

Fiber to the Getting Beyond

CEDRIC F. LAM

Bandwidth-hungry applications are driving adoption of fiber-based “last-mile” connections—and raising the challenge of moving access-network capacity to the next level.

Home: 10 Gb/s



F

or years, the global Internet’s “last mile”—the access networks that connected end users to the central offices (COs) of telecom operators—represented the key bottleneck in end-to-end network infrastructures. That has begun to change with the widespread adoption of broadband access technologies, including digital subscriber line (DSL), cable modem and, most recently, fiber-to-the-home (FTTH). These technologies have helped enable the Internet we know and use today, and have led to the proliferation of applications such as cloud computing, the “Internet of Things” and video streaming.

FTTH in particular, and initiatives such as Google Fiber, are spurring the development of true “IPTV”—video streaming over the Internet Protocol (IP) network that carries other network traffic. And these developments are nudging video streaming further and further toward bandwidth-hungry, time-shifted “unicast” models of consumption.

The drive toward increasing video streaming will inevitably lead to greater adoption of FTTH—and greater use of the passive optical networks (PONs) that have emerged as the dominant technology for implementing FTTH. That very demand, however, has presented current PON technologies with a scaling problem, much like the one faced by the long-haul networks of 20 years ago. While PON capacity has expanded substantially in recent years to the 10-Gb/s

systems that represent today’s state-of-the-art, those systems are now mature, and the community must now develop new PON systems to address the needs of upcoming years.

This article takes a look at the current state and future prospects of FTTH and, in particular, PON systems. It also explores some potential design challenges with the current favored architecture for next-generation PONs, and suggests some plausible alternative paths.

FTTH today

In industry jargon, FTTH is a subset of “fiber-to-the-x” (FTTx), which denotes any broadband architecture in which optical fiber is used for last-mile connections—that is, the “x” might stand for curb, premises, building, home, desktop or something else. A smooth broadband user experience requires a ubiquitous, affordable end-to-end broadband infrastructure, and FTTx is generally regarded as the ultimate last-mile component of that infrastructure, because of optical fiber’s virtually unlimited bandwidth for an access network.

In 2014, worldwide FTTx deployment passed 100 million connections, and rapid growth continues today. In the United States, FTTH implementation has lagged relative to Asia; nonetheless, more than 100 cities are now connected with gigabit-capable broadband networks through Google Fiber and competing initiatives, mostly with FTTH technologies.

Architecturally, the most straightforward approach to achieving FTTH is a “home run” setup, which connects each end user to the CO with a single dedicated fiber. This brute-force approach, while simple, requires a large number of active transceivers and fiber strands, and thus

The drive toward increasing video streaming will inevitably lead to greater adoption of FTTH—and greater use of passive optical networks.

THE EVOLUTION OF TV

Video streaming is rapidly evolving from traditional, bandwidth-efficient broadcast models to user-driven “unicast” consumption, demanding ever-greater bandwidth.



