Ethernet Passive Optical Network (EPON): Building a Next-Generation Optical Access Network

Glen Kramer and Gerry Pesavento, Alloptic, Inc.

ABSTRACT

This article describes Ethernet passive optical networks, an emerging local subscriber access architecture that combines low-cost point-to-multipoint fiber infrastructure with Ethernet. EPONs are designed to carry Ethernet frames at standard Ethernet rates. An EPON uses a single trunk fiber that extends from a central office to a passive optical splitter, which then fans out to multiple optical drop fibers connected to subscriber nodes. Other than the end terminating equipment, no component in the network requires electrical power, hence the term passive. Local carriers have long been interested in passive optical networks for the benefits they offer: minimal fiber infrastructure and no powering requirement in the outside plant. With Ethernet now emerging as the protocol of choice for carrying IP traffic in metro and access networks, EPON has emerged as a potential optimized architecture for fiber to the building and fiber to the home.

INTRODUCTION

While in recent years the telecommunications backbone has experienced substantial growth, little has changed in the access network. The tremendous growth of Internet traffic has accentuated the aggravating lag of access network capacity. The "last mile" still remains the bottleneck between high-capacity local area networks (LANs) and the backbone network. The most widely deployed broadband solutions today are digital subscriber line (DSL) and cable modem (CM) networks. Although they are an improvement over 56 kb/s modems, they are unable to provide enough bandwidth for emerging services such as IP telephony, video on demand (VoD), interactive gaming, or two-way videoconferencing. A new technology is required; one that is inexpensive, simple, scalable, and capable of delivering bundled voice, data, and video services to an end-user subscriber over a single network. Ethernet passive optical networks (EPONs), which represent the convergence of low-cost Ethernet equipment and low-cost fiber infrastructure, appear to be the best candidate for the next-generation access network.

EVOLUTION OF THE FIRST MILE

The first mile? Once called the last mile, the Ethernet community has renamed this network section to the first mile, to symbolize its priority and importance. The first mile connects the service provider central offices to business and residential subscribers. Also referred to as the subscriber access network or local loop, it is network infrastructure at the neighborhood level. Residential subscribers demand first mile access solutions that are broadband, offer Internet media-rich services, and are comparable in price to existing networks.

Incumbent telephone companies responded to Internet access demand by deploying DSL technology. DSL uses the same twisted pair as telephony lines and requires a DSL modem at the customer premises and a digital subscriber line access multiplexor (DSLAM) in the central office. The data rate provided by DSL is typically offered in a range of 128 kb/s–1.5 Mb/s. While this is significantly faster than an analog modem, it is well shy of being considered broadband, in that it cannot support full-service voice, data, and video. In addition, the physical area one central office can cover with DSL is limited to distances less than 5.5 km, which constitutes approximately 60 percent of end-user subscribers. As a result, network operators are now deploying remote DSLAMS closer to subscribers; however, in general, service providers do not provide DSL services to subscribers located more than a few miles from a local exchange office due to costs [1].

Cable television companies responded to Internet service demand by integrating data services over their coaxial cable networks, which were originally designed for analog video broadcast. Typically, these hybrid fiber coax (HFC) networks have fiber running between a video

0163-6804/02/$17.00 © 2002 IEEE

head-end or hub to a curbside optical node, with the final drop to the subscriber being coaxial cable, repeaters, and tap couplers. The drawback of this architecture is that each shared optical node has less than 36 Mb/s effective data throughput, which is typically divided between 2000 homes, resulting in frustrating slow speed during peak hours. To alleviate bandwidth bottlenecks, optical fibers, and thus optical nodes, are penetrating deeper into the first mile.

The next wave of local access deployment promises to bring fiber to the building (FTTB) and fiber to the home (FTTH). Unlike previous architectures, where fiber is used as a feeder to shorten the lengths of copper and coaxial networks, these new deployments use optical fiber throughout the access network. New optical fiber network architectures are emerging that are capable of supporting gigabit per second speeds, at costs comparable to DSL and HFC networks.

TRAFFIC GROWTH

Data traffic is increasing at an unprecedented rate. Sustainable data traffic growth rate of over 100 percent per year is observed since 1990. There were periods when a combination of economical and technological factors resulted in even larger growth rates (1000 percent increase per year in 1995 and 1996) [2]. This trend is likely to continue in the future. Simply put, more and more users are getting online, and those who are already online are spending more time online. Market research shows that after upgrading to a broadband connection users spend about 35 percent more time online than before [3]. Voice traffic is also growing, but at a much slower rate of 8 percent annually. According to most analysts, data traffic has already surpassed voice traffic. More and more subscribers telecommute, and require the same network performance as they see on corporate LANs. More services and new applications will become available as bandwidth per user increases (Fig. 1).

