STARGATE: The Next Evolutionary Step toward Unleashing the Potential of WDM EPONs

Martin Maier and Martin Herzog, Institut National de la Recherche Scientifique (INRS)
Martin Reisslein, Arizona State University

This work was supported by the NSERC Discovery Grants Program (Grant 327320-06) and FQRNT Etablissement de Nouveaux Chercheurs Program (Grant 114954).

ABSTRACT

WDM EPONs are rapidly becoming mature. In this article we briefly review the state of the art of cost reduction, colorless ONUs, and WDM PONs. We then lift the veil of future WDM EPONs, and elaborate on evolutionary cost-effective upgrades of WDM EPONs and their all-optical WDM integration with Ethernet-based metropolitan area networks to provide transparent connections at the wavelength and sub-wavelength granularity on demand between ONUs residing in different WDM EPONs.

INTRODUCTION

Future broadband passive optical networks (PONs) not only have to unleash their economic potential and societal benefits by opening up the first/last mile bandwidth bottleneck between bandwidth-hungry end users and high-speed backbone networks; they also must enable the support of a wide range of new and emerging services and applications, such as triple play, video on demand, videoconferencing, peer-to-peer (P2P) audio/video file sharing, multichannel high-definition television, multimedia/multiparty online gaming, telemedicine, telecommuting, and surveillance, to get back on the road to prosperity [1]. Due to their longevity, low attenuation, and huge bandwidth, asynchronous transfer mode (ATM) and Ethernet-based PONs are already widely deployed in today's operational access (e.g., fiber to the premises [FTTP] and fiber to the home [FTTH]) networks. Typically, these PONs are time-division multiplexing (TDM) single-channel systems, where the fiber infrastructure carries a single upstream wavelength channel and a single downstream wavelength channel. To support the aforementioned emerging services and applications in a cost-effective and future-proof manner, and to unleash the full potential of FTTx networks, PONs need to evolve by addressing the following three tasks [2].

Cost reduction: Cost is key in access networks due to the small number of cost-sharing subscribers compared to metro and wide area networks. Devices and components that can be mass produced and widely applied to different types of equipment and situations must be developed. It is important that installation costs, which largely contribute to overall costs, be reduced. A promising example for cutting installation costs is NTT's envisioned do-it-yourself (DIY) installation deploying a user-friendly hole-assisted fiber that exhibits negligible loss increase and sufficient reliability, even when it is bent at right angles, clinched, or knotted, and can be produced economically.

Colorless ONU: The next target is to make the optical network unit (ONU), which connects one or more subscribers to the PON, colorless (i.e., wavelength-independent). Colorless ONUs require either no light source at all or only a broadband light source, resulting in decreased costs, simplified maintenance, and reduced stock inventory issues.

WDM PONs: The third and final target is to increase the number of wavelength channels by means of wavelength-division multiplexing (WDM). The use of WDM technologies allows access network operators to respond to user requests for service upgrades and network evolution. Deploying WDM adds a new dimension to current TDM PONs. The benefits of the new wavelength dimension are manifold. Among others, it may be exploited:
- To increase network capacity
- To improve network scalability by accommodating more end users
- To separate services
- To separate service providers

All three tasks are currently addressed by various research groups worldwide [3]. For example, in the United States, FTTH costs per connected home dropped by a factor of almost 5 between 1993 and 2004 [1].

PONs come in a number of flavors, such as ATM-based APONs, broadband PONs

 0163-6804/07/$20.00 © 2007 IEEE

(BPONs), and gigabit PONs (GPONs). In this article we focus on Ethernet PONs (EPONs), standardized by the IEEE 802.3ah Ethernet in the First Mile (EFM) Task Force. EPONs aim at converging the low-cost equipment and simplicity of Ethernet and the low-cost infrastructure of PONs. Given the fact that 95 percent of LANs use Ethernet, EPONs and their WDM upgraded descendants are likely to become increasingly the norm due to their capability of natively carrying variable-size IP packets in a simpler and more efficient way than their ATM-based counterparts, which suffer from a severe cell tax and 125 μs framing overhead [1]. Besides access networks, Ethernet is also gaining ground in metropolitan and even wide area networks, giving rise to end-to-end Ethernet networks. The recently approved IEEE 802.17 resilient packet ring (RPR) standard aims at combining Ethernet's statistical multiplexing gain, low equipment cost, and simplicity with synchronous optical network/synchronous digital hierarchy's (SONET/SDH's) carrier-class functionalities of high availability, reliability, and profitable TDM (voice) support to build high-performance metro edge and metro core ring networks that interconnect multiple access networks [4].

In this article we consider WDM EPONs and, arguing that the three aforementioned tasks will be addressed successfully in the near term, elaborate on the question "WDM EPON — what's next?" Our focus is on evolutionary upgrades and further cost reductions of WDM EPONs, and the all-optical integration of Ethernet-based WDM EPON and WDM upgraded RPR networks. The resultant Ethernet-based optical access MAN, which we call STARGATE, is described at length and evaluated by means of simulation.