1927
Analog TV



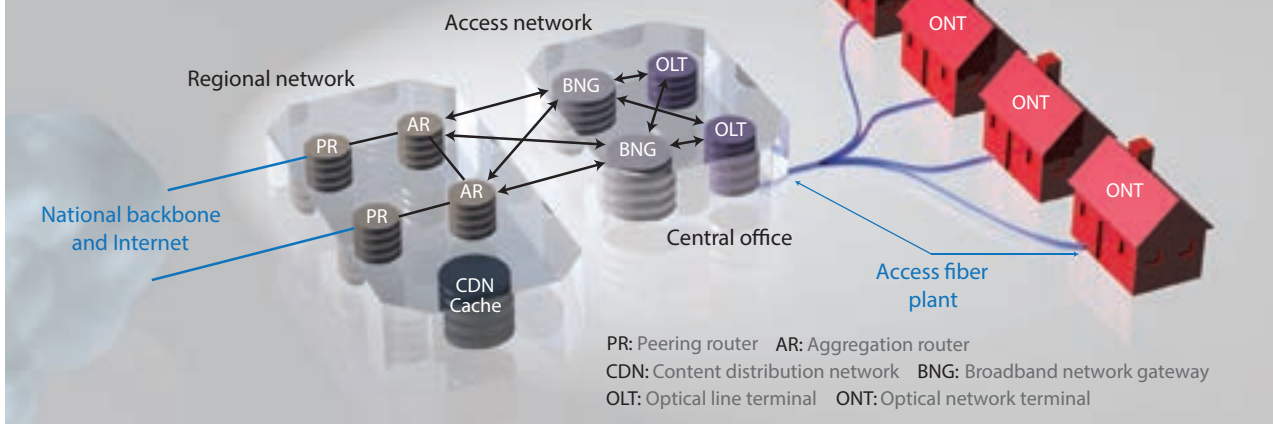
1990
Digital TV



2010
HDTV / 3-D TV



END-TO-END BROADBAND ACCESS NETWORK ARCHITECTURE



Phil Saunders

substantially increases infrastructure costs. As a result, the vast majority of FTTH implementations have adopted a PON architecture, in which an unpowered optical splitter enables a single fiber from the CO to serve multiple endpoints.

Most PON systems deployed to date have been time-division-multiplexed (TDM) systems. In this architecture, the signal from a single line emanating from the CO's optical line terminal (OLT) is broadcast through a power splitter at a remote node in the field, and the bandwidth is shared among 32 to 64 users (or, in some cases, 128) using a TDM protocol.

TDM-PON systems allow operators to take advantage of statistical multiplexing gains and make efficient use of last-mile bandwidth, and have enjoyed significant commercial success: the current G-PON (2.5-Gb/s downstream and 1.25-Gb/s upstream capacity) and E-PON (1-Gb/s symmetric capacity) standards both belong to the TDM-PON architecture, and higher-capacity installations using the IEEE's newer 10-Gb/s standard, 10G-EPON, are also being deployed. (See sidebar on p. 28 for more information on the competing, and sometimes confusing, array of PON

standards.) Yet scaling TDM-PON systems for the expected future explosion in bandwidth demand presents a number of potential challenges:

No unicast advantage. TDM-PON's broadcast architecture, while quite efficient for delivering broadcast TV signals to end users using a native PON multicast mode, offers no advantage for the unicast traffic that will become increasingly dominant.


Potentially costly upgrade path. One characteristic of TDM-PON systems is that each end user's optical network terminal (ONT) needs to operate at the aggregate speed shared among all the PON members. Thus, to upgrade a given PON system to 10-Gb/s capacity, every user needs to be equipped with a 10-Gb/s-capable transceiver—a potentially expensive proposition.

In addition, making a new, higher-speed OLT at the CO backward-compatible with legacy ONTs among end users is not a simple task. As a result, 10-Gb/s TDM-PONs overlay



The Internet's "Last Mile"


FOUR FTTH ARCHITECTURES



[Current Generation]

THE HOME RUN


A single fiber connects each end user to the telecom central office (CO). It's simple, but expensive, with high infrastructure costs that make it difficult to scale.



[Next Generation]

WDM-PON ARCHITECTURE


Scalability could come from using a wavelength splitter to assign each end user a dedicated wavelength—a WDM-PON. Could potentially boost network capacity without upgrading all end-user network terminals to overall network speed, a possible advantage in scaling.



[Current Generation]

TDM-PON ARCHITECTURE

Data traveling along a single fiber is optically split, separated in time and then shared among 32 to 64 different end users. Commercially successful, but equipment requirements, especially for end users, will make it difficult to scale past 10 Gb/s.



[Next Generation]

NG-PON2 (A TWDM-PON ARCHITECTURE)

Combines WDM-PON and TDM-PON to allow higher bandwidth and backward compatibility with current-gen installations. Four parallel transceivers at the CO terminal (OLT), each operating at 10 Gb/s, allow for total network capacity of 40 Gb/s; a tunable filter at the end-user terminal (ONT) selects the user's downstream wavelength. May be too complex to be economically feasible at present.