Neither DSL nor CMs can keep up with such demand. Both technologies are built on top of existing copper communication infrastructure not optimized for data traffic. In CM networks, only a few RF channels are dedicated to data, while the majority of bandwidth is tied up servicing legacy analog video. DSL copper networks do not allow sufficient data rates at required distances. Most network operators have come to the realization that a new data-centric solution is necessary. Such a technology would be optimized for Internet Protocol (IP) data traffic. The remaining services, such as voice and video, will converge into a digital format, and a true full-service network will emerge.

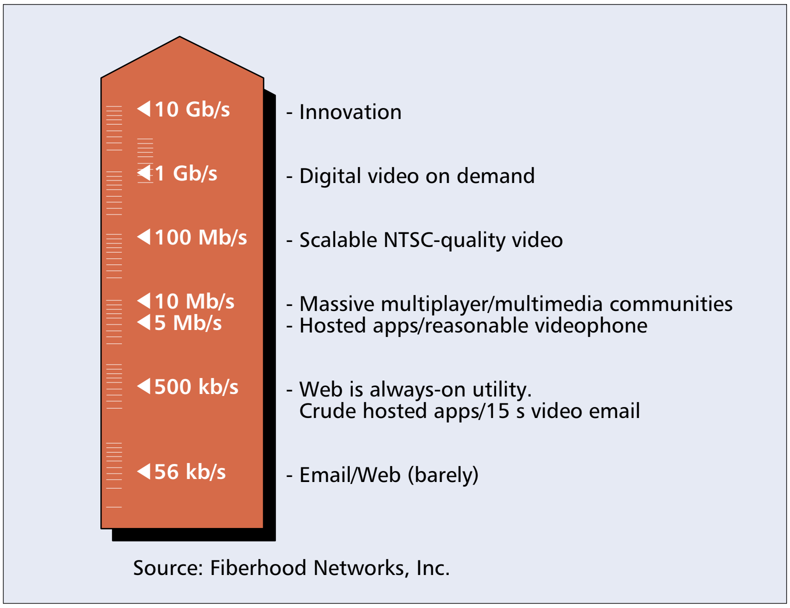

Source: Fiberhood Networks, Inc.

Figure 1. *Per-user bandwidth requirements for new services.*

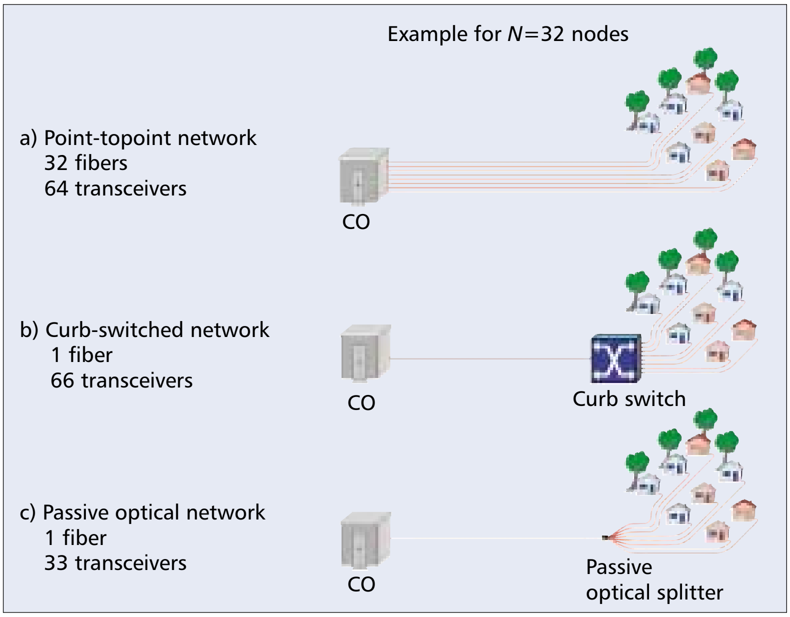

Figure 2. *Fiber to the home (FTTH) deployment scenarios.*

THE NEXT-GENERATION ACCESS NETWORK

Optical fiber is capable of delivering bandwidth-intensive integrated, voice, data, and video services at distances beyond 20 km in the subscriber access network. A logical way to deploy optical fiber in the local access network is using a point-to-point (P2P) topology, with dedicated fiber runs from the local exchange to each end-user subscriber (Fig. 2a). While this is a simple architecture, in most cases it is cost prohibitive due to the fact that it requires significant outside plant fiber deployment as well as connector termination space in the local exchange. Considering N subscribers at an average distance L km from the central office, a P2P design requires $2N$ transceivers and $N * L$ total fiber length (assuming single fiber is used for bidirectional transmission).

To reduce fiber deployment, it is possible to deploy a remote switch (concentrator) close to the neighborhood. This reduces fiber consumption to only L km (assuming negligible distance between the switch and customers), but actually increases the number of transceivers to $2N + 2$, since there is one more link added to the network (Fig. 2b). In addition, a curb-switched architecture requires electrical power as well as backup power at the curb unit. Currently, one of the highest costs for local exchange carriers is providing and maintaining electrical power in the local loop.

