photodiode converts the optical signal into an electrical signal irrespective of the optical carrier frequency (wavelength).

Because wavelength stacking takes W time slots, a node needs to decide in advance when to access the medium. A separate wavelength is used as a control channel for the reservations. Time slots are grouped into cycles of length W slots. Each node may transmit and receive at most one packet within each cycle. The switches T and R in Fig. 14 synchronize wavelength stacking and unstacking. Wavelength stacking is completed in the last time slot of a given cycle and the packet is stored in the delay line by setting T in the cross state. A packet is stored as long as switch T is in the bar state. The packet is transmitted to the network by setting switches T and S in the cross state exactly $2W$ time slots after the reservation. Whenever a node recognizes its address on the control channel, it stores the packet in the delay line by setting switches S and R in the cross state $2W$ time slots after the address notification. The node starts unstacking the packet at the beginning of the next cycle by setting switch R in the cross state. Each node removes a packet that it receives as well as its reservation.

Wavelength stacking/unstacking allows a node to simultaneously send/receive data at different wavelengths in the same time slot despite the fact that the node has only one transceiver. The presented node architecture can be used to realize photonic slot routing (PSR) metro WDM networks, where all wavelengths in a given slot (the *photonic slot*) are switched together rather than separately on a per-wavelength basis [78, 79]. However, the quality of the optical signal may suffer from passing the numerous delay lines and switches in a node.

Virtual Circles with DWADMs — In the unidirectional slotted ring WDM network presented by Cho and Mukherjee in [80], each node is equipped with a *dynamic wavelength add-*



■ **FIGURE 15.** Virtual circles comprising nodes whose DWADMs are tuned to the same wavelength.

drop multiplexer (DWADM). As opposed to tunable transmitters and receivers that can operate independently, the input and output wavelengths of a DWADM must be the same, that is, if the wavelength to receive at a given node s is λ_i , the wavelength to transmit must be the same wavelength λ_i . Furthermore, if the node has to send to another node d , then node d has to use wavelength λ_i to receive and to send a packet. Thus, virtual circles are created (as depicted in Fig. 15) which change dynamically according to varying traffic demands.

The ring network uses W data wavelength channels and a separate control wavelength channel. Nodes communicate over the control channel in a TDM fashion to exchange transmission requests and acknowledgments. The $(W + 1)$ wavelengths are divided into three cycles, which are repeated periodically. In the first cycle a control packet sent by a server node collects transmission requests from all nodes. These are processed by the server node, and wavelength assignments/acknowledgments are sent back to the nodes in the second cycle. In the third cycle each node that has been assigned a wavelength tunes the DWADM appropriately and starts the data transmission.

Owing to their relatively simple structure DWADMs are less expensive than tunable transceivers [80]. However, due to their reduced flexibility wavelength utilization is typically smaller than in TT-TR systems, where transmitters and receivers can be tuned to any arbitrary wavelength independently.

MULTITOKEN RINGS

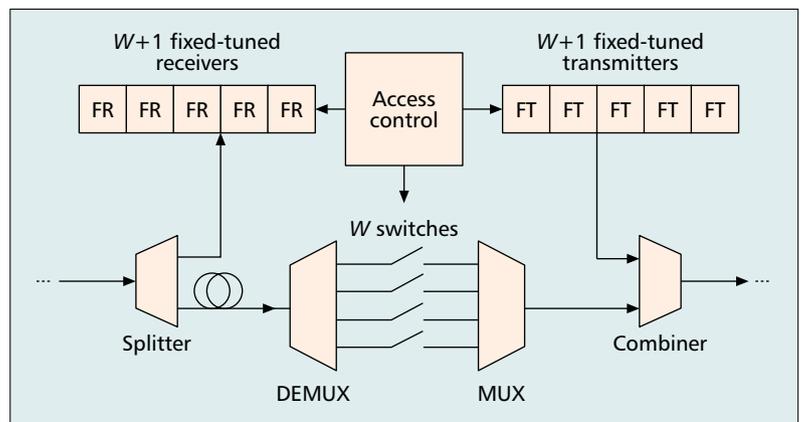
Slotted WDM ring networks have a number of advantages, such as easy synchronization of nodes even at high data rates. Also, they can achieve high channel utilization and low access delay and allow for relatively simple access schemes. However, variable-size packets are difficult to accommodate and, as discussed later, explicit fairness control is needed, which can complicate the medium access significantly. In contrast, variable-size packets can be transported in a reasonably fair manner in token rings where the access is controlled by means of a special control packet — the token — which circulates around the ring. The token is passed downstream from node to node. Each node can hold the token up to a prescribed amount of time, during

which the node is allowed to send (fixed-size or variable-size) packets. Due to the limited token holding time, fairness is achieved. Furthermore, as opposed to slotted rings, nodes do not have to be synchronized. On the other hand, immediate channel access is not possible and the token rotation time (ring propagation time) may decrease the channel utilization efficiency in high-speed optical networks.

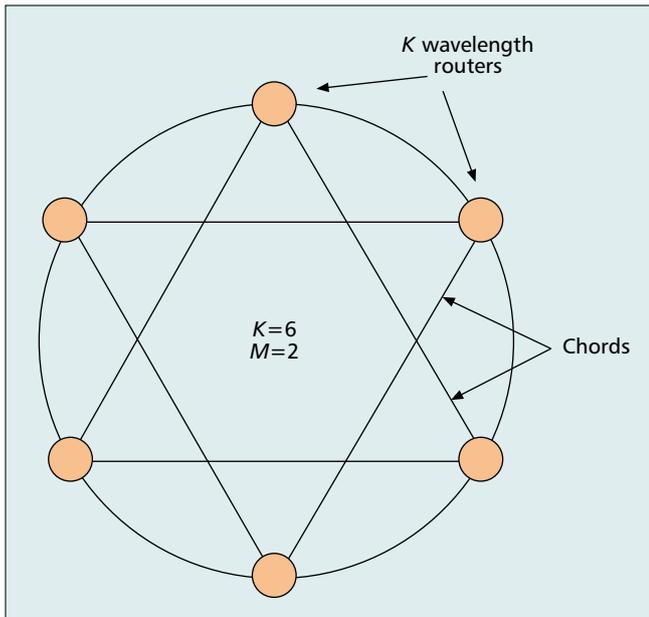
A token-based access scheme for a $CC\text{-}FT^{W+1}\text{-}FR^{W+1}$ unidirectional WDM ring network, the Multitoken Interarrival Time (MTIT) access protocol, was examined in [81, 82]. For each data channel, every node has one fixed-tuned transmitter, one fixed-tuned receiver, and one on-off optical switch, as shown in Fig. 16. The on-off switches are used to control the flow of optical signals through the ring and prevent re-circulation of the same packet on the ring. Once transmitted by the source node, the packet makes one round trip in the ring and is removed from the network by the same source node, that is, MTIT employs source stripping. A separate wavelength is used as the control channel for the purpose of access control and ring management. The optical signal on the control channel is separately handled by an additional fixed-tuned transceiver.