Phil Saunders

slightly different transmission wavelengths—1,270 nm for upstream and 1,577 nm for downstream transmission, versus 1,310 upstream and 1,490 nm downstream for G-PON and E-PON—to allow the 10-Gb/s PON ONUs and OLTs to coexist on the same fiber plant with the slower G-PON and E-PON architectures. Scaling the speed of TDM-PONs by 10 times also will require 10 times the power budget, and 10-Gb/s TDM-PONs have required forward error correction to overcome that challenge. Chromatic dispersion effects also become significant at speeds above 10 Gb/s.

Uplink bottleneck. Another challenge for TDM-PON is in the uplink direction, as multiple end-user ONTs all transmit to a single OLT at the central office. Thus TDM-PONs must support high-speed, burst-mode operation, with the end-user ONT only transmitting during an assigned time slot and the OLT receiver at the central

office quickly synchronizing its clock and nimbly adjusting to accept the upstream packet burst.

Such high-speed burst-mode OLT receivers are challenging to implement. As a result—especially since video streaming, which requires relatively little upstream bandwidth, has been the dominant broadband application to date—a number of asymmetric PON standards have been proposed to ease upstream implementation and reduce system costs for next-gen passive optical networks.

Wavelength-division-multiplexed (WDM) PON

Given the difficulty of scaling TDM-PONs, the community is actively exploring alternative approaches. One such alternative is a WDM-PON, which connects each end user via a dedicated wavelength, assigned by a wavelength splitter in the remote node. From a transmission perspective, WDM-PONs are much more scalable than TDM-PONs, because

each end-user ONT can operate at its own speed (rather than, as with TDM-PON, requiring each end-user ONT to operate at the speed of the entire network).

Hence, for example, a 40-channel WDM-PON, with each channel having a speed of 1 Gb/s, will have an aggregate PON capacity of 40 Gb/s, whereas in a TDM-PON, each end-user ONT will need to operate at 40 Gb/s for the same aggregate PON capacity, which could be significantly harder to achieve and scale in the long run. Moreover, a WDM-PON does not require burst-mode operation, which makes transceiver electronics much simpler. And the small insertion loss of a WDM splitter significantly reduces requirements in laser output power, receiver sensitivity and dynamic range.

Yet WDM-PONs have challenges of their own—one of the main issues being the cost of components. To make efficient use of the optical spectrum inside fibers, WDM-PONs require dense WDM (DWDM), and DWDM lasers need thermoelectric cooler (TEC) controllers to maintain wavelength stability. DWDM lasers for conventional transport systems are packaged in a comparatively expensive flatbed butterfly package; PON lasers are packaged in a lower-cost TO-CAN, similar to transistors. A scheme that packaged both the laser and TEC controller in a single TO-CAN would help to make WDM-PON cost-effective. In addition, the availability of low-cost tunable lasers will be important in realizing practical WDM-PON systems, as each end-user ONT needs to be equipped with a tunable laser.

TWDM-PON architectures

Perhaps the biggest disadvantage of WDM-PON, however, is that the required fiber plant is not compatible with traditional TDM-PONs already deployed in the field. As a result, the 40-Gb/s next-generation standard proposed by ITU-T, called NG-PON2, combines TDM and WDM into a single TWDM architecture, to keep the system compatible with existing TDM-PON fiber installations.

NG-PON2 represents the first time that WDM has been introduced into a PON standard to increase PON capacity. Because transmission of serial 40-Gb/s signals in a pure TDM manner is very difficult, the NG-PON2 standard uses four parallel 10-Gb/s WDM transceivers at the CO's OLT to achieve an aggregate PON capacity of 40 Gb/s. The four wavelengths are multiplexed and demultiplexed with a 4×10 -Gb/s WDM array transceiver at the OLT, and all wavelengths are broadcast to all end-user ONTs, each of which is equipped with a single-channel, 10-Gb/s transceiver. To make the ONTs wavelength agnostic, a tunable filter is used at the ONT to select one of the OLT downstream wavelengths to be received. As with pure WDM-PONs,

each ONT needs to be equipped with a tunable laser. Although NG-PON2 has advertised itself as a 40-Gb/s PON, the end-user ONTs can only burst up to 10 Gb/s—which is, however, substantially more than adequate for current broadband applications.