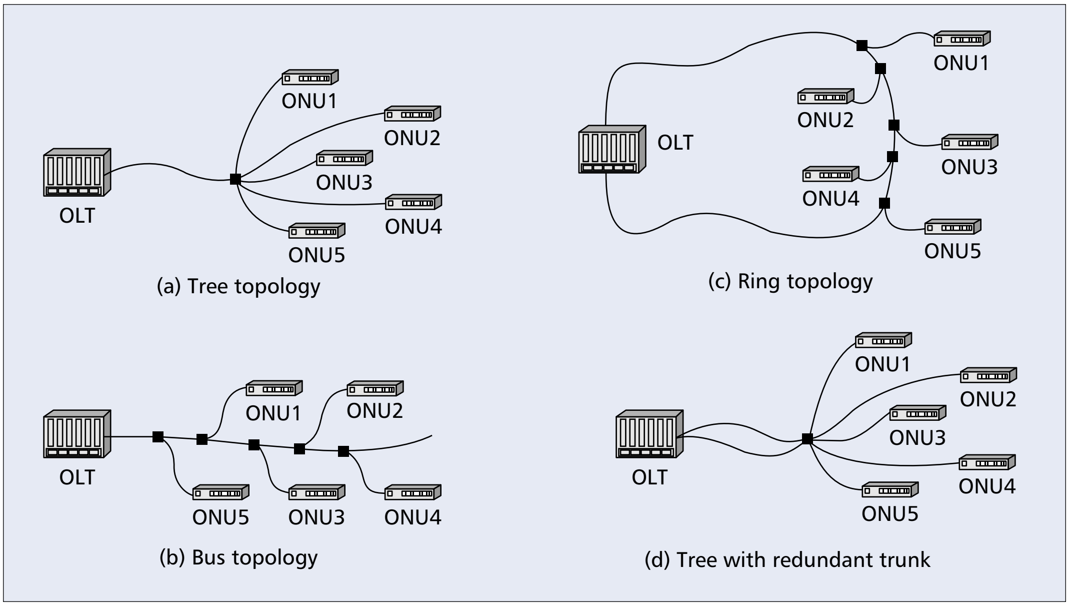

(a) Tree topology

(c) Ring topology

(b) Bus topology

(d) Tree with redundant trunk

Figure 3. *PON topologies.*

Logically, the first mile is point-to-multipoint, with a Central Office typically servicing thousands of subscribers. There are several multipoint topologies suitable for the access network, including tree, tree-and-branch, ring, or bus.

Therefore, it is logical to replace the hardened active curb-side switch with an inexpensive passive optical component. Passive optical network (PON) is a technology viewed by many as an attractive solution to the first mile problem [4, 5]; PONs minimize the amount of optical transceivers, central office terminations, and fiber deployment. A PON is a point-to-multipoint optical network with no active elements in the signals' path from source to destination. The only interior elements used in a PON are passive optical components, such as optical fiber, splices, and splitters. Access networks based on single-fiber PON only require $N + 1$ transceivers and L km of fiber (Fig. 2c).

PON TOPOLOGIES

Logically, the first mile is point-to-multipoint (P2MP), with a central office typically servicing thousands of subscribers. There are several multipoint topologies suitable for the access network, including tree, tree-and-branch, ring, and bus (Fig. 3). Using 1:2 optical tap couplers and 1:N optical splitters, PONs can be flexibly deployed in any of these topologies. In addition, PONs can be deployed in redundant configurations such as double rings or double trees; or redundancy may be added only to a part of the PON, say the trunk of the tree (Fig. 3d).

All transmissions in a PON are performed between an optical line terminal (OLT) and optical network units (ONUs). The OLT resides in the local exchange (central office), connecting the optical access network to the metro backbone. The ONU is located at either the curb (FTTC solution) or the end-user location (FTTH and FTTB), and provides broadband voice, data, and video services. In the downstream direction (from OLT to ONUs), a PON is a P2MP network, and in the upstream direction it is a multipoint-to-point network.

The advantages of using PONs in subscriber access networks are numerous:

- PONs allow for long reach between central offices and customer premises, operating at distances over 20 km.
- PONs minimizes fiber deployment in both the local exchange office and the local loop.
- PONs provides higher bandwidth due to deeper fiber penetration, offering gigabit per second solutions.
- Operating in the downstream as a broadcast network, PONs allow for video broadcasting as either IP video or analog video using a separate wavelength overlay.
- PONs eliminate the necessity to install active multiplexers at splitting locations, thus relieving network operators of the gruesome task of maintaining active curb-side units and providing power to them. Instead of active devices in these locations, PONs use small passive optical splitters, located in splice trays and deployed as part of the optical fiber cable plant.
- Being optically transparent end to end, PONs allow upgrades to higher bit rates or additional wavelengths.