Channel access is regulated by a multitoken approach. Each channel is associated with one token that circulates among the nodes on the control channel and regulates access to the corresponding data channel. The MTIT protocol controls the token holding time by means of a *target token interarrival time* with value $TTIT$. The $TTIT$ is agreed upon by all nodes connected to the ring at the configuration time of the system. The *token interarrival time* $TIAT$ is defined as the time elapsed between two consecutive token arrivals at the node. Upon a token arrival, the node is allowed to hold the token for a period of time equal to $TTIT - TIAT$. When the token holding time is up, the node must release the token as soon as the currently ongoing packet transmission is completed. A token can also be released earlier if no more packets are left in the node's transmission buffer. Note that concurrent transmissions on distinct channels are possible at the same node when two or more tokens are simultaneously held at the node.

With the $FT^W - FR^W$ node structure, MTIT avoids receiver collisions and allows each node to simultaneously use multiple data wavelength channels. However, the number of transceivers at each node is rather large. MTIT achieves low access delay due to the fact that a node has the opportunity to grab a token more frequently than in conventional token rings where a node has to wait one round-trip time for the next token. A unique feature of MTIT is its capability to self-adjust the relative positions of tokens along the ring circumference and maintain an even distribution of the token position. As a



■ **FIGURE 16.** MTIT node architecture.



■ FIGURE 17. SMARTNet: Meshed ring with $K = 6$ wavelength routers, each connected to its $M = 2$ nd neighboring routers.

result, the variance of the token inter-arrival time is low, guaranteeing to every node a consistent channel access delay in support of high-priority traffic. On the other hand, the capacity of MTIT is smaller than that of destination-stripping ring networks since source stripping does not allow for spatial wavelength reuse. For uniform traffic it was shown that MTIT achieves high bandwidth efficiency and low access delay for varying packet sizes even in relatively large (thousands of kilometers) networks. Both bandwidth efficiency and access delay improve with the number of wavelengths used in the ring.

MESHED RINGS

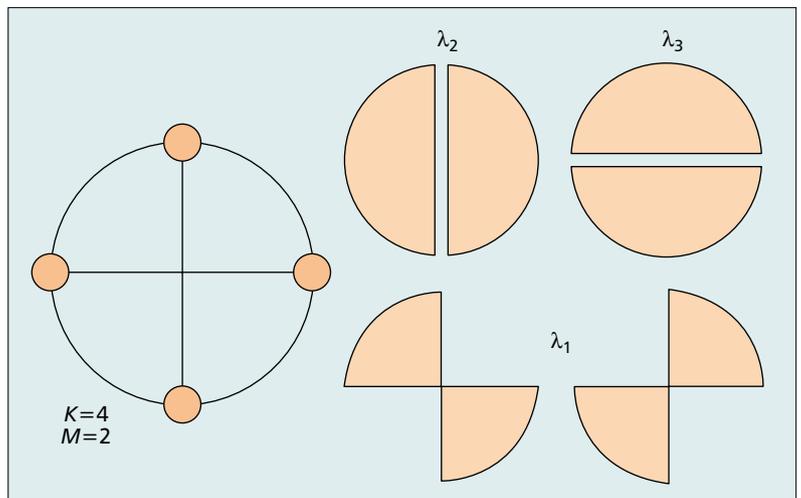
In unidirectional ring WDM networks with source stripping, packets are removed by the source node and each transmission requires a full circulation of the packet on the ring. The network capacity is limited by the aggregate capacity of all wavelengths. The network capacity of unidirectional ring networks can be increased with destination stripping where a transmission is propagating only on the ring segment between the corresponding pair of source and destination nodes. Due to spatial reuse, multiple simultaneous transmissions can take place on each wavelength. For uniform traffic, the mean distance between source and destination is half the ring circumference. As a consequence, two simultaneous transmissions can take place at each wavelength on average, resulting in a network capacity that is 200 percent as large as that of unidirectional rings with source stripping. In bidirectional rings, the network capacity can be further increased by means of *shortest path routing*, where a given packet is sent on that ring which provides the shortest distance to the corresponding destination. For uniform traffic the mean distance between source and destination is only a quarter of the ring circumference. Therefore, the aggregate capacity of bidirectional destination-stripping ring networks is increased by 400 percent compared to unidirectional source-stripping ring networks. The capacity of bidirectional ring WDM networks can be further increased by meshing the ring, which is discussed next.

The Scalable Multi-channel Adaptable Ring Terabit Network (SMARTNet) [83–86] achieves a significant increase in network capacity over the bidirectional destination-stripping ring by adding fiber short-cuts that connect certain nodes. More precisely, SMARTNet is based on a bidirectional slotted ring network with shortest path routing and destination stripping. Each node is connected to both rings and has a $FT^W - FR^W$ structure, which allows a node to simultaneously transmit and receive data on different wavelengths. All wavelengths are divided into fixed-size slots whose length is equal to the transmission time of a fixed-size packet plus a header for indicating the slot status. Medium access is governed by means of an empty-slot protocol.

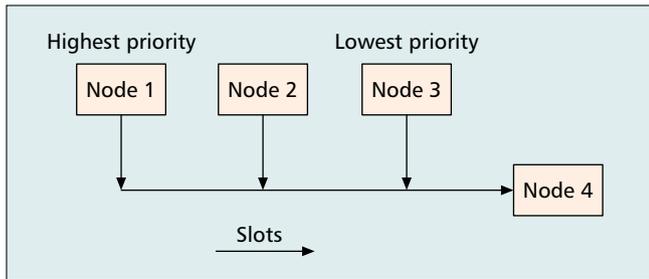
In addition to the N nodes, K equally spaced wavelength routers, each with four pairs of input/output ports, are deployed in the bidirectional ring. Wavelength routers are used to provide short-cuts, in that data packets do not have to pass through the ring nodes that are between two interconnected routers. Specifically, two input/output ports of each wavelength router are used to insert the router into the bidirectional ring; the other two pairs of ports are used for creating bidirectional links (chords) to the two M th neighboring routers. Routers $r_{[(k+M) \bmod K]}$ and $r_{[(k-M) \bmod K]}$ are said to be the M th neighboring routers of router r_k on the ring, where $k = 0, 1, \dots, K - 1$. Figure 17 depicts a meshed ring with $K = 6$ wavelength routers, each connected to its two $M = 2$ nd neighboring routers.

Each wavelength router is characterized by a wavelength routing matrix that determines to which output port each wavelength from a given input port is routed. The wavelength routing matrix is chosen such that the average distance between each source-destination pair is minimized with a minimum number of required wavelengths. For example, an optimal set of wavelength paths for $K = 4$, $M = 2$, and $W = 3$ is shown in Fig. 18.

SMARTNet is able to significantly increase the capacity of a bidirectional ring network with shortest path routing and destination stripping. For uniform traffic it was shown that a meshed ring with $K = 6$ wavelength routers and $M = 2$ increases network capacity by 720 percent compared to unidirectional source-stripping rings. Thus, the capacity of meshed rings is 80 percent larger than that achieved by nonmeshed bidirectional ring networks with destination stripping at the expense of additional wavelength routers and chords that add to network costs.



■ FIGURE 18. Wavelength paths in a meshed ring with $K = 4$ and $M = 2$, using $W = 3$ wavelengths.



■ FIGURE 19. Medium access priorities in ring networks.

We defer the summary of the results of this section to later on in the article.

FAIRNESS CONTROL AND QOS SUPPORT

Several of the aforementioned access protocols were extended in order to achieve fairness and quality of service (QoS) support. In this section, we discuss these protocol extensions in greater detail.