The remainder of the article is organized as follows. We describe WDM upgraded EPON and RPR networks, respectively. We explain STARGATE in greater detail. Results are then provided, and we conclude the article.

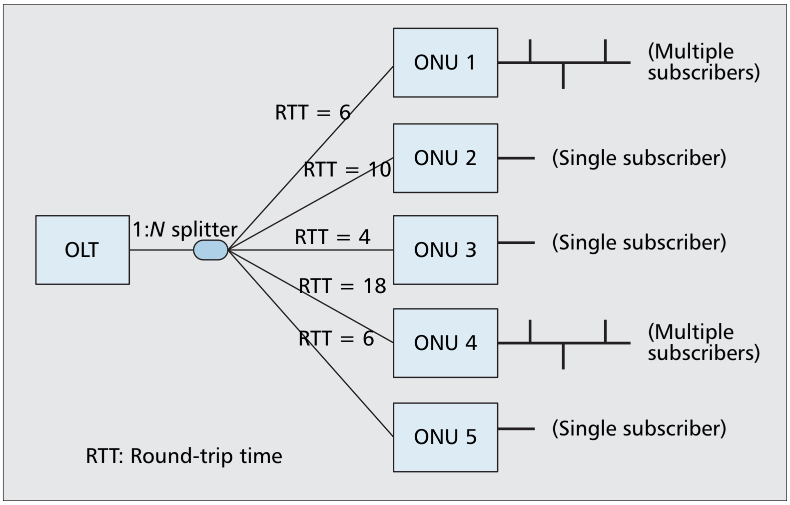

Figure 1. *EPON architecture with one optical line terminal (OLT) and* N = *5 optical network units (ONUs), each with a different round-trip time.*

WDM EPON

EPONs are single-channel TDM systems deploying a single downstream wavelength channel and a single upstream wavelength channel. Typically, EPONs have a physical tree topology with the central office (CO) located at the root and the subscribers connected to the leaf nodes of the tree, the ONUs. As shown in Fig. 1, at the root of the tree resides an optical line terminal (OLT), which is collocated with the CO. The EPON connects the OLT to *N* ONUs through a 1:*N* optical splitter/combiner. Each ONU buffers data received from its attached subscriber(s). In general, the round-trip time (RTT) between an OLT and each ONU is different. Due to the directional properties of the optical splitter/combiner, the OLT is able to broadcast data to all ONUs in the downstream direction. In the upstream direction, however, ONUs cannot communicate directly with one another.

To avoid collisions and increase bandwidth utilization in the upstream direction, the OLT dynamically allocates a variable time slot to each ONU based on the instantaneous bandwidth demands of the ONUs. To facilitate dynamic bandwidth allocation (DBA) and arbitrate the upstream transmission of multiple ONUs, the so-called MultiPoint Control Protocol (MPCP) is deployed in EPONs. Besides auto-discovery, registration, and ranging (RTT computation) operations for newly added ONUs, MPCP provides the signaling infrastructure (control plane) for coordinating the data transmission from the ONUs to the OLT. MPCP uses two types of messages to facilitate arbitration: REPORT and GATE. Each ONU has a set of queues, possibly prioritized, holding Ethernet frames ready for upstream transmission to the OLT. The REPORT message is used by an ONU to report bandwidth requirements (typically in the form of queue occupancies) to the OLT. Upon receiving a REPORT message, the OLT passes it to the DBA algorithm module. The DBA module calculates the upstream transmission schedule of all ONUs such that channel collisions are avoided, bandwidth utilization is increased, and QoS requirements are met. After executing the DBA algorithm, the OLT transmits GATE messages to issue transmission grants. Each transmission grant contains the transmission start time and transmission length of the corresponding ONU. Note that MPCP does not specify any particular DBA algorithm. MPCP simply provides a framework for the implementation of various DBA algorithms. We refer the interested reader to [5] for further information on single-channel TDM EPONs and a survey of DBA algorithms.

In [6] we described an evolutionary WDM upgrade of EPONs. The proposed future-proof approach does not impose any particular WDM architecture on the ONU, thus allowing these decisions to be dictated by economics, state-of-the-art transceiver manufacturing technology, traffic demands, and service provider preferences. The proposed WDM upgrade approach allows for cautious pay-as-you-grow WDM upgrades of individual ONUs and thus helps operators realize their survival strategy for highly cost-sensitive access networks. The recommended WDM extensions to MPCP guarantee compliance with IEEE 802.3ah and enable the OLT