APON TO EPON

Passive optical networking has been considered for the access network for quite some time, even well before the Internet spurred bandwidth demand. The Full Service Access Network (FSAN) Recommendation (ITU G.983) defines a PON-based optical access network that uses asynchronous transfer mode (ATM) as its layer 2 protocol. In 1995, when the FSAN initiative started, ATM had high hopes of becoming the prevalent technology in the LAN, MAN, and backbone. However, since that time, Ethernet technology has leapfrogged ATM. Ethernet has become a universally accepted standard, with over 320 million port deployments worldwide, offering staggering economies of scale [6]. High-speed Gigabit Ethernet deployment is widely accelerating and 10 Gigabit Ethernet products are available. Ethernet, which is easy to scale and manage, is win-

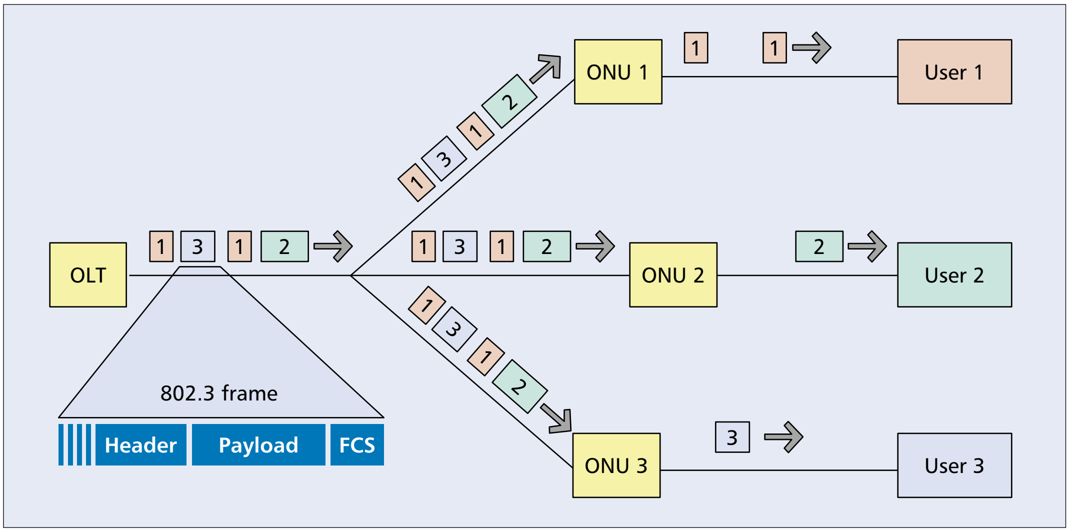

■ **Figure 4.** *Downstream traffic in EPON.*

Newly adopted Quality of Service (QoS) techniques have made Ethernet networks capable of supporting voice, data and video. Ethernet is an inexpensive technology, that is ubiquitous and interoperable with a variety of legacy equipment.

ning new ground in MANs and WANs. Considering that 95 percent of LANs use Ethernet, ATM PON may not be the best choice to interconnect two Ethernet networks.

One of ATM's shortcomings is the fact that a dropped or corrupted cell will invalidate the entire IP datagram. However, the remaining cells carrying the portions of the same IP datagram will propagate further, thus consuming network resources unnecessarily. Also, ATM imposes a cell tax on variable-length IP packets. For example, for the trimodal packet size distribution reported in [7], the cell tax is approximately 13 percent; that is, to send the same amount of user's data an ATM network must transmit 13 percent more bytes than an Ethernet network (counting 64-bit preamble and 96-bit IPG in Ethernet and 12 bytes of overhead associated with AAL-5 in ATM). Finally, perhaps most important, ATM did not live up to its promise of becoming an inexpensive technology; vendors are in decline and manufacturing volumes relatively low. ATM switches and network cards are significantly (roughly 8 ×) more expensive than Ethernet switches and network cards [6].

On the other hand, Ethernet looks like a logical choice for an IP data-optimized access network. An Ethernet PON (EPON) is a PON that carries all data encapsulated in Ethernet frames. 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). Ethernet is an inexpensive technology that is ubiquitous and interoperable with a variety of legacy equipment.

An EPON Network

The IEEE 802.3 standard defines two basic configurations for an Ethernet network. In one case it can be deployed over a shared medium using carrier sense multiple access with collision detection (CSMA/CD) protocol. In another case stations may be connected through a switch using full-duplex links. Properties of an EPON are such that it cannot be considered either shared medium or a point-to-point network; rather, it is a combination of both.

In the downstream direction, Ethernet frames transmitted by OLT pass through a 1:*N* passive splitter and reach each ONU. Splitting ratios are typically between 4 and 64. This behavior is similar to a shared media network. Because Ethernet is broadcasting by nature, in the downstream direction (from network to user), it fits perfectly with the EPON architecture: packets are broadcast by the OLT and extracted by their destination ONU based on the media access control (MAC) address (Fig. 4).

In the upstream direction, due to directional properties of a passive combiner (optical splitter), data frames from any ONU will only reach the OLT, not other ONUs. In that sense, in the upstream the behavior of EPON is similar to that of a P2P architecture. However, unlike in a true P2P network, in EPON frames from different ONUs transmitted simultaneously still may collide. Thus, in the upstream direction (from user to network), the ONUs need to share the trunk fiber channel capacity and resources.