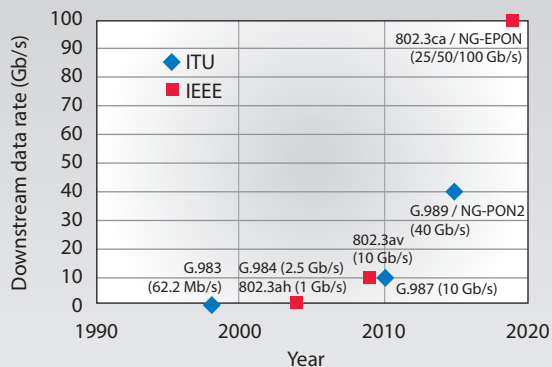
The NG-PON2 design philosophy has several problems, however. For one, the new standard aims at supporting multistage splitting and high splitting ratios (perhaps as great as 1,024 users per CO connection), to efficiently use the vast bandwidth offered by NG-PON2 and to cut system costs through better sharing of the expensive OLT optics and through reducing the number of fibers required to connect to the CO. But not many of the embedded PON systems currently deployed have such large splitting ratios. Moreover, to support the very high splitting ratio, the OLT burst-mode receiver requires a high dynamic range (greater than 20 dB).

A second problem is the potentially very high power budget needed for NG-PON2 implementation. The NG-PON2 standard specifies a loss budget of up to 35 dB between the OLT and ONT, and this does not account for the extra losses from the WDM multiplexers inside the OLT and the tunable filter in the ONT optical module. As a result, lasers with very high transmitting power and receivers with high sensitivity will be needed, at a potentially significant increase to the system costs—even if the technology is feasible. And having to support very high splitting ratios will strain the power budget requirement still further.

Can TWDM-PON work economically?

In the author's opinion, optical-component technologies have not advanced to the point that NG-PON2 implementation will be economically feasible any time soon—especially given that there is no real demand for such systems at present. The step from G-PON to 10-Gb/s TDM-PON was incremental, as the optical systems are quite similar apart from the speeds of the optical transceivers. However, in moving from current 10-Gb/s TDM-PON to NG-PON2, the structures of the optical transceivers are much more

**NG-PON2
represents the first
time that WDM has
been introduced
into a PON standard
to increase PON
capacity.**



Understanding the PON standards

Standards in passive optical networks can be the source of confusion, as they involve two standard-setting bodies—the Institute of Electrical and Electronics Engineers (IEEE) and the Telecommunication Standardization Sector of the International Telecommunication Union (ITU-T). Here’s a look at the main standards.

G-PON and E-PON. Most commercial PON systems deployed today are based on the ITU-T G.984 “G-PON” asymmetric standard (2.5 Gb/s downstream, 1.25 Gb/s upstream) or the IEEE 802.3ah “E-PON” symmetric standard (1 Gb/s downstream and upstream), both of which are based on time-division-multiplexed (TDM) PON architectures.

10 Gb/s PON systems. In 2009, the IEEE introduced a next-generation standard for 10-Gb/s PONs, standardized as IEEE 802.3av and known as 10G-EPON. This TDM-PON has proved the most successful (and, now, mature) of the 10-Gb/s PON systems, employed especially in Asia for fiber-to-the-building installations to connect multiple-dwelling-unit buildings.

ITU-T answered by starting work on two different standards, the asymmetric XG-PON1 (with 10-Gb/s downstream and 2.5-Gb/s upstream capacity) and the symmetric XG-PON2 (with 10-Gb/s downstream and upstream speeds), which collectively came to be called NG-PON (Next Generation PON). While prototypes of these systems were built, the difficult-to-meet requirement for burst-mode timing, coupled with a lack of real demand, delayed commercialization and completion of these standards.