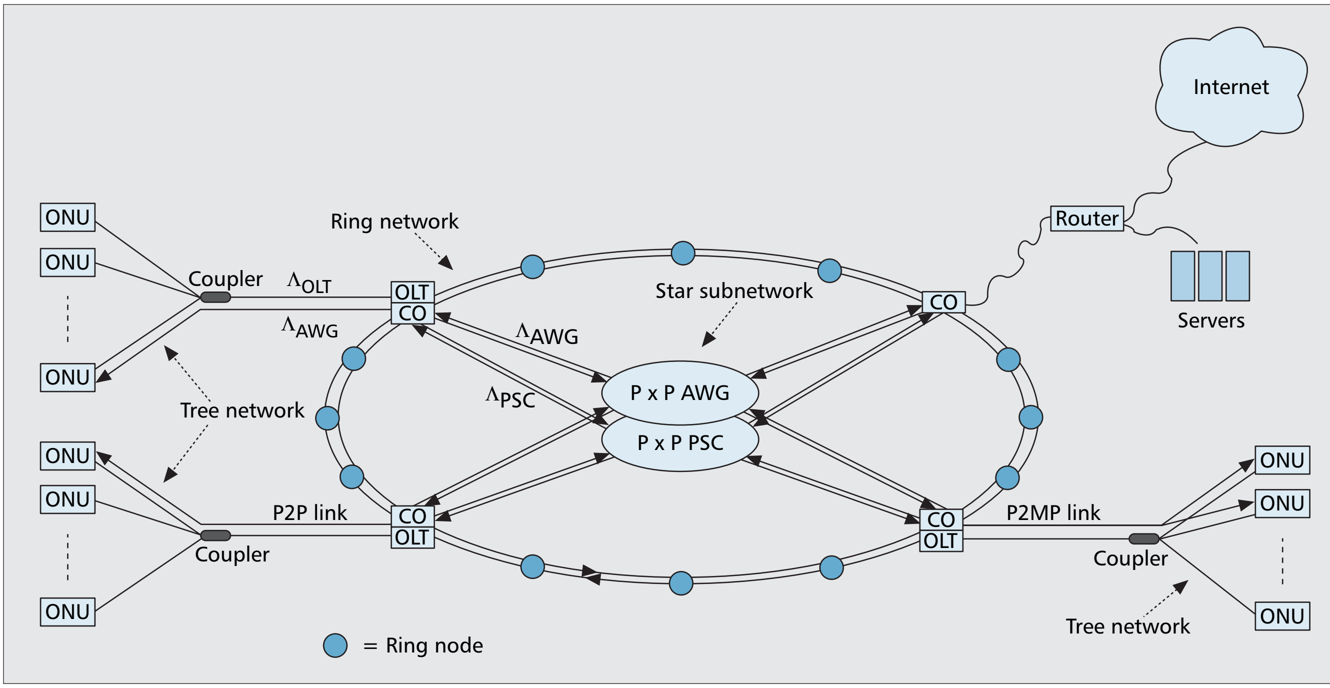

Figure 2. *STARGATE network architecture comprising* P = *4 central offices (COs) and* N_r = *12 RPR ring nodes.*

to schedule transmissions to and receptions from ONUs on any wavelength channel(s) supported by both the OLT and the respective ONU. The discovery and registration of ONUs take place on the original wavelength channel of TDM EPONs, resulting in backward compatibility. We refer the interested reader to [6] for further technical details and additional information on DBA algorithms for WDM EPONs.

WDM RPR

RPR is an optical dual-fiber bidirectional ring network where each fiber ring carries a single wavelength channel. Destination stripping in conjunction with shortest path routing is deployed to improve the spatial reuse of bandwidth. Each node is equipped with two fixed-tuned transmitters and two fixed-tuned receivers, one for each fiber ring. Each node has separate (electrical) transit and station queues for either ring. In RPR, in-transit ring traffic is given priority over station traffic such that in-transit packets are not lost due to buffer overflow. Furthermore, RPR provides resilience against any single link or node failure. For more detailed information on RPR the interested reader is referred to [4].

In [7] we described a novel evolutionary WDM upgrade of RPR that builds on its node architecture and protocols. In the proposed WDM upgrade approach, only a subset of RPR nodes need to be WDM upgraded and interconnected by a dark fiber star WDM subnetwork in a pay-as-you-grow manner according to given traffic demands and cost constraints, as opposed to conventional WDM-upgraded rings that require all RPR nodes to be WDM upgraded at the same time. Unlike WDM rings, the resultant hybrid ring-star network improves the resilience, spatial reuse, and bandwidth efficiency of RPR dramatically. We refer the interested reader to [7] for more detailed information on the novel evolutionary WDM upgrade approach and its underlying techniques.

STARGATE

RPR can easily bridge to Ethernet networks such as EPON and may also span into metropolitan area networks (MANs) and wide area networks (WANs). This makes it possible to perform layer 2 switching from access networks far into backbone networks [4]. It remains to be seen whether end-to-end Ethernet networks turn out to be practical. From an all-optical integration point of view, however, end-to-end optical islands of transparency are not feasible, and are expected to be of limited geographical coverage due to physical transmission impairments as well as other issues such as management, jurisdiction, and billing issues. As a matter of fact, islands of transparency with optical bypassing capability are key in MANs in order to support not only various legacy but also future services in an easy and cost-effective manner.

The proposed STARGATE architecture all-optically integrates Ethernet based access and metro networks. The rationale behind STARGATE is based on the following three principles.