FAIRNESS CONTROL

In general, the bandwidth of a network is shared by all nodes. Each node ready to send data should have the same opportunity to transmit data. As we have seen in the preceding section, most of the packet-switched ring WDM networks are based on a unidirectional ring (see Fig. 9). In this architecture, each wavelength can be considered a unidirectional bus terminating at a prescribed destination (see Fig. 19). In an empty-slot access protocol, upstream nodes have a better than average chance to receive an empty slot for transmission, while downstream nodes have a worse than average chance. At heavy traffic this can lead to *starvation* of downstream nodes since they “see” slots that are mostly used by upstream nodes. To avoid starvation, the transmission rate of nodes has to be controlled in order to achieve fairness among all nodes. However, restricting nodes in their transmission decreases channel utilization. In general, there is a tradeoff between fairness and channel utilization.

We now comprehensively survey the fairness mechanisms that have been developed for slotted-ring networks.

MMR — Since SRR is not able to enforce fairness, a fairness-control algorithm is typically superimposed on SRR. The Multi-MetaRing (MMR) fairness algorithm is used on top of SRR in [87]. The MMR algorithm adapts a mechanism originally proposed for the MetaRing high-speed electronic metropolitan area network [32, 33, 88]. Fairness in the MetaRing is achieved by circulating a control message, named SAT (short for SATisfied). Each node is assigned a maximum number of packets to be transmitted between two SAT visits; this maximum number of packets is the node’s *quota* or *credit*. Each node normally forwards the SAT message on the ring with no delay, unless it is not SATisfied in the sense that it has not transmitted the permitted number of packets since the last time it forwarded the SAT. The SAT is delayed at unSATisfied nodes until SATisfaction is obtained, that is, either the node packet buffer is empty or the permitted number of packets has been transmitted.

In the MMR Single SAT (MMR-SS) scheme, a single SAT message regulates the transmissions of all nodes on all wavelength channels. Each node can transmit up to K packets to each destination since the last SAT visit. Each SATisfied node forwards the SAT to the upstream node. Thus, the SAT logically rotates in the opposite direction with respect to data

(although the physical propagation is co-directional). With this scheme the SAT propagation delays are very large since the SAT message has to traverse almost the entire network to reach the upstream node. Alternatively, the MMR Multiple SAT (MMR-MS) scheme uses one SAT message for each wavelength. It was shown in [89] that this MMR-MS scheme is generally the preferable extension of the MetaRing fairness-control scheme to a WDM ring.

M-ATMR — The access protocol from [73, 74] discussed earlier suffers from fairness problems due to destination stripping. In [90] Bengi and van As adopted an extension of the well-established Asynchronous Transfer Mode Ring (ATMR) fairness protocol to the multiple channel WDM ring case. This extension is M-ATMR. In M-ATMR each node receives a prescribed number of transmission credits for each destination. When a node has used all its credits or has nothing to send, it transitions into the inactive state. In order to properly apply the credit reset mechanism, every node has to know which node was the last active node. To achieve this, each active node overwrites a so-called *busy address field* in the header of every incoming slot with its own address. (The busy address field may be included in the SCM header of each WDM wavelength channel.) Thus, a node receiving a slot with its own busy address knows that all the other nodes are inactive. If the last active node detects inactivity of all the other nodes, it generates a reset immediately after its own transmission. The reset mechanism causes the nodes to reset their credits to the predefined values. In this manner, it is guaranteed that every node uses a maximum number of slots between two subsequent reset cycles. It was shown in [90] that the M-ATMR fairness protocol applied for best-effort traffic provides throughput and delay fairness for both uniform and client/server traffic scenarios.

DQBR — The Distributed Queue Bidirectional Ring (DQBR) fairness protocol [70] for the control channel based-HORNET of an earlier section is an adaptation of the distributed queue dual bus (DQDB) protocol. The DQBR fairness protocol works as follows. In each control-channel frame, a bit stream of length W bits, called the *request bit stream*, follows the wavelength-availability information. When a node on the network receives a packet in VOQ w , the node notifies the upstream nodes about the packet by setting bit w in the request bit stream in the control channel that travels upstream with respect to the direction the packet will travel. All upstream nodes take note of the requests by incrementing a counter called a *request counter* (RC). Each node maintains a separate RC for each wavelength. Thus, if bit w in the request bit stream is set, RC w is incremented. Each time a packet arrives at VOQ w , the node stamps the value in RC w onto the packet and then clears the RC. The value of this stamp is called the *wait counter* (WC). After the packet reaches the head of the VOQ, if the WC equals n it must allow n empty frames to pass by for downstream packets that were generated earlier. When an empty frame passes by the node on wavelength w , the WC for the packet at the head of VOQ w is decremented (if the WC equals zero, the RC w is decremented). Not until the WC equals zero can the packet be transmitted. The counting system ensures that the packets are sent in the order in which they arrived in the network.

The performance of the DQBR fairness-control scheme was investigated for a 25-node HORNET network by means of simulation for two traffic scenarios in [70]. In the first traffic scenario, variable-size packet traffic was uniformly randomly generated by the nodes. The traffic generated for node 18 was 1.5 times the capacity of wavelength 18. It was demon-

strated that with DQBR the throughput is equal for all nodes, whereas without DQBR the nodes close to node 18 have difficulty sending packets to node 18. In the second traffic scenario, unbalanced traffic was considered. Specifically, node 10 had 9.33 Gb/s of traffic arriving at its queue destined for node 18, node 11 had 4.67 Gb/s destined for node 18, and all other nodes had little traffic. The wavelength could support only 10 Gb/s, so it was heavily overloaded. It was found that without DQBR the nodes close to node 18 are unable to transmit packets on wavelength 18, whereas with DQBR all nodes have an equal ratio of throughput to load for wavelength 18.

QoS SUPPORT

Many applications, for example, multimedia traffic, require quality of service (QoS) with respect to throughput, delay, and jitter. To meet these requirements, networks typically provide different service classes, for example, constant bit rate (CBR) and variable bit rate (VBR). In general, in WDM networks traffic with stringent throughput, delay, and jitter requirements is supported by means of circuit switching via reservation of network resources, resulting in *guaranteed* QoS. On the other hand, to provide QoS to bursty traffic more efficiently, nodes process and forward packets with different priorities while benefiting from statistical multiplexing, leading to *statistical* QoS. In the following sections, we review different approaches for providing QoS in metro WDM ring networks.

SR³ — Synchronous Round Robin with Reservations (SR³) is derived from the SRR and MMR protocols and allows nodes to reserve slots, thereby achieving stronger control over access delays [91]. The SR³ protocol can be used in conjunction with SRR and MMR, requiring a marginal algorithmic complexity increase with no additional signaling messages.

In SR³, time is subdivided into successive periods called *reservation frames*. Each reservation frame consists of P SRR frames. Each node can reserve up to P slots for a given destination per reservation frame, that is, at most one slot per destination per SRR frame. Reservations are effective when all network nodes have become aware of the other nodes' reservations. SAT messages are used to broadcast the reservation information. Each SAT distributes information regarding current reservations on the channel it regulates. Each SAT contains a *reservation field* (SAT-RF) which is subdivided into $(N - 1)$ subfields; each subfield is assigned to a particular node for reservations. If node i needs to reserve h ($1 \leq h \leq P$) slots per reservation frame on channel j , it waits until it receives the j -SAT; it then forwards the reservation request after properly setting the i th SAT-RF subfield to the value h . The j -SAT visits all nodes during the next tour of the multi-ring. By the time node i again receives the j -SAT, all nodes in the network are aware of the request of node i . Node i can thus update its reservation request on channel j every time it releases the j -SAT.