Beyond 10 Gb/s. More recently, to stay ahead of IEEE in competition, ITU-T has proposed a new standard, NG-PON2 (or ITU-T G.989), which would leapfrog the earlier unsuccessful NG-PON standard to aim at capacities of 40 Gb/s, by combining time-division and wavelength-division multiplexing (TWDM) in the same architecture. (Because of the technical difficulties of putting that standard into operation, ITU-T is working on an “initial stage” or bridge standard, XGS-PON, with 10-Gb/s symmetric capacity, that can support fixed-wavelength operation of existing technologies but can adapt to multi-wavelength implementations when the optical technologies are ready.)

IEEE set up a study group for its own competing next-gen standard, Next-Generation E-PON (NG-EPON), in November 2015. Also a TWDM approach, NG-EPON aims at supporting ONTs with one, two or four wavelengths and at system capacities of 25 Gb/s, 50 Gb/s or 100 Gb/s.

complex, and packaging complexity becomes exponentially higher.

Small, low-loss, low-cost and low-power tunable filters, required for end-user ONTs in an NG-PON2 system, are not easier to manufacture than semiconductor tunable lasers (which are mostly monolithic structures). In the long run, innovations in photonic integration circuits could potentially solve these challenges.

NG-PON2 also creates problems from the perspective of the medium access control (MAC) protocol, which manages PON bandwidth allocation to individual ONTs. That’s because NG-PON2 adds a layer of optical wavelengths to manage, while still retaining the complexities of dynamic time slot management and burst-mode transmission that are challenges in TDM-PON. The millisecond-scale tuning speeds of the tunable laser and filters used in NG-PON2 will make it economically difficult to do fast wavelength switching on the packet level. Coordinating wavelengths with time slots together will add complexity to NG-PON2 dynamic bandwidth allocation (DBA) algorithms.

On the physical layer, burst-mode operation of a WDM laser causes transient wavelength drift, as the sudden injection of current into the laser heats up the laser structure (unless an external modulator is used, which increases cost and optical losses). The drift increases with the laser bias current and laser output power, and can be as large as 20 to 30 GHz, with time constants on the order of milliseconds. Such drift not only causes a penalty in OLT receiver sensitivity at the central office, but also crosstalk to other wavelength channels in a broadcast-and-select TDM-PON fiber plant. A pure WDM-PON without TDM overlay, on the other hand, does not require burst-mode operation and, thus, does not suffer from these challenges. These attributes could make WDM-PON much easier to implement and scale in the long run—despite the industry’s current preference for TWDM-PON architectures.

The outlook for next-generation PON systems

So—where does all of this leave the outlook for next-gen PONs? Access networks, like datacom systems, are very cost sensitive—and, in the datacom world, a tenfold performance improvement is usually required to justify a twofold cost increase

for deploying a next-generation system. NG-PON2 may not meet that test: The NG-PON2 optical layer not only is complex on paper, but also has proven very difficult to implement through R&D activities in the past three to four years.

In the foreseeable future, 10-Gb/s PON networks should provide adequate bandwidth to meet existing demands, so there is no immediate justification for a need to dynamically adjust wavelength loadings, with the additional cost and complexity that would bring. That fact may underlie ITU-T's decision in 2015 to create the XGS-PON (10-Gb/s symmetric PON with single wavelength) standard as an "initial stage" of NG-PON2. (See sidebar at left exploring these and other PON standards.)

As PON speeds increase beyond 10 Gb/s, however, some form of WDM becomes necessary to handle the scaling. The view beyond 10 Gb/s has prompted the other standard-setting body, IEEE, to envision next-generation E-PON (NG-E-PON), with a study group begun in November 2015 and with the initial use case targeting business applications and cellular backhubs. NG-E-PON is, like ITU-T's NG-PON2, a TWDM standard. It starts with a single-wavelength speed of 25 Gb/s, implemented with duo-binary modulation to reduce the transmission bandwidth requirement and dispersion penalties. Other modulation methods such as PAM-4 have also been proposed. The idea is to control the signal bandwidth, so that 10-Gb/s optics could be used for transmission, thereby potentially smoothing the upgrade path from the current technology standard.