Evolutionary downstream SDM upgrades: Eventually, when per-user bandwidth needs grow, incrementally upgrading existent EPON tree networks with additional fibers may prove attractive or even become mandatory. In fact, some providers are already finding this option attractive since long runs of multifiber cable are almost as economical in both material and installation costs as the same lengths of cables with one or a few fibers [1]. Interestingly, the standard IEEE 802.3ah supports not only the commonly used point-to-multipoint (P2MP) topology but also a hybrid EPON topology consisting of point-to-point (P2P) links in conjunction with P2MP links. In STARGATE we explore the merits of deploying an additional P2P or P2MP fiber link in EPON tree networks to connect the OLT with a subset of one or more ONUs in an evolutionary fashion according to given traffic demands and/or cost constraints. It is important to note, that STARGATE requires an additional P2P or P2MP fiber link only in the downstream direction from OLT to ONU(s) and none in the upstream direction. Thus, STARGATE makes use of evolutionary downstream space-division multiplexing (SDM) upgrades of WDM/TDM EPONs.

Optical bypassing: The problem with using SDM in EPONs is the increased electro-optic port count at the OLT. To avoid this, STARGATE makes use of optical bypassing. Specifically, all wavelengths on the aforementioned additional P2P or P2MP downstream fiber link coming from the metro edge ring are not terminated at the OLT, thus avoiding the need for optical-electronic-optical (OEO) conversion and additional transceivers at the OLT, as explained in greater detail shortly. Note that OEO conversion usually represents the major part of today's optical networking infrastructure costs. Due to the small to moderate distances of STARGATE's access-metro networks, optical bypassing and the resultant transparency can easily be implemented, thereby avoiding OEO conversion and resulting in major cost savings.

Passive optical networking: Finally, the last design principle of STARGATE is based on the idea of letting low-cost passive optical networking technologies follow low-cost Ethernet technologies from access networks into metro networks. In doing so, not only PONs but also MANs benefit from passivity, which is a powerful tool to build low-cost high-performance optical networks [8]. As we will see shortly, STARGATE makes use of an athermal (temperature-insensitive) arrayed waveguide grating (AWG) wavelength router, which eliminates the need for temperature control and monitoring the wavelength shift of the AWG, and thus leads to simplified network management and reduced costs.

We note that passive optical networking in all-optical wavelength-routing WDM networks has recently begun to gain momentum within the so-called time-domain wavelength interleaved networking (TWIN) concept that enables cost-effective and flexible optical networks to be built using readily available components [9]. In TWIN, fast TDM switching and packet switching in the passive optical wavelength-selective WDM network core are emulated through the use of emerging fast tunable lasers at the optical network edge, thus avoiding the need for fast optical switching and optical buffering. The original TWIN did not scale well because the number of nodes N was limited by the number of available wavelengths W (i.e., $N = W$). Very recently, TWIN with wavelength reuse (TWIN-WR) was proposed, where the number of nodes is independent of the number of wavelengths (i.e., $N > W$ [10]). Both TWIN and TWIN-WR require network-wide scheduling of transmissions in order to avoid channel collisions. Unlike in TWIN, a source node in TWIN-WR may not be able to send traffic directly to any destination node in an optical single hop, resulting in multihopping via intermediate electrical gateways. As we will see shortly, STARGATE differs from TWIN and TWIN-WR in a number of ways. First, STARGATE supports extensive wavelength reuse while providing optical *single-hop* communication among all ONUs. Second, STARGATE requires only *local* scheduling of transmissions within each separate EPON in order to completely avoid channel collisions throughout the network, thus avoiding the need for network-wide scheduling. Third, STARGATE does not require any time-of-day synchronization. Finally, while TWIN and TWIN-WR are designed to support WANs of arbitrary topology, STARGATE targets access and metro networks, whose *regular* topologies (tree, ring, star) help simplify scheduling significantly.

In the following, the STARGATE architecture and operation are described in greater detail.

STARGATE ARCHITECTURE

The network architecture of STARGATE is shown in Fig. 2. STARGATE consists of an RPR metro edge ring that interconnects multiple WDM EPON tree networks among each other as well as to the Internet and server farms. The RPR network consists of P COs and N_r RPR ring nodes, and its RTT equals τ_{ring}. The CO in the upper right corner of the figure is assumed to be attached to the Internet and a number of servers via a common router. We refer to this CO as the *hotspot* CO. The P COs are interconnected via a single-hop WDM star subnetwork whose hub is based on a wavelength-broadcasting $P \times P$ passive star coupler (PSC) in parallel with an athermal wavelength-routing $P \times P$ AWG. The end-to-end propagation delay of the star subnetwork equals τ_{star}. Each CO is attached to a separate input/output port of the AWG and PSC by means of two pairs of counterdirectional fiber links. Each fiber going to and coming from the AWG carries $\Lambda_{\text{AWG}} = P$ wavelength channels. Each fiber going to and coming from the PSC carries $\Lambda_{\text{PSC}} = 1 + H + (P - 1)$ wavelength channels, consisting of one control channel λ_c, $1 \leq H\ P - 1$ dedicated home channels for the hotspot CO, and $(P - 1)$ dedicated home channels, one for each of the remaining $(P - 1)$ COs. The home channels are fixedly assigned to the COs. Data destined for a certain CO is sent on its corresponding home channel. All COs except the hotspot CO are collocated with a separate OLT of an attached WDM EPON. Let Λ_{OLT} denote the number of used wavelengths in each WDM EPON in both directions between the ONUs and the corresponding OLT; that is, there is one set of Λ_{OLT} upstream wavelength channels and another set of Λ_{OLT} downstream wavelength channels in each WDM EPON. Furthermore, each WDM EPON deploys an additional P2P or P2MP downstream fiber link from the CO to a single ONU or multiple ONUs, respectively. Each downstream fiber link carries the Λ_{AWG} wavelength channels coming from the AWG of the star subnetwork.