It was shown in [91] that SR³ guarantees a throughput-fair access to each node. Moreover, the bandwidth left unused by guaranteed services can be shared by best-effort traffic very effectively. Even for the basic best-effort service that requires no service guarantee, the reservation scheme can be very beneficial; the average and variability of the access delays are greatly reduced when slots are reserved, leading to improved performance and fairness. The reservation scheme can also be extended to *multiple* service classes. It was shown in [92] that in an unbalanced multiclass traffic scenario, a very good separation of the different traffic classes is obtained. The performance of higher-priority traffic is largely unaffected by lower-priority traffic, even when the lower-priority traffic grows to overload conditions.

Reservation Scheme for QoS Support — For QoS support in the WDM ring network [73, 74] discussed in an earlier section, a connection-oriented protocol based on connection setup and termination was proposed in [65]. In order to enable connection-oriented packet transmission for real-time services, the ring is subdivided into so called *connection frames*. The real-time connections are established by reserving equally spaced slots within successive connection frames such that each destination node can be reached by a prescribed slot on the corresponding wavelength. Best-effort data traffic is transmitted in slots that are unreserved and empty. It was shown that this QoS approach is able to meet delay requirements almost deterministically. Note that this scheme allows for reserving only one fixed-size slot per frame, that is, only fixed-size packets are supported, similar to the SR³ scheme.

A similar reservation scheme for providing QoS was presented in [65, 90]. In addition to the W normal VOQs, each node has W real-time VOQs. Packets in the real-time VOQs are transmitted via connections in equally spaced, reserved slots. At each wavelength the ring is subdivided into frames, each consisting of N/W slots, one slot per destination node receiving on that wavelength. A single reservation slot carries a connection setup and a connection termination field, each consisting of N bits on the subcarrier. When a node sets a bit in the setup field, the slot to the corresponding destination is reserved in each frame. After one circulation of the reservation slot, all nodes are aware of the reservation and the setup flag is cleared. All nodes keep track of the reservations by maintaining a table that is updated when the reservation slot passes. To free the reserved slots, the same set/circulation/reset procedure is performed with the corresponding bit in the termination field.

MTIT-QoS With Lightpaths — The MTIT protocol discussed in a previous section can be extended to support not only packet switching but also circuit switching with guaranteed QoS [93]. The proposed solution allows for the all-optical transmission of packets with source stripping and circuits via a *tell-and-go* establishment of lightpaths (wavelength routes) with destination stripping. The lightpath establishment technique sets up a point-to-point connection between the source and the destination as follows. The on-off switches (Fig. 16) at both the source and the destination corresponding to the lightpath wavelength are set in the off state. As a consequence, the data transmission is restricted to the ring segment between the source and destination nodes. This allows downstream nodes following the destination node to spatially reuse the wavelength channel.

Each node maintains a *local lightpath table* (LLT) for all active lightpaths that is updated each time a token passes. A *token lightpath table* (TLT) is transmitted with each token to broadcast the changes of lightpath deployment on the ring on the wavelength associated with the token. Each token consists of two lists, the so called *add list* for circuit setup and the so called *delete list* for circuit teardown. Specifically, a node holding a token can set up a lightpath to a destination node at the token's wavelength by making an entry in the add list of the token. The path to the destination must not be occupied by another lightpath. A lightpath is torn down by the source by making an entry in the delete list of the token. Assuming uniform traffic with Poisson arrivals and exponentially distributed message lengths, it was analytically shown that an acceptable throughput/access delay performance can be achieved and that the achievable system throughput grows and access delay decreases as the number of wavelengths increases.

	MAWSON	RINGO	SRR	HORNET	Bengi <i>et al.</i>	Jelger <i>et al.</i>
Research focus	Testbed + protocol	Testbed	Protocol	Testbed + protocol	Protocol	Protocol
Special feature	Technically simple	—	—	—	—	—
Node structure	FT ^W -FR	FT ^W -FR, TT-FR	TT-FR	TT-FR	HORNET	FT-FR ^W , FT-TR
Scalability	$N = W$	$N = W$	$N \geq W$	$N \geq W$	$N \geq W$	$N \geq W$
Packet removal	(Dest.)	Dest.	Dest.	Dest.	Dest.	Dest.
Var. packet size	Reservation	—	—	Var. size slots	—	—
Fairness control	Not required	—	MMR	—	M-ATMR	—
QoS support	—	—	CBR + VBR	—	—	—
Perf. evaluation	Simulation	—	Analy. + sim.	Simulation	Analy. + Sim.	Analy. + sim.
References	[53, 54]	[42, 43, 56]	[58, 59, 89] [87, 92, 91]	[44, 45, 62]	[65, 66, 90]	[13, 67, 68]
	CC HORNET	CC Bengi	Smiljanic <i>et al.</i>	Cho <i>et al.</i>	MTIT	SmartNET
Research focus	Testbed + protocol	Protocol	Architec. + prot.	Concept	Protocol	Concept
Special feature	Bidirectional	—	Wavel. stacking	Virtual circles	Token ring	Meshed ring
Node structure	CC-TT ² /FT ² -FR ⁴	CC-TT/FT ² -FR ²	CC-TT/FT-FR ²	CC-DWADM	CC-FT ^{W+1} -FR ^{W+1}	FT ^W -FR ^W
Scalability	$N \geq W$	$N \geq W$	$N \geq W$	$N \geq W$	$N \geq W$	$N \geq W = 5$
Packet removal	Dest.	(Dest.)	Dest.	(Dest.)	Source	Dest.
Var. packet size	Reduced fragment.	Reservation	—	—	Yes	—
Fairness control	DQBR	—	—	—	Not required	—
QoS support	—	CBR	CBR	—	(CBR)	—
Perf. evaluation	Simulation	Simulation	Analysis	Simulation	Analy. + sim.	Analysis
References	[70]	[73, 74]	[76, 77]	[80]	[81, 82, 93]	[83, 84]

■ **Table 1.** Overview of surveyed packet-switched ring WDM networks.

CONCLUSIONS

In this article, we have provided a comprehensive up-to-date survey of studies of packet-switched ring WDM networks. The current goal of the research on ring WDM networks is to develop designs that overcome the emerging metro gap between high-speed local clients (and networks) and the very-high-speed backbone networks. To overcome this metro gap, the ring networks need to efficiently use the wavelength resources, to be easily upgradeable (and scalable), and to flexibly support varying traffic loads and packet formats in a fair and cost effective manner. In this survey we attempt to provide a qualitative assessment of how the developed networks address the metro gap issues and to outline open areas for future research efforts. Toward this end, in Table 1 we contrast the surveyed networks in terms of node structure, scalability, and packet removal, as well as support for variable-size packet fairness and QoS. We also consider the focus and perspective of the research efforts and the method of performance evaluation. If the packet removal is not explicitly addressed for a particular network, but could be done by the destination, we denote (Dest.) for the packet removal. For the HORNET and Bengi networks we consider the versions with and without control channel separately.