The role of new applications

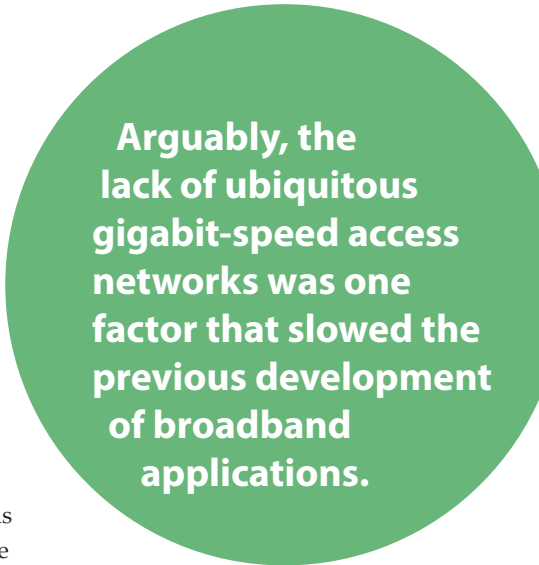
It is interesting to note that, although standards for 10-Gb/s PONs were completed in 2009, and although these systems, have been deployed for fiber-to-the-building in some Asian countries, large-scale rollout of 10-Gb/s-PON FTTH has never really happened. That's because, quite simply, bandwidth demand has not materialized—indeed, according to IEEE, the capacities provided by the current G-PON or E-PON technologies should suffice till 2020. Video streaming remains by far the biggest bandwidth driver for broadband networks, and a 1080p HD video stream consumes only about 10 to 15 Mb/s bandwidth. Thus, the 2.5-Gb/s downstream bandwidth of G-PON, shared among 32 users in standard deployment, is more than sufficient to support current video streaming.

Yet these arguments don't undermine the case for action now on the next-generation PON standards. Arguably, the lack of ubiquitous gigabit-speed access networks was one factor that slowed the previous development of

broadband applications that could take advantage of the vast bandwidths available from FTTH. With those gigabit-capable networks now proliferating, we can expect new "killer" applications to start to emerge—some of which may use hundreds of megabits per second bandwidths by themselves, and in combination may drive large increases in total bandwidth demand per household.

The immersive experiences of virtual reality, for example, will demand bandwidths of 50 to 100 Mb/s per application. Video surveillance for home security, high-definition video conferencing, telemedicine and other Internet of Things applications will generate significant upstream bandwidth demands, on the order of a few hundred megabits per second. These applications, and others we can't yet foresee, will drive ever-greater needs for capacity and speed—and for a new generation of optical access systems. [OPEN](#)

Cedric Lam (cflam@ieee.org) is Engineering Director at Google Access.



References and Resources

- ▶ W. Harrop. "Quantifying the broadband access bandwidth demands of typical home users," Australian Telecommunication Networks & Applications Conference (December 2006).
- ▶ C.F. Lam. *Passive Optical Networks: Principles and Practice*, Academic Press (2007).
- ▶ C.F. Lam. "FTTH look ahead—technologies and architectures," European Conference and Exhibition on Optical Communication, Turin, Italy (2010).
- ▶ P. Z. Dashti et al. "Optoelectronic integration for broadband optical access networks," IEEE Photonics Conference Burlingame, California (2012).
- ▶ R. Urata et al. "Silicon photonics for optical access networks," 9th International Conference on Group IV Photonics (GFP), San Diego, Calif. (2012).
- ▶ X.-Z. Qiu. "Burst-mode receiver technology for short synchronization," Optical Fiber Communications Conference Tutorial, OW3G.4 (2013).
- ▶ R. Bonk et al. "The underestimated challenges of burst-mode WDM transmission in TWDM-PON," *Opt. Fib. Technol.* **26**(A), 59 (December 2015).
- ▶ IEEE 802.3 Industry Connections NG-E-PON Ad Hoc report, March 2015.
- ▶ IEEE 802.3 NG-E-PON study group meeting, 10-12 November 2015, www.ieee802.org/3/NGEPONSG/public/2015_11/index.html
- ▶ International Telecommunications Union (ITU), ITU-T work program G.XGS-PON, www.itu.int/itu-t/workprog/wp_item.aspx?isn=10591