In TWIN, fast TDM switching and packet switching in the passive optical wavelength-selective WDM network core are emulated through the use of emerging fast tunable lasers at the optical network edge.

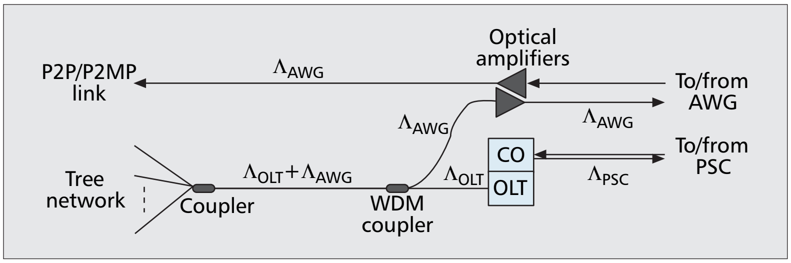

■ **Figure 3.** *Optical bypassing of optical line terminal and central office.*

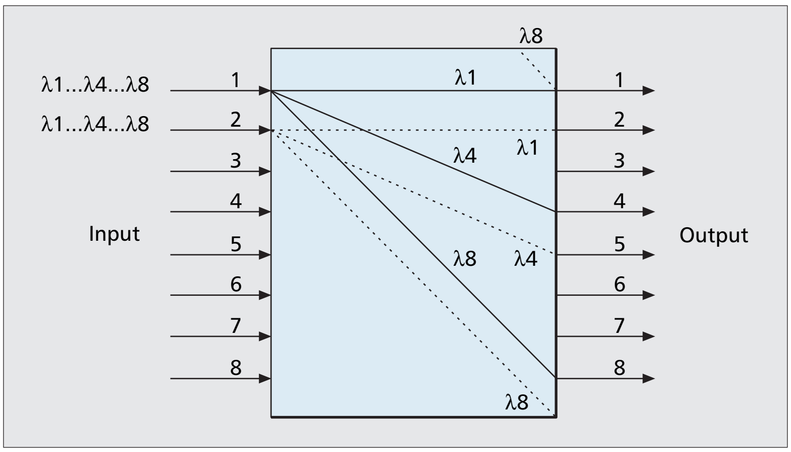

■ **Figure 4.** *Wavelength routing of an 8 × 8 arrayed waveguide grating.*

Figure 3 depicts the interconnection of a given WDM EPON and the star subnetwork in greater detail, illustrating the optical bypassing of the collocated OLT and CO. Note that in the figure, Λ_{OLT} comprises both upstream and downstream wavelength channels that run in opposite directions on the tree network to and from the OLT, respectively. In contrast, the Λ_{AWG} wavelength channels are carried on the tree network only in the upstream direction, while in the downstream direction they are carried on the separate P2P or P2MP downstream fiber link. As shown in Fig. 3, a WDM coupler is used on the tree network in front of the OLT to separate the Λ_{AWG} wavelength channels from the Λ_{OLT} wavelength channels and to guide them directly onward to the AWG of the star subnetwork, optically amplified if necessary. In doing so, the Λ_{AWG} wavelength channels are able to optically bypass the CO and OLT. Similarly, the Λ_{AWG} wavelength channels coming from the AWG optically bypass both CO and OLT and directly travel on the P2P or P2MP link onward to the subset of attached ONUs. As a result, the ONU(s) as well as the hotspot CO that send and receive on any of the Λ_{AWG} wavelength channels are able to communicate all-optically with each other in a single hop across the AWG of the star subnetwork. In other words, the star forms a gate for all-optically interconnecting multiple WDM EPONs. Accordingly, we call the network STARGATE.

The wavelength-routing AWG allows for the spatial reuse of all Λ_{AWG} wavelength channels at each AWG port, as shown in Fig. 4 for an 8 × 8 AWG ($P = 8$) and $\Lambda_{AWG} = P = 8$ wavelengths. Observe that the same wavelength can be simultaneously deployed at two (and more) AWG input ports without resulting channel collisions at the AWG output ports. For instance, wavelength λ_4 incident on input ports 1 and 2 is routed to different output ports 4 and 5, respectively. The wavelength-routing characteristics of the AWG have the following implications: First, due to the fact that the AWG routes wavelengths arriving at a given input port independent from all other AWG input ports, no network-wide scheduling but only local scheduling at each AWG input port is necessary to avoid channel collisions on the AWG. Second, note that in Fig. 4 each AWG input port reaches a given AWG output port on a different wavelength channel. Consequently, under full spatial wavelength reuse, Λ_{AWG} different wavelengths arrive at each AWG output port simultaneously. To avoid receiver collisions (destination conflicts), each AWG output port must be equipped with a receiver operating on all Λ_{AWG} wavelengths. A receiver collision occurs if none of the destination node's receivers is tuned to the wavelength on which data arrives.