We see from the table that among the networks not having a control channel, the TT-FR node structure is most common. Indeed, this node structure is relatively simple and effective for unicast packet transmissions. For multicast traffic, which requires multiple transmissions on the different drop wavelengths of the fixed-tuned receivers, the FT^W structure has the advantage that these transmissions can be conducted simultaneously. With the TT structure, on the other hand, multiple sequential transmissions are required. As noted in the table, for a control channel-based network, an FT-FR transceiver that is used exclusively for control is added. It may be worthwhile to investigate the cost effectiveness tradeoffs between

operating these dedicated control components and control wavelength channel, on the one hand, and conducting the control over the data transceiver and data channels, on the other hand. In conjunction with this question it may be of interest to explore whether the control channel and control components could be efficiently used to also carry some data traffic, for example, multicast and broadcast data traffic that has to reach a large number of receivers, similar to control traffic. As we observe from the table, the entire single-fiber node structure is duplicated for the dual-fiber HORNET. An important direction for future research is to investigate effective protection strategies for such multi-fiber rings, as well as the scaling to additional rings for very high capacity networks, for example, similar to [94].

We see from the table that all protocol-oriented and concept-oriented research efforts (as well as the HORNET testbed) allow for easy scalability in the number of nodes. The proof-of-concept testbeds MAWSON and RINGO, on the other hand, are at present limited to as many nodes as there are wavelength channels. There appears to be a need for more testbed activity on scalable networks.

All networks, except for the token ring network, allow for destination removal (stripping) and can thus exploit spatial wavelength reuse. Spatial wavelength reuse is not possible in the source-stripping token ring. However, the source-stripping in conjunction with token-passing does have several advantages, such as easy support for fairness and QoS. Clearly the challenge for ring networks is to achieve the efficiency of spatial wavelength reuse while at the same time providing QoS and fairness for variable-size packets. As surveyed in this article and indicated in Table 1, a number of techniques have recently been developed to support some combination of support for variable-size packets, fairness, and QoS in the different destination-stripping networks, and this area appears to continue to be a very active research area.

We note from Table 1 that the developed networks have been evaluated either by analysis, simulation, or experimentation, or a combination of analysis/simulation or simulation/experiment. There appears to be a need to complement the experiment (and experiment/simulation) evaluations with formal analysis, which may lead to fundamental insights that can enhance the considered testbed implementations. Similarly, it may be worthwhile to test the concept and protocol developments that have so far been evaluated by analysis and simulation in future testbeds.

We finally note that in our opinion, an attractive way to deploy WDM ring networks is the upgrade of RPR networks. We believe that this upgrade will become necessary because RPR networks with their inefficiencies (due to the single-channel nature of RPR and the OEO conversion of all traffic at each node) will no longer cope with the increasing amount of data traffic. Given the fiber structure of RPR, WDM is a promising and natural candidate to increase the capacity of installed RPR networks in an economical manner. We think that the surveyed WDM ring network and node architectures provide valuable insights and guidelines as to how to extend current single-channel RPR networks to multichannel WDM ring networks. In particular, some of the surveyed OOO ring node architectures can be used to replace RPR's opaque OEO node structure with an all-optical structure that enables not only transparency but also future-proofness with respect to protocols, data rate, and modulation format. In addition, the mechanisms and techniques surveyed in the previous section may be useful to extend RPR's fairness and QoS features to a WDM environment. In summary, current RPR ring networks belong to the family of first-generation (opaque) optical networks (Fig. 2) and the surveyed material may help to migrate toward second-generation (transparent) RPR ring networks.

ACKNOWLEDGMENT

We are indebted to Prof. Adam Wolisz, head of the telecommunication networks (TKN) institute at the Technical University of Berlin, Germany, for providing an environment that allowed for this work. The support of Martin Herzog by the graduate, interdisciplinary engineering research program (Graduieratenkolleg) in Stochastic Modelling and Quantitative Analysis of Complex Systems in Engineering at the Technical University of Berlin, funded by the Deutsche Forschungsgemeinschaft (DFG), is gratefully acknowledged.