Multiple Access

One possible way of separating the channels is to use wavelength division multiplexing (WDM), in which each ONU operates at a different wavelength. While a simple (from a theoretical perspective) solution, it remains cost prohibitive for an access network. A WDM solution would require either a tunable receiver or a receiver array at the OLT to receive multiple channels. An even more serious problem for network operators would be wavelength-specific ONU inventory: instead of having just one type of ONU, there would be multiple types of ONUs based on their laser wavelength. It would also be more problematic for an unqualified user to replace a defective ONU. Using tunable lasers in ONUs is too expensive at the current state of technology. For these reasons a WDM PON network is not an attractive solution in today's environment.

We believe time-sharing is the preferred method of optical channel sharing in an access network as it allows for a single upstream wavelength, such as 1310 nm, and a single transceiver in the OLT, resulting in a cost-effective solution.

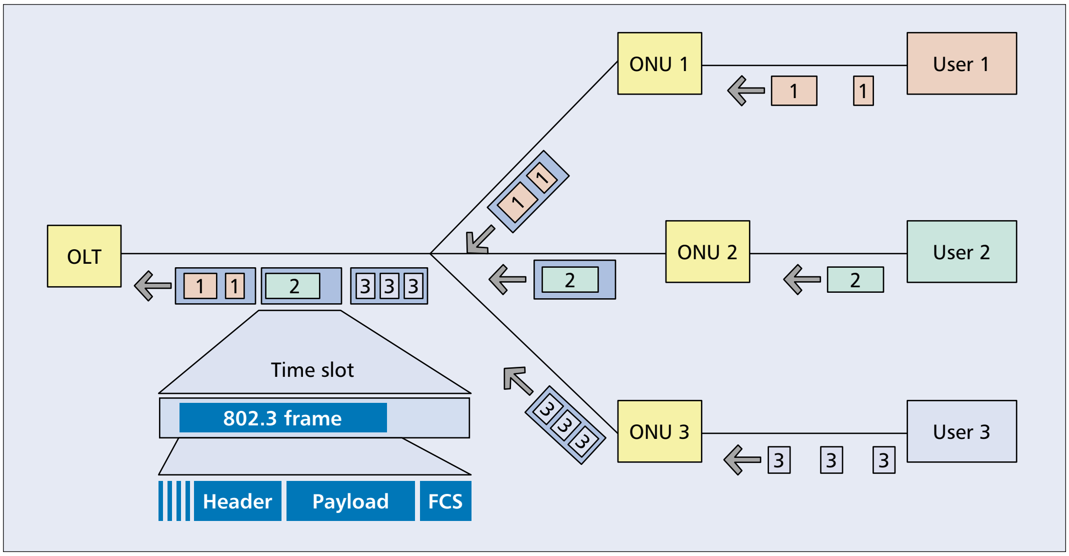

Figure 5. *Upstream traffic in EPON.*

Contention-based media access (something similar to CSMA/CD) is difficult to implement because ONUs cannot detect a collision at the OLT (because of directional properties of optical splitter/combiner). An OLT could detect a collision and inform ONUs by sending a jam signal, however, propagation delays in PON, which can exceed 20 km in length, greatly reduce the efficiency of such a scheme. Contention-based schemes also have a drawback of providing a nondeterministic service; that is, node throughput and channel utilization may be described as statistical averages. On a small scale (comparable to propagation delay) there is no guarantee of a node getting access to the media. It is not a problem for CSMA/CD-based enterprise networks where links are short and typically overprovisioned, and traffic predominantly consists of data. Subscriber access networks, however, in addition to data, must support voice and video services, and thus must provide some guarantees on timely delivery of these traffic types.

We believe time-sharing is the preferred method of optical channel sharing in an access network because it allows for a single upstream wavelength (e.g., 1310 nm) and a single transceiver in the OLT, resulting in a cost-effective solution.

To introduce determinism in the frame delivery, different noncontention schemes were proposed. Figure 5 illustrates an upstream time-shared data flow in EPON.

All ONUs are synchronized to a common time reference, and each ONU is allocated a time slot. Each time slot is capable of carrying several Ethernet frames. An ONU should buffer frames received from a subscriber until its time slot arrives. When its time slot arrives, the ONU would "burst" all stored frames at full channel speed (standard Ethernet rate). If there are no frames in the buffer to fill the entire time slot, idles are transmitted. The possible time slot allocation schemes could range from a static allocation (fixed time-division multiple access, TDMA) to a dynamically adapting scheme based on instantaneous queue size in every ONU (statistical multiplexing scheme). There are more allocation schemes possible, including schemes utilizing traffic priority and QoS, service level agreements (SLAs), and oversubscription ratios.