Similar to an RPR node, each CO is equipped with two fixed-tuned transceivers, one for each direction of the dual fiber ring. In addition, each CO has one transceiver fixedly tuned to the control channel λ_c of the star subnetwork. For data reception on its PSC home channel, each CO (except the hotspot CO) has a single fixed-tuned receiver. The hotspot CO is equipped with $1 \leq H \leq P - 1$ fixed-tuned receivers. For data transmission on the PSC, each CO (except the hotspot CO) deploys a single transmitter that can be tuned over the $(P - 1) + H$ home channels of the COs. The hotspot CO deploys H tunable transmitters whose tuning range covers the home channels of the remaining $(P - 1)$ COs as well as the Λ_{AWG} wavelengths. Unlike the remaining COs, the hotspot CO is equipped with an additional multiwavelength receiver operating on Λ_{AWG}. In each WDM EPON, the OLT is equipped with an array of fixed-tuned transmitters and fixed-tuned receivers, operating at the Λ_{OLT} downstream and Λ_{OLT} upstream wavelength channels, respectively. Similar to [6], STARGATE does not impose any particular WDM node structure on the ONUs except for ONUs that receive data over the AWG. Those ONUs must be equipped with a multiwavelength receiver operating on the Λ_{AWG} wavelengths in order to avoid receiver collisions, as explained above.

STARGATE OPERATION

The STARGATE network operates as follows.

Discovery and registration: In each WDM EPON, the IEEE 802.3ah REGISTER REQ MPCP message with WDM extensions described in [6] is deployed for the discovery and registration of ONUs. The REGISTER REQ message is sent from each ONU to its OLT and carries the MAC address as well as detailed information about the WDM node structure of the ONU. In doing so, the OLT of each WDM EPON learns about the MAC address and WDM node structure of each of its attached ONUs. After registration, the OLTs exchange via the PSC (to be described shortly) the MAC addresses of their attached ONUs that are able to receive data over the AWG. As a result, the OLTs know not only which MAC addresses can be reached via the AWG but also to which AWG output ports the corresponding ONUs are attached, and thus on which of the Λ_{AWG} wavelengths they can be reached from a given AWG input port.

Piggyback REPORT MPCP message: The IEEE 802.3ah REPORT MPCP message can carry one or more queue sets, each set comprising up to eight queues, as shown in Fig. 5. In STARGATE, we let ONUs use the first queue set to report bandwidth requirements on the Λ_{OLT} upstream wavelengths for sending data to the OLT. To report bandwidth requirements on any of the Λ_{AWG} wavelength channels to an ONU located in a different EPON, a given ONU uses an additional queue set and writes the MAC address of the destination ONU in the reserved field of the REPORT MPCP message and sends it to the OLT. Thus, the bandwidth requirements on Λ_{AWG} ride piggyback on those on Λ_{OLT} within the same REPORT message.

STARGATE MPCP message: The WDM extended IEEE 802.3ah GATE message in [6] is used to coordinate the upstream transmission on the Λ_{OLT} wavelengths within each WDM EPON and also to coordinate the all-optical transmission on any of the Λ_{AWG} wavelengths across the star subnetwork between two ONUs residing in different WDM EPONs, provided both ONUs support the Λ_{AWG} wavelengths. Based on the MAC address of the destination ONU carried piggyback in the REPORT message, the OLT of the source WDM EPON uses the extended GATE MPCP message, which we call the STARGATE message, to grant the source ONU a time window on that wavelength which the AWG routes to the destination ONU according to the DBA algorithm in use at the OLT.

STARGATING service: Similar to an EPON, the STARGATE network is not restricted to any specific DBA algorithm. However, DBA algorithms for STARGATE should be able to dynamically set up transparent all-optical circuits across the AWG at the wavelength and subwavelength granularity with predictable QoS in terms of bounded delay and guaranteed bandwidth between ONUs in different WDM EPONs. Each OLT uses its DBA module to provide gated service across the AWG-based star network, a service we correspondingly call STARGATING.

Access control on ring and PSC: ONUs unable to send and receive data across the AWG as well as RPR ring nodes send their data on the tree, ring, and/or PSC along the shortest path in terms of hops. Channel access on the dual fiber ring is governed by RPR protocols, described earlier. On the PSC, time is divided into periodically recurring frames. Each frame on control channel λ_c consists of P control slots, each dedicated to a different CO. Each CO stores data packets to be forwarded on the PSC. For each stored data packet the CO broadcasts a control packet on the PSC to all COs in its assigned control slot. A control packet consists of two fields:

- Destination address
- Length of the corresponding data packet

After τ_{star}, all COs receive the control packet and build a common distributed transmission schedule for collision-free transmission of the corresponding data packet on the home channel of the destination CO at the earliest possible time. The CO forwards the received data packet toward the final destination node.