REFERENCES

- [1] B. Mukherjee, "WDM Optical Communication Networks: Progress and Challenges," *IEEE JSAC*, vol. 18, no. 10, Oct. 2000, pp. 1810–24.
- [2] P.-H. Ho and H. T. Mouftah, "A Framework of Scalable Optical Metropolitan Networks for Improving Survivability and Class of Service," *IEEE Network*, vol. 16, no. 4, July/Aug. 2002, pp. 29–35.
- [3] S. Yao, S. J. B. Yoo, and B. Mukherjee, "All-Optical Packet Switching for Metropolitan Area Networks: Opportunities and Challenges," *IEEE Commun. Mag.*, vol. 39, no. 3, Mar. 2001, pp. 142–148.
- [4] G. Barish and K. Obraczka, "World Wide Web Caching: Trends and Techniques," *IEEE Commun. Mag.*, vol. 38, no. 5, May 2000, pp. 178–185.
- [5] L. G. Kazovsky et al., "High Capacity Metropolitan Area Networks for the Next-Generation Internet," *Proc. 35th Asilomar Conf. Signals, Systems, and Comp.*, vol. 1, Pacific Grove, CA, Nov. 2001, pp. 3–7.
- [6] L. G. Kazovsky et al., "Optical Metropolitan Area Networks," *Proc. Optical Fiber Communication Conf. and Exhibit (OFC)*, paper WU1, vol. 3, Anaheim, CA, Mar. 2001, pp. WU1-1–WU1-3.
- [7] I. Chlamtac and A. Ganz, "A Multibus Train Communication (AMTRAC) Architecture for High-Speed Fiber Optic Networks," *IEEE JSAC*, vol. 6, no. 6, July 1988, pp. 903–12.
- [8] B. Mukherjee, "WDM-Based Local Lightwave Networks Part I: Single-Hop Systems," *IEEE Network*, vol. 6, no. 3, May 1992, pp. 12–27.
- [9] —, "WDM-Based Local Lightwave Networks Part II: Multihop Systems," *IEEE Network*, vol. 6, no. 4, July, 1992, pp. 20–32.
- [10] M. Maier, *Metropolitan Area WDM Networks: An AWG-Based Approach*, Kluwer Academic Publishers, Norwell, MA, 2003.
- [11] M. Maier, M. Reisslein, and A. Wolisz, "Toward Efficient Packet Switching Metro WDM Networks," *Opt. Net. Mag.*, vol. 3, no. 6, Nov./Dec. 2002, pp. 44–62.
- [12] K. Bengi, *Optical Packet Access Protocols for WDM Networks*, Kluwer Academic Publishers, Norwell, MA, 2002.
- [13] C. S. Jelger and J. M. H. Elmirghani, "Photonic Packet WDM Ring Networks Architecture and Performance," *IEEE Commun. Mag.*, vol. 40, no. 11, Nov. 2002, pp. 110–15.
- [14] L. Wuttisittikulij and M. J. O'Mahony, "Design of a WDM Network using a Multiple Ring Approach," *Proc. IEEE GLOBECOM*, vol. 1, Kobe, Japan, Apr. 1997, pp. 551–55.
- [15] H. Woesner, "Primenet: A Concept for a WDM-Based Fiber Backbone," *Proc. Optical Network Design and Modelling (ONDM)*, Rome, Italy, Apr. 1998, pp. 98–106.
- [16] C.-J. Chae and R. S. Tucker, "Virtual Multiple Optical Ring Networks for Interconnecting Metropolitan IP/WDM Routers using an $N \times N$ Arrayed Waveguide Grating," *Proc. the 14th Annual Meeting of the IEEE Lasers and Electro-Optics Society (LEOS)*, vol. 1, San Diego, California, Nov. 2001, pp. 95–96.
- [17] M. N. Ransom and D. R. Spears, "Applications of Public Gigabit Networks," *IEEE Network*, vol. 6, no. 2, Mar, pp. 30–40. 1992.
- [18] P. E. Green, *Fiber Optic Networks*, Prentice Hall, 1993.
- [19] R. Ramaswami, "Optical Fiber Communication: From Transmission to Networking," *IEEE Commun. Mag.*, vol. 40, no. 5, May 2002, pp. 138–47.
- [20] S. Mokbel, "Canada's Optical Research and Education Network: CA*net3," *Proc. Design of Reliable Communication Networks (DRCN)*, Munich, Germany, Apr. 2000, pp. 10–32.
- [21] Y. Cai, R. M. Fortenberry, and R. S. Tucker, "Demonstration of Photonic Packet-Switched Ring Network with Optically Transparent Nodes," *IEEE Photonics Technology Letters*, vol. 6, no. 9, Sept. 1994, pp. 1139–41.
- [22] D. Guo and A. S. Acampora, "Scalable Multihop WDM Passive Ring with Optimal Wavelength Assignment and Adaptive Wavelength Routing," *IEEE/OSA J. Lightwave Tech.*, vol. 14, no. 6, June 1996, pp. 1264–77.
- [23] E. Modiano and P. J. Lin, "Traffic Grooming in WDM Networks," *IEEE Commun. Mag.*, vol. 39, no. 7, July 2001, pp. 124–29.
- [24] R. Jain, "Optical Networking: Recent Developments, Issues, and Trends," *Tutorial at IEEE Infocom 2003*, San Francisco, CA, Mar. 2003, <http://www.cis.ohio-state.edu/~jain>.
- [25] R. D. Doverspike, S. J. Phillips, and J. R. Westbrook, "Transport Network Architectures in an IP World," *Proc. IEEE Infocom*, Tel Aviv, Israel, Mar. 2000, pp. 305–314.
- [26] I. M. White, "A New Architecture and Technologies for High-Capacity Next-Generation Metropolitan Networks," Ph.D. dissertation, Stanford University, 2002.
- [27] ITU-T Rec. G.7041, "Generic Framing Procedure (GFP)," Oct. 2001.
- [28] ITU-T Rec. G.707, "Network Node Interface for the Synchronous Data Hierarchy," Oct. 2000.
- [29] ITU-T Rec. G.7042, "Link Capacity Adjustment Scheme (LCAS) for Virtual Concatenation," Oct. 2001.
- [30] S. S. Gorshe and T. Wilson, "Transparent Generic Framing Procedure (GFP): A Protocol for Efficient Transport of Block-coded Data through SONET/SDH Networks," *IEEE Commun. Mag.*, vol. 40, no. 5, May 2002, pp. 88–95.
- [31] D. E. Hubber, W. Steinlin, and P. J. Wild, "SILK: An Implementation of a Buffer Insertion Ring," *IEEE JSAC*, vol. SAC-1, no. 5, Nov. 1983, pp. 1214–23.
- [32] I. Cidon and Y. Ofek, "Metaring: A Full-Duplex Ring with Fairness and Spatial Reuse," *Proc. IEEE INFOCOM*, San Francisco, CA, June 1990, pp. 969–78.