Decentralized approaches to implementing a dynamic slot assignment scheme are also possible, in which ONUs decide when to send data and for how long. These schemes are somewhat similar to a token ring, except that in this case it is a passive ring. In such a scheme, every ONU, before sending its data, will send a special message announcing how many bytes it is about to send. The ONU scheduled next (say, in round-robin fashion) will monitor the transmission of the previous ONU and time its transmission so that it arrives to the OLT right after transmission from the previous ONU. Thus, there will be no collision, and no bandwidth will be wasted. This scheme is similar to hub polling [8]. However, this scheme has a major limitation: it requires connectivity (communicability) between ONUs. This imposes some constraints on PON topology; namely, the network should be deployed as a ring or a broadcasting star. This requirement is not desirable since:

- It may require more fiber to be deployed.
- Fiber plant with different topology might be already predeployed.

In general, a preferred algorithm shall support any P2MP PON topology.

In an optical access network, we can count only on connectivity from the OLT to every ONU (downstream traffic) and every ONU to the OLT (upstream traffic); this is true for all PON topologies. Therefore, the OLT remains the only device that can arbitrate time-division access to the shared channel.

The challenge of implementing an OLT-based dynamic arbitration scheme is in the fact that the OLT does not know how many bytes of data each ONU has buffered. The burstiness of data traffic precludes a queue occupancy prediction of any reasonable accuracy. If the OLT is to make an accurate time slot assignment, it should know the state of a given ONU exactly.

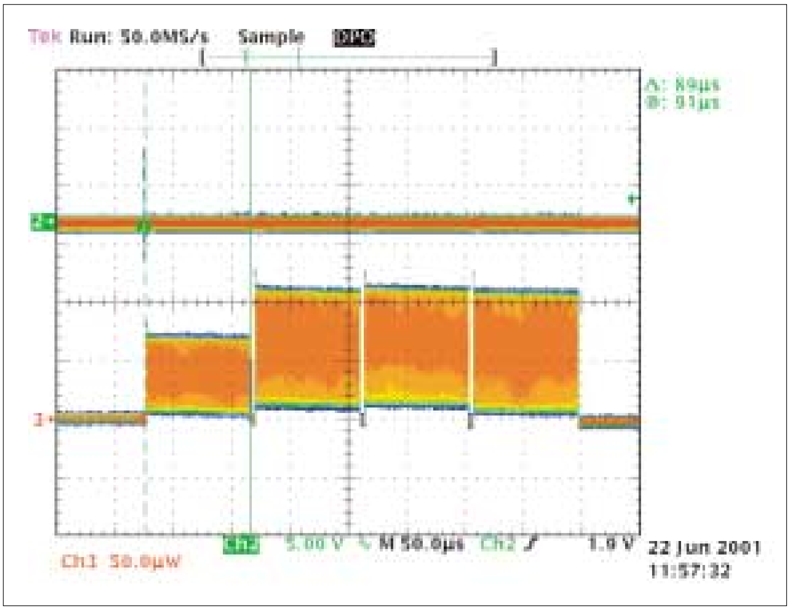

■ **Figure 6.** *An illustration of the near-far problem: a snapshot of received power level from four ONUs.*

One solution may be to use a polling scheme based on Grant and Request messages. Requests are sent from an ONU to report changes in an ONU's state. The OLT processes all Requests and allocates different transmission windows (time slots) to ONUs. Slot assignment information is delivered to ONUs using Grant messages.

The advantage of having centralized intelligence for the slot allocation algorithm is that the OLT knows the state of the entire network and can switch to another allocation scheme based on that information; the ONUs don't need to negotiate or acknowledge new parameters, nor switch to new settings synchronously, making ONUs simpler and cheaper, and the entire network more robust.

Choosing the best allocation scheme, however, is not a trivial task. If all users belong to the same administrative domain (say a corporate or campus network), full statistical multiplexing would make sense; network administrators would like to get the most out of the available bandwidth. However, subscriber access networks are not private LANs, and the objective is to ensure SLA compliance for each individual user. Using statistical multiplexing mechanisms to get each user best effort bandwidth may complicate billing and potentially may offset the user's drive to upgrade to a higher bandwidth. Also, subscribers may get used to and expect the performance they get during low-activity hours when lots of best-effort bandwidth is available. Then, at peak hours, the same users would perceive service as unsatisfactory, even though they get what is guaranteed by their SLA. An optimized bandwidth allocation algorithm will ultimately depend on the future SLA and billing model used by service providers.

TRANSCEIVER ISSUES

Due to unequal distances between central office and ONUs, optical signal attenuation in the PON is not the same for each ONU. The power level received at the OLT will be different for each ONU (the so-called *near-far* problem). Figure 6 depicts power levels received by an OLT from four different ONUs. As shown, one ONU's signal strength is lower at the OLT, most likely due to its longer distance.

To properly detect the incoming bitstream, the OLT receiver must be able to quickly adjust its zero-one threshold at the beginning of each received time slot (i.e., it should operate in *burst mode)*. A burst mode receiver is necessary only in the OLT. The ONUs read a continuous bitstream (data or idles) sent by the OLT and do not need to readjust quickly.