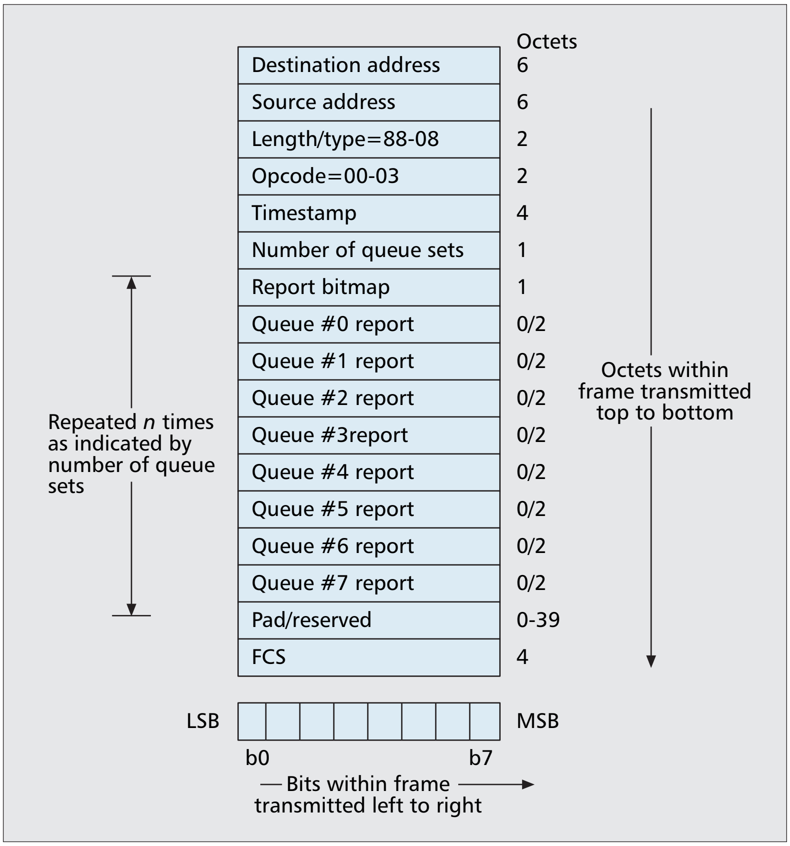

Figure 5. *REPORT MPCP message.*

Results

In our illustrative simulation example, we set the STARGATE parameters as follows: $P = 4$, $N_r = 12$, $H = 1$, $\Lambda_{PSC} = 5$, and $\Lambda_{AWG} = \Lambda_{OLT} = 4$. Each WDM EPON is assumed to have 32 ONUs, each located 20 km away from its OLT ($\tau_{tree} = 0.1$ ms). The perimeter of the RPR ring is assumed to be 100 km ($\tau_{ring} = 0.5$ ms, $\tau_{star} = 0.16$ ms). The tree networks and ring network operate at 1 Gb/s, while the star subnetwork operates at 10 Gb/s. We assume that ONUs generate Poisson traffic with the following trimodal packet length distribution: 50 percent 40-byte packets, 30 percent 552-byte packets, and 20 percent 1500-byte packets. A packet generated by a given ONU is uniformly destined for any of the remaining ONUs or the hotspot CO. The hotspot CO is assumed to send and receive traffic equivalent to the total traffic of a tree network (i.e., 32 ONUs).

In Fig. 6 we consider two different types of ONUs. ONUs of type 1 are equipped with a single transmitter that can be tuned over Λ_{OLT}; ONUs of type 2 use a single transmitter tunable over both Λ_{OLT} and Λ_{AWG}. As DBA we use limited service with a maximum transmission window size of $W_{max} = 5000$ bytes/ONU. The

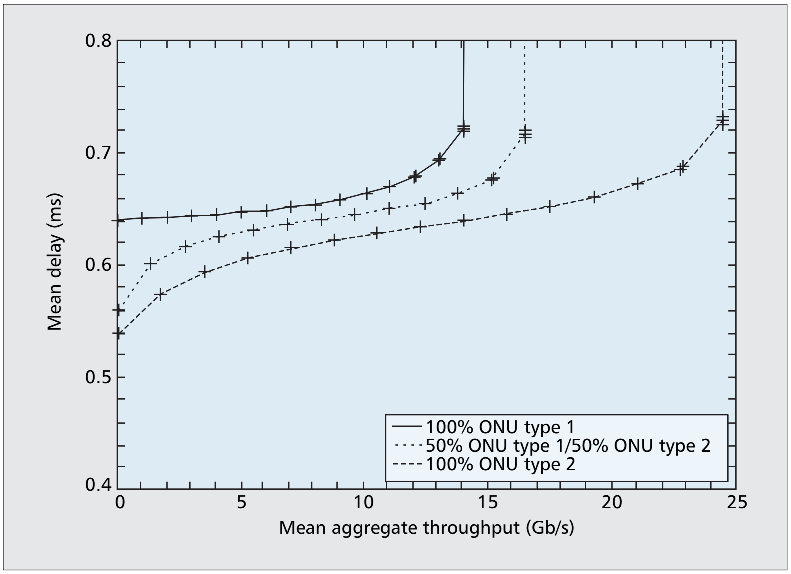

■ **Figure 6.** *Mean delay vs. mean aggregate throughput of STARGATE for limited service.*