- [33] —, "MetaRing: A Full-Duplex Ring with Fairness and Spatial Reuse," *IEEE Trans. Commun.*, vol. 41, no. 1, Jan. 1993, pp. 110–20.
- [34] V. Gamberoza et al., "Design, Analysis, and Implementation of DVSR: A Fair, High Performance Protocol for Packet Rings," *IEEE/ACM Trans. Net.*, vol. 12, no. 1, Feb. 2004.
- [35] G. Kramer and G. Pesavento, "Ethernet Passive Optical Network (EPON): Building a Next-Generation Optical Access Network," *IEEE Commun. Mag.*, vol. 40, no. 2, Feb. 2002, pp. 66–73.
- [36] T. Shan, J. Yang, and C. Sheng, "EPON Upstream Multiple Access Scheme," *Proc. IEEE Int'l. Conf. Infotech and Infonet (ICII)*, vol. 2, Beijing, China, October/November 2001, pp. 273–278.
- [37] G. Kramer, B. Mukherjee, and G. Pesavento, "IPACT: A Dynamic Protocol for an Ethernet PON (EPON)," *IEEE Commun. Mag.*, vol. 40, no. 2, Feb. 2002, pp. 74–80.
- [38] A. Gumaste and I. Chlamtac, "A Protocol to Implement Ethernet Over PON," *Proc. IEEE Int'l. Conf. Commun. (ICC)*, vol. 2, Anchorage, AK, May 2003, pp. 1345–1349.
- [39] M. C. Nuss, "Optical Ethernet in the Metro," *Proc. 14th Annual Meeting of the IEEE Lasers and Electro-Optics Society (LEOS)*, vol. 1, San Diego, CA, Nov. 2001, p. 289.
- [40] A. Richter et al., "Germany-Wide DWDM Field Trial: Transparent Connection of a Long Haul Link and a Multiclient Metro Network," *Proc. Optical Fiber Commun. Conf. and Exhibit (OFC)*, paper ML3, vol. 1, Anaheim, CA, Mar. 2001, pp. ML3/1–ML3/3.
- [41] D. Stoll et al., "Metropolitan DWDM: A Dynamically Configurable Ring for the KomNet Field Trial in Berlin," *IEEE Commun. Mag.*, vol. 39, no. 2, Feb. 2001, pp. 106–13.
- [42] R. Gaudino et al., "RINGO: A WDM Ring Optical Packet Network Demonstrator," *Proc. 27th European Conf. Optical Commun. (ECOC)*, vol. 4, Amsterdam, Netherlands, Oct. 2001, pp. 620–21.
- [43] R. Gaudino, "RINGO: Demonstration of a WDM Packet Network Architecture for Metro Applications," *Proc. 4th Int'l. Conf. Transparent Opt. Net. (ICTON)*, vol. 1, Warsaw, Poland, Apr. 2002, pp. 77–80.
- [44] S. M. Gemelos et al., "WDM Metropolitan Area Network Based on CSMA/CA Packet Switching," *IEEE Photonics Technology Letters*, vol. 11, no. 11, Nov. 1999, pp. 1512–14.
- [45] K. V. Shrikhande et al., "HORNET: A Packet-Over-WDM Multiple Access Metropolitan Area Ring Network," *IEEE JSAC*, vol. 18, no. 10, Oct. 2000, pp. 2004–16.
- [46] Y. Fukashiro et al., "Fast and Fine Tuning of a GCSR Laser using a Digitally Controlled Driver," *Proc. Optical Fiber Communication Conference and Exhibit (OFC)*, paper WM43, vol. 2, Baltimore, MD, Mar. 2000, pp. 338–40.
- [47] K. Shrikhande et al., "Performance Demonstration of a Fast-Tunable Transmitter and Burst-Mode Packet Receiver for HORNET," *Proc. Opt. Fiber Commun. Conf. and Exhibit (OFC)*, paper ThG2, vol. 4, Anaheim, CA, Mar. 2001, pp. ThG2-1–ThG2-3.
- [48] D. Wonglumson et al., "HORNET: A Packet-Switched WDM Metropolitan Area Ring Network: Optical Packet Transmission and Recovery, Queue Depth, and Packet Latency," *Proc. IEEE Lasers and Electro-Optics Society (LEOS)*, vol. 2, San Francisco, CA, Nov. 1999, pp. 653–54.
- [49] D. Wonglumson et al., "Experimental Demonstration on an Access Point for HORNET: A Packet-Over-WDM Multiple-Access MAN," *IEEE/OSA J. Lightwave Tech.*, vol. 18, no. 12, Dec. 2000, pp. 1709–17.
- [50] I. M. White et al., "Experimental Demonstration of a Media Access Protocol for HORNET: A WDM Multiple Access Metropolitan Area Ring Network," *Proc. Optical Fiber Commun. Conf. and Exhibit (OFC)*, paper WD3, vol. 2, Baltimore, MD, Mar. 2000, pp. 50–52.
- [51] I. M. White et al., "Experimental Demonstration of a Novel Media Access Protocol for HORNET: A Packet-Over-WDM Multiple-Access MAN Ring," *IEEE Photonics Tech. Letters*, vol. 12, no. 9, Sept. 2000, pp. 1264–66.
- [52] S. Johansson et al., "A Cost Effective Approach to Introduce an Optical WDM Network in the Metropolitan Environment," *IEEE JSAC*, vol. 16, no. 7, Sept. 1998, pp. 1109–22.
- [53] M. A. Summerfield, "MAWSON: A Metropolitan Area Wavelength Switched Optical Network," *Proc. Asia Pacific Conf. Commun. (APCC)*, vol. 1, Sydney, Australia, Nov. 1997, pp. 327–31.
- [54] J. Fransson et al., "Design of a Medium Access Control Protocol for a WDMA/TDMA Photonic Ring Network," *Proc. IEEE GLOBECOM*, vol. 1, Sydney, Australia, Nov. 1998, pp. 307–12.
- [55] M. J. Spencer and M. A. Summerfield, "WRAP: A Medium Access Control Protocol for Wavelength-Routed Passive Optical Networks," *IEEE/OSA J. Lightwave Tech.*, vol. 18, no. 12, Dec. 2000, pp. 1657–76.
- [56] A. Carena et al., "RINGO: A Demonstrator of WDM Optical Packet Network on a Ring Topology," *Proc. Optical Network Design and Modelling (ONDM)*, Torino, Italy, Feb. 2002, pp. 183–97.
- [57] A. Bianco et al., "Variable-Size Packets in Slotted WDM Ring Networks," *Proc. Optical Network Design and Modelling (ONDM)*, Torino, Italy, Feb. 2002, pp. 151–66.
- [58] M. A. Marsan et al., "On the Capacity of MAC Protocols for All-Optical WDM Multi-Rings with Tunable Transmitters and Fixed Receivers," *Proc. IEEE INFOCOM*, vol. 3, Mar. 1996, pp. 1206–16.
- [59] —, "MAC Protocols and Fairness Control in WDM Multirings with Tunable Transmitters and Fixed Receivers," *IEEE J. Lightwave Tech.*, vol. 14, no. 6, June 1996, pp. 1230–44.
- [60] M. A. Marsan et al., "Modeling Slotted WDM Rings with Discrete-Time Markovian Models," *Computer Networks*, vol. 32, no. 5, May 2000, pp. 599–615.
- [61] A. Bianco et al., "A-Posteriori Access Strategies in All-Optical Slotted WDM Rings," *Proc. IEEE GLOBECOM*, vol. 1, Sydney, Australia, Nov. 1998, pp. 300–06.
- [62] K. Shrikhande et al., "CSMA/CA MAC Protocols for IP-HORNET: An IP over WDM Metropolitan Area Ring Network," *Proc. IEEE GLOBECOM*, vol. 2, San Francisco, CA, Nov./Dec. 2000, pp. 1303–07.
- [63] W.-P. Chen and W.-S. Hwang, "A Packet Pre-Classification CSMA/CA MAC Protocol for IP over WDM Ring Networks," *Proc. IEEE Int'l. Conf. Commun. Systems (ICCS)*, vol. 2, Singapore, Nov. 2002, pp. 1217–21.
- [64] C.-C. Li, S.-W. Kau, and W.-S. Hwang, "A CSMA/CP MAC Protocol for IP over WDM Metropolitan Area Ring Networks," *Proc. IEEE Int'l. Conf. Commun. Systems (ICCS)*, vol. 2, Singapore, Nov. 2002, pp. 1212–16.
- [65] K. Bengi and H. R. van As, "Efficient QoS Support in a Slotted Multihop WDM Metro Ring," *IEEE JSAC*, vol. 20, no. 1, Jan. 2002, pp. 216–27.
- [66] K. Bengi, "An Analytical Model for a Slotted WDM Metro Ring with A-Posteriori Access," *Proc. Optical Network Design and Modelling (ONDM)*, Torino, Italy, Feb. 2002.
- [67] C. S. Jelger and J. M. H. Elmirghani, "A Simple MAC Protocol for WDM Metropolitan Access Ring Networks," *Proc. IEEE GLOBECOM*, vol. 3, San Antonio, TX, Nov. 2001, pp. 1500–04.
- [68] —, "Performance of a Slotted MAC Protocol for WDM Metropolitan Access Ring Networks under Self-Similar Traffic," *Proc., IEEE Int'l. Conf. Commun. (ICC)*, vol. 5, New York, NY, Apr. 2002, pp. 2806–11.
- [69] E. W. M. Wong, A. Fumagalli, and I. Chlamtac, "Performance Evaluation of CROWNS: WDM Multi-Ring Topologies," *Proc. IEEE Int'l. Conf. Commun. (ICC)*, vol. 2, Seattle, WA, June 1995, pp. 1296–1301.
- [70] I. M. White et al., "Design of a Control-Channel-Based Media-Access-Control Protocol for HORNET," *J. Opt. Net.*, vol. 1, no. 12, Dec. 2002, pp. 460–73.
- [71] I. M. White et al., "Architecture and Protocols for HORNET: A Novel Packet-over-WDM Multiple-Access MAN," *Proc. IEEE GLOBECOM*, vol. 2, San Francisco, CA, Nov./Dec. 2000, pp. 1298–02.
- [72] I. M. White et al., "Experimental Demonstration of the HORNET Survivable Bi-Directional Ring Architecture," *Proc. Opt. Fiber Commun. Conf. and Exhibit (OFC)*, paper WW1, Anaheim, CA, Mar. 2002, pp. 346–49.
- [73] K. Bengi, "An Optical Packet-Switched IP-over-WDM Metro Ring Network," *Proc., 27th Annual IEEE Conf. Local Comp. Net. (LCN)*, Tampa, FL, Nov. 2002, pp. 43–52.