An alternative approach may be to allow ONUs to adjust their transmitter power such that power levels received by the OLT from all ONUs become the same. This method is not particularly favored by transceiver designers because it makes the ONU hardware more complicated, requires special signaling protocol for feedback from the OLT to each ONU, and, most important, may degrade the performance of all ONUs to that of the most distant unit.

Another issue is that it is not enough just to disallow ONUs to send any data (i.e., send only zeros). The problem is that even in the absence of data, lasers generate spontaneous emission noise. Spontaneous emission noise from several ONUs located close to the OLT can easily obscure the signal from a distant ONU. Thus, an ONU must shut down its laser between time slots. Also, it is important that the laser be able to stabilize quickly after being turned on.

SECURITY

Security has never been a strong part of Ethernet networks. In P2P full-duplex Ethernet security is not a critical issue because there are only two communicating stations using a private channel. In shared half-duplex Ethernet, security concerns are minimized because users belong to a single administrative domain and are subject to same set of policies.

P2MP Ethernet, however, has a different set of requirements. EPON has a broadcasting downstream channel and serves noncooperative users. In fact, an EPON cannot be considered a peer-to-peer network in that ONUs cannot communicate directly with each other (unless purposely provisioned to) or even learn of each other's existence. Since a malicious ONU may be placed in promiscuous mode and read all downstream packets, encryption mechanisms are necessary. Encryption and decryption may be implemented at the physical layer, data link layer, or higher layers. Implementing encryption above the MAC layer will encrypt the MAC frame payload only, and leave headers in plain text. Frame check sequence (FCS) is calculated for the encrypted payload. In this scenario, a MAC sublayer will verify the received frame integrity before passing the payload to a higher sublayer for decryption. This scheme prevents malicious ONUs from reading the payload, but they may still learn other ONUs' MAC addresses.

Alternatively, encryption can be implemented in the physical layer (abbreviated as PHY) below the MAC. In this scheme, the PHY layer will encode the entire bitstream, including the frame headers and CRC. At the receiving end, the PHY layer will decrypt the data before passing it to the MAC for verification. Since encryption keys are different for different ONUs, frames not destined to a given ONU will not decrypt into a properly formed frame and will be rejected by the MAC. In this scheme no information may be learned by a malicious ONU. But the difficulty is that PHY by definition is a connectionless layer. Requiring the PHY in an OLT to apply different keys for different ONUs will make it connection-aware. Encryption in EPON still remains an open question.

In this scheme no information may be learned by a malicious ONU. But the difficulty is that PHY by definition is a connection-less layer. Requiring the PHY in an OLT to apply different keys for different ONUs will make it connection-aware. Encryption in EPON still remains an open question.

SERVICE REQUIREMENTS

The migration of TDM circuit-switched networks to IP packet-switched networks is progressing at a rapid pace. However, although the next-generation access network will be optimized for IP data traffic, legacy equipment (RF set-top boxes, analog TV sets, TDM private branch exchangess, etc.) and legacy services (T1/E1s, ISDN, POTS, analog video, etc) will remain in use in the foreseeable future. Therefore, it is critical for next-generation access networks, such as EPONs, to be able to provide both IP-based services and jitter-sensitive and time-critical legacy services that have traditionally not been the focus of Ethernet.

Fortunately, EPONs can be designed to carry best-effort data as well as time-critical voice and video. This is typically accomplished using QoS techniques such as prioritization, VLAN tagging, reservations, and bandwidth provisioning. Several carriers, ranging from cable companies to large incumbent carriers, have deployed Gigabit EPONs delivering voice, data, and video.

The issue in implementing a circuit-over-packet emulation scheme is mostly related to clock distribution. In one scheme, users provide a clock to their respective ONUs, which is in turn delivered to the OLT. But since the ONUs cannot transmit all the time, the clock information must be delivered in packets. The OLT will regenerate the clock using this information. It is somewhat trivial to impose a constraint that the OLT should be a clock master for all downstream ONU devices. In this scenario, the ONUs will recover the clock from its receive channel, use it in its transmit channel, and distribute it to all legacy devices as a secondary clock reference.

IEEE P802.3AH STATUS

The standards work for Ethernet in the local subscriber access network is being done in the IEEE P802.3ah Ethernet in the First Mile (EFM) Task Force. This group received approval to operate as a Task Force from the IEEE-SA Standards Board in September 2001.

The P802.3ah EFM Task Force is bringing Ethernet to the local subscriber loop, focusing on both residential and business access networks. While at first glance this appears a simple task, in reality the requirements of local exchange carriers are vastly different than those of enterprise managers for which Ethernet was designed. In order to "evolve" Ethernet for local subscriber networks, P802.3ah is focused on four primary standards definitions:

- Ethernet over copper
- Ethernet over P2P fiber
- Ethernet over P2MP fiber
- Operation, administration, and maintenance (OAM)

Thus, the EFM Task Force is focused on both copper and fiber standards, optimized for the first mile and glued together by a common OAM system. This is a particularly strong vision, since it allows a local network operator a choice of Ethernet flavors using a common hardware and management platform. In each of these subject areas, new PHY specifications are being discussed to meet the requirements of service providers while preserving the integrity of Ethernet. Standards for EFM are anticipated by September 2003, with baseline proposals emerging as early as March 2002.