limited service DBA assigns the requested bandwidth, but not more than W_{max}, to ONUs of type 1 on a first come first served basis on any of the Λ_{OLT} wavelengths at the earliest possible time. For ONUs of type 2, the limited service DBA schedules eligible traffic on Λ_{AWG} and the remaining traffic on Λ_{OLT} at the earliest possible time. Figure 6 depicts the mean delay (in milliseconds) vs. mean aggregate throughput (in gigabits per second) of STARGATE together with 95 percent confidence intervals for different combinations of ONU types 1 and 2. We observe that by providing more ONUs with access to the optical bypassing wavelengths Λ_{AWG}, the throughput-delay performance of STARGATE is improved significantly, especially at low and high traffic loads.

Conclusions

STARGATE lets low-cost PON technologies follow low-cost Ethernet technologies from access networks into metro networks. By adding the space dimension and optical bypassing to conventional WDM/TDM EPONs, STARGATE provides a cost-effective all-optical WDM integration of Ethernet-based access and metro networks. The proposed STARGATING service allows dynamic setup of fine-granularity transparent connections between ONUs in different WDM EPONs in support of emerging and future applications (e.g., P2P file sharing). Exciting future research avenues include the design and evaluation of DBA algorithms for STARGATE.

References

[1] P. E. Green, *Fiber To The Home-The New Empowerment*, Wiley, 2006.

[2] H. Shinohara, "Broadband Access in Japan: Rapidly Growing FTTH Market," *IEEE Commun. Mag.*, vol. 43, no. 9, Sept. 2005, pp. 72–78.

[3] T. Koonen, "Fiber to the Home/Fiber to the Premises: What, Where, and When," *Proc. IEEE*, vol. 94, no. 5, May 2006, pp. 911–34.

[4] F. Davik *et al.*, "IEEE 802.17 Resilient Packet Ring Tutorial," *IEEE Commun. Mag.*, vol. 42, no. 3, Mar. 2004, pp. 112–18.

[5] M. P. McGarry, M. Maier, and M. Reisslein, "Ethernet PONs: A Survey of Dynamic Bandwidth Allocation (DBA) Algorithms," *IEEE Commun. Mag.*, vol. 42, no. 8, Aug. 2004, pp. S8–S15.

[6] —, "WDM Ethernet Passive Optical Networks," *IEEE Commun. Mag.*, vol. 44, no. 2, Feb. 2006, pp. S18–S25.

[7] M. Herzog and M. Maier, "RINGOSTAR: An Evolutionary Performance-Enhancing WDM Upgrade of IEEE 802.17 Resilient Packet Ring," *IEEE Commun. Mag.*, vol. 44, no. 2, Feb. 2006, pp. S11–S17.

[8] P. E. Green, "Paving the Last Mile With Glass," *IEEE Spectrum*, vol. 39, no. 12, Dec. 2002, pp. 13–14.

[9] I. Widjaja *et al.*, "Light Core and Intelligent Edge for a Flexible, Thin-Layered, and Cost-Effective Optical Transport Network," *IEEE Commun. Mag.*, vol. 41, no. 5, May 2003, pp. S30–S36.

[10] C. Nuzman and I. Widjaja, "Time-Domain Wavelength Interleaved Networking with Wavelength Reuse," *Proc. IEEE INFOCOM*, Apr. 2006.

Biographies

MARTIN MAIER (maier@ieee.org) is an associate professor at INRS. He was educated at the Technical University of Berlin, and received Dipl.-Ing. and Dr.-Ing. degrees (both with distinctions) in 1998 and 2003, respectively. In the summer of 2003 he was a postdoc fellow at the Massachusetts Institute of Technology (MIT), Cambridge. He was a visiting professor at Stanford University, California, October 2006 through March 2007. Currently, his research activities focus on evolutionary WDM upgrades of optical access and metro networks. He is the author of the book *Metropolitan Area WDM Networks — An AWG Based Approach* (Springer, 2003).

MARTIN HERZOG (herzog@emt.inrs.ca) is a postdoctoral researcher at the Institut National de la Recherche Scientifique (INRS), Montreal, Canada. He received Dipl.-Ing. and Dr.-Ing. degrees from the Technical University of Berlin, Germany, in 2002 and 2006, respectively. His current research focuses on architectures and protocols for optical access and metropolitan area networks, and switching techniques for wide area networks.

MARTIN REISSLEIN (reisslein@asu.edu) is an associate professor in the Department of Electrical Engineering at Arizona State University, Tempe. He received his Ph.D. in systems engineering from the University of Pennsylvania in 1998. From July 1998 through October 2000 he was a scientist with the German National Research Center for Information Technology (GMD FOKUS), Berlin, and a lecturer at the Technical University of Berlin. 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.