- [74] —, "Access Protocols for an Efficient Optical Packet-Switched Metropolitan Area Ring Network Supporting IP Datagrams," *Proc. 11th Int'l. Conf. Comp. Commun. and Networks (ICCCN)*, Miami, FL, Oct. 2002, pp. 284–89.
- [75] M. Boroditsky et al., "Experimental Demonstration of Composite Packet Switching on a WDM Photonic Slot Routing Network," *Proc., Optical Fiber Commun. Conf. and Exhibit (OFC)*, paper ThG6, vol. 4, Anaheim, CA, Mar. 2001, pp. ThG6-1–ThG6-3.
- [76] A. Smiljanic, M. Boroditsky, and N. J. Frigo, "Optical Packet-Switched Ring Network with Flexible Bandwidth Allocation," *Proc. IEEE Wksp. High Perf. Switching and Routing (HPSR)*, Dallas, TX, May 2001, pp. 83–87.
- [77] —, "High-Capacity Packet-Switched Optical Ring Network," *IEEE Communications Letters*, vol. 6, no. 3, Mar. 2002, pp. 111–13.
- [78] I. Chlamtac et al., "Scalable WDM Access Network Architecture Based on Photonic Slot Routing," *IEEE/ACM Trans. Net.*, vol. 7, no. 1, Feb. 1999, pp. 1–9.
- [79] V. Elek, A. Fumagalli, and G. Wedzinga, "Photonic Slot Routing: A Cost Effective Approach to Designing All-Optical Access and Metro Networks," *IEEE Commun. Mag.*, vol. 39, no. 11, Nov. 2001, pp. 164–72.
- [80] W. Cho and B. Mukherjee, "Design of MAC Protocols for DWADM-Based Metropolitan-Area Optical Ring Networks," *Proc. IEEE GLOBECOM*, vol. 3, San Antonio, TX, Nov. 2001, pp. 1575–79.
- [81] A. Fumagalli, J. Cai, and I. Chlamtac, "The Multi-Token Inter-Arrival Time (MTIT) Access Protocol for Supporting IP over WDM Ring Network," *Proc. IEEE Int'l. Conf. Commun.*, vol. 1, Vancouver, Canada, June 1999, pp. 586–90.
- [82] J. Cai, A. Fumagalli, and I. Chlamtac, "The Multitoken Interarrival Time (MTIT) Access Protocol for Supporting Variable Size Packets Over WDM Ring Network," *IEEE JSAC*, vol. 18, no. 10, pp. 2094–2104, Oct. 2000.
- [83] I. Rubin and H.-K. Hua, "An All-Optical Wavelength-Division Meshed-Ring Packet-Switching Network," *Proc. IEEE INFOCOM*, vol. 3, Boston, MA, Apr. 1995, pp. 969–76.
- [84] —, "SMARTNet: An All-Optical Wavelength-Division Meshed-Ring Packet-Switching Network," *Proc. IEEE GLOBECOM*, vol. 3, Singapore, Nov. 1995, pp. 1756–60.
- [85] I. Rubin and H.-K. H. Hua, "Synthesis and Throughput Behavior of WDM Meshed-Ring Networks under Nonuniform Traffic Loading," *IEEE/OSA J. Lightwave Tech.*, vol. 15, no. 8, Aug. 1997, pp. 1513–21.
- [86] I. Rubin and J. Ling, "All-Optical Cross-Connect Meshed-Ring Communications Networks using a Reduced Number of Wavelengths," *Proc. IEEE INFOCOM*, vol. 2, New York, NY, Mar. 1999, pp. 924–31.
- [87] M. A. Marsan et al., "An Almost Optimal MAC Protocol for All-Optical WDM Multi-Rings with Tunable Transmitters and Fixed Receivers," *Proc. IEEE Int'l. Conf. Commun. (ICC)*, vol. 1, Montreal, June 1997, pp. 437–42.
- [88] J. Chen, I. Cidon, and Y. Ofek, "A Local Fairness Algorithm for the MetaRing, and its Performance Study," *Proc. IEEE GLOBECOM*, vol. 3, Orlando, FL, Dec. 1992, pp. 1635–41.
- [89] M. A. Marsan, A. Bianco, E. Leonardi, F. Neri, and S. Toniolo, "MetaRing Fairness Control Schemes in All-Optical WDM Rings," *Proc. IEEE INFOCOM*, vol. 2, Kobe, Japan, Apr. 1997, pp. 752–60.
- [90] K. Bengi and H. R. van As, "QoS Support and Fairness Control in a Slotted Packet-Switched WDM Metro Ring Network," *Proc. IEEE GLOBECOM*, vol. 3, San Antonio, TX, Nov. 2001, pp. 1494–99.
- [91] M. A. Marsan et al., "SR³: A Bandwidth-Reservation MAC Protocol for Multimedia Applications over All-Optical WDM Multi-Rings," *Proc. IEEE INFOCOM*, vol. 2, Kobe, Japan, Apr. 1997, pp. 761–768.
- [92] —, "All-Optical WDM Multi-Rings with Differentiated QoS," *IEEE Commun. Mag.*, vol. 37, no. 2, Feb. 1999, pp. 58–66.
- [93] A. Fumagalli, J. Cai, and I. Chlamtac, "A Token-Based Protocol for Integrated Packet and Circuit Switching in WDM Rings," *Proc. IEEE GLOBECOM*, vol. 4, Sydney, Australia, Nov. 1998, pp. 2339–44.
- [94] A. Lee et al., "Heuristic for Setting Up a Stack of WDM Rings with Wavelength Reuse," *IEEE J. Lightwave Tech.*, vol. 18, no. 4, Apr. 2000, pp. 521–29.

BIOGRAPHIES

MARTIN HERZOG received the Dipl.-Ing. degree (with distinctions) in computer engineering at the Technical University of Berlin in 2002. He is currently a Ph.D. student at the Telecommunication Networks Group, also at the TU Berlin, and participates in a graduate, interdisciplinary engineering research program. His research interests lie in the area of optical wavelength division multiplexing (WDM) networks with a focus on architectures and access protocols for metro networks.

MARTIN MAIER received the Dipl.-Ing. and 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 (USC), Los Angeles, and Arizona State University (ASU), Tempe. In the summer of 2003 he was a postdoc fellow at Massachusetts Institute of Technology (MIT), Cambridge. 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 particular focus on metropolitan area networks. Recently, his research interests have been concentrated on evolutionary WDM upgrades, control, and management of optical ring networks. He was a recipient of the two-year Deutsche Telekom Doctoral Scholarship from June 1999 through May 2001. He is also a co-recipient of the Best Paper Award presented at the SPIE Photonics East 2000-Terabit Optical Networking Conference. He is the author of the book *Metropolitan Area WDM Networks: An AWG Based Approach* (Norwell, MA: Kluwer, 2003).

MARTIN REISSLEIN (reisslein@asu.edu) is an assistant professor in the Department of Electrical Engineering at Arizona State University, Tempe. He received the Dipl.-Ing. (FH) degree from the Fachhochschule Dieburg, Germany, in 1994, and the M.S.E. degree from the University of Pennsylvania, Philadelphia, in 1996. both in electrical engineering. He received his Ph.D. in systems engineering from the University of Pennsylvania in 1998. 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. While in Berlin he was teaching courses on performance evaluation and computer networking at the Technical University Berlin. He 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 Computers 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 co-recipient of the Best Paper Award of the SPIE Photonics East 2000 - Terabit Optical Networking conference. His research interests are in the areas of Internet QoS, video traffic characterization, wireless networking, and optical networking.