The Ethernet over P2MP track is focusing on the lower layers of an EPON network. This involves a PHY specification, with possibly minimal modifications to the 802.3 MAC. The standards work for P2MP fiber-based Ethernet is in progress, with a P2MP protocol framework emerging. This emerging protocol uses MAC control messaging (similar to the Ethernet PAUSE message) to coordinate multipoint-to-point upstream Ethernet frame traffic. Materials concerning the P802.3ah standards effort can be found at www.ieee802.org/3/efm and presentation materials at www.ieee802.org/ 3/efm/public.

THE MARKET FOR EPONS

Although many sectors of the telecommunications industry are suffering from the slowdown in service provider capital expenditures, analysts still expect the optical access market to grow rapidly. CIBC forecasts the market for PON access systems to reach $1 billion by 2004 from $23 million in 2000 [9]. Unlike the backbone network, which received an abundance of investment in long-haul fiber routes during the Internet boom, optical technology has not been widely deployed in the access network. It is possible that EPONs and P2P optical Ethernet offer the best possibility of a turnaround in the telecom sector. Service providers investing in optical access technologies will enable new applications, stimulating revenue growth and driving more traffic onto backbone routes. The large increase in access network bandwidth provided by EPONs and P2P optical Ethernet will eventually stimulate renewed investment in metro and long-haul fiber routes.

CONCLUSION

The subscriber access network is constrained by equipment and infrastructure not originally designed for high-bandwidth IP data. Whether riding on shorter copper drops or optical fiber, Ethernet is emerging as the future broadband protocol of choice, offering plug and play simplicity, IP efficiency, and low cost. Of particular interest are Ethernet PONs, which combine low-cost point-to-multipoint optical infrastructure with low-cost high-bandwidth Ethernet. The future broadband access network is likely to be a combination of point-to-point and point-to-multipoint Ethernet, optimized for transporting IP data, as well as time critical voice and video.

REFERENCES

[1] "Access Network Systems: North America — Optical Access. DLC and PON Technology and Market Report," RHK-RPT-0548, RHK Telecommun. Industry Analysis, San Francisco, CA, June 2001.

[2] K. G. Coffman and A. M. Odlyzko, "Internet Growth: Is There a "Moore's Law" for Data Traffic?" *Handbook of Massive Data Sets*, J. Abello, P. M. Pardalos, and M. G. C. Resende, Eds., Kluwer, 2001.

[3] JP Morgan Securities, Inc., Broadband 2001, A Comprehensive Analysis of Demand, Supply, Economics, and Industry Dynamics in the U.S. Broadband Market," Apr. 2001.

[4] G. Pesavento and M. Kelsey, "PONs for the Broadband Local Loop," *Lightwave*, vol. 16, no. 10, Sept. 1999, pp. 68–74.

[5] B. Lung, "PON Architecture 'Futureproofs' FTTH," *Lightwave*, vol. 16, no. 10, Sept. 1999, pp. 104-7.

[6] S. Clavenna, "Metro Optical Ethernet," *Lightreading* (www.lightreading.com), Nov. 2000.

[7] K. Claffy, G. Miller, and K. Thompson, "The Nature of the Beast: Recent Traffic Measurements from an Internet Backbone," *Proc. INET '98*, Geneva, Switzerland, July 1998; http://www.isoc.org/inet98/proceedings/6g/6g_3.htm

[8] J. L. Hammond and P. J. P. O'Reilly, *Performance Analysis of Local Computer Networks*, Addison Wesley, 1987.

[9] CIBC World Markets, Inc., "Passive Optical Networks — Is There Light at the End of Access Tunnel?" Jan. 2001.

BIOGRAPHIES

GLEN KRAMER (glen.kramer@alloptic.com) is an R&D engineer in Advanced Technology Lab in Alloptic, Inc. He received a B.S. degree in computer engineering from Kishinev Polytechnic Institute, Moldova, in 1992, and an M.S. degree from University of California at Davis in 2000. He is a member of Networks Research Lab at UC Davis. His research interests are in the areas of passive optical networks, broadband access network architectures and protocols, QoS in packet networks, traffic characterization, and modeling.

GERRY PESAVENTO (gerry.pesavento@alloptic.com) is the founder and vice president of R&D at Alloptic, Inc. He is also the current chair of the EPON track within the IEEE P802.3ah EFM standards effort. His research interests include fiber optic components and optical Ethernet networks. He received a B.S.E.E. from Rensselaer Polytechnic Institute (RPI) and an M.B.A. from the University of California at Berkeley.

The future broadband access network is likely to be a combination of point-to-point and point-to-multipoint Ethernet, optimized for transporting IP data, as well as time critical voice and video.