Evolution of Ethernet Standards in the IEEE 802.3 Working Group

David Law, Hewlett-Packard Ltd.
Dan Dove, Applied Micro
John D'Ambrosia, Dell
Marek Hajduczenia, ZTE Corporation and Universidade de Coimbra
Mark Laubach, Broadcom Corporation
Steve Carlson, High Speed Design, Inc.

ABSTRACT

Ethernet is constantly evolving, adapting to the needs of the networking world, addressing the requirements of both operators and end users, while making sure that the resulting technology is cost-efficient, reliable, and operates in a plug-and-play manner. The IEEE 802.3 Working Group has been working for the last 30+ years, pushing the boundaries on the speed and capacity of wireline Ethernet links, migrating from shared medium CSMA/CD systems to switched point-to-point Ethernet and then introducing multilane technology and point-to-point emulation over shared media of passive optical networks. In this article, we look at the latest projects adding new features and capabilities to the family of wired Ethernet standards, enabling the exponential growth of the Ethernet ecosystem, driven by technical maturity, cost effectiveness, and broad market support.

INTRODUCTION

The total amount of data created or replicated on the planet in 2010 exceeded 1 zettabyte (1 zettabyte is 10^{21} bytes), or 143 Gbytes for each of the 7 billion people on the planet [1]. This volume of information requires high-speed links between server farms, cloud storage, and end users to make sure that it can be processed in a timely and reliable fashion. The relentless growth of the number of end stations connected to the network, whether permanent or nomadic (computer terminals, mobile devices, automated devices generating machine-to-machine traffic), has led to explosive growth in the volume of information exchanged at all levels of the networking infrastructure. The popularity of Ethernet and its widespread use in access, aggregation, transport, core networks, and data centers, combined with the unprecedented demand for advanced data connectivity services, fuel the development of new Ethernet standards, providing higher-speed links to address the market demand.

Ethernet is also venturing into brand new application areas, and is adding support for synchronization protocols or even potentially becoming a de facto standard for in-vehicle data networks, providing a common transport platform for control and multimedia applications.

This article will examine the evolution of Ethernet standards taking place in the IEEE 802.3 Working Group. There are a number of exciting new projects, pushing the boundaries of Ethernet into new application areas and markets.

EVOLUTION OF ETHERNET STANDARDS

The IEEE Std 802.3 Ethernet standard was first published in 1985, specifying a half-duplex carrier sense multiple access with collision detection (CSMA/CD) medium access control (MAC) protocol operating at 10 Mb/s, and a medium attachment unit (MAU) for operation on a coaxial cable medium, supporting a bus topology between the attached end stations.

Amendments to the IEEE 802.3 standard then added specifications for, among other items, a repeater to extend topologies supported, MAUs for operation over fiber optic cabling, a MAU for operation over twisted pair cabling, 10BASE-T, and layer management. In 1995 amendment IEEE Std 802.3u was published adding operation at 100 Mb/s (fast Ethernet). This included a number of physical layer (PHY) specifications for operation over fiber optic and twisted pair cabling (100BASE-TX).

Amendment IEEE Std 802.3x published in 1997 added full duplex operation to the MAC and a flow control protocol to take advantage of the full duplex capable medium, such as twisted pair and fiber, for which PHYs were already specified in IEEE 802.3, as well as support switching, which was becoming more cost effective due to increased device integration.

In 1998 amendment IEEE Std 802.3z was

 0163-6804/13/$25.00 © 2013 IEEE

published, adding operation at 1000 Mb/s (Gigabit Ethernet), and subsequently in 1999 amendment IEEE Std 802.3ab was published, adding 1000BASE-T PHY specifications to support 1000 Mb/s operation over twisted pair cabling.

Amendment IEEE 802.3ad (link aggregation) was published in 2000, adding the ability to aggregate multiple full duplex point-to-point links in to a single logical link from the perspective of the MAC client. Since link aggregation has application beyond Ethernet, as well as its architectural positioning, it was subsequently moved the the IEEE 802.1 Working Group in 2008 and is now titled IEEE Std 802.1AX Link Aggregation.

In 2002 amendment IEEE Std 802.3ae was published adding operation at 10 Gb/s (10 Gigabit Ethernet), and in 2006 amendment IEEE Std 802.3an was published adding the 10GBASE-T PHY specifications to support 10 Gb/s operation over twisted pair cabling. It was followed in 2010 by the amendment IEEE Std 802.3ba, adding operation at 40 Gb/s and 100 Gb/s (40 Gigabit Ethernet and 100 Gigabit Ethernet). The development of 40 Gb/s and 100 Gb/s Ethernet was done in close cooperation with International Telecommunications Union Telecommunications Standardization Sector (ITU-T) SG15 to ensure transparent connectivity into the optical transport network (OTN). Operation at 10 Gb/s, 40 Gb/s, and 100 Gb/s only supports full duplex operation.

Amendment IEEE Std 802.3ah published in 2004 first added support for subscriber access network Ethernet (Ethernet in the first mile, or EFM for short). As well as the addition of a number of fiber optic and voice grade copper PHYs, it also specified a fiber optic point-to-multipoint network topology using passive optical splitters known as Ethernet passive optical networks (EPONs).

Amendment IEEE Std 802.3ap first added support for backplane Ethernet in 2007.

A summary of the speed and distance for various MAUs and PHYs supported by the approved IEEE 802.3 standard (at the time of writing this article) and amendments is shown in Fig. 1.

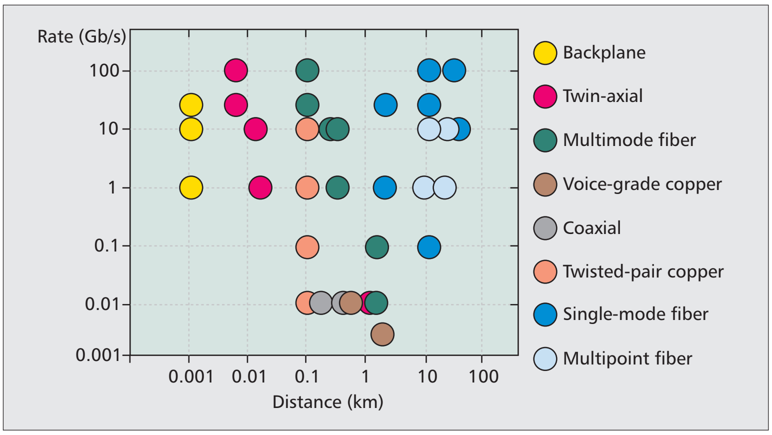

Figure 1. *Speed and reach for various IEEE Std 802.3 MAUs and PHYs.*

Other additions include IEEE Std 802.3af, "DTE Power via MDI," published in 2003, also known as power over Ethernet, which enables power to be supplied on the same cabling as the data transmission, and IEEE Std 802.3at, published in 2009 which enhanced the maximum power available and the classification mechanism.

In addition, in 2010 amendment IEEE Std 802.3az added support for energy-efficient Ethernet (EEE) to, among others, the 100BASE-T, 1000BASE-T, and 10GBASE-T PHYs. This not only reduces the power consumption of the PHYs, but also specifies signaling that can enable the reduction of the power consumption of the attached device.

High-Speed Optical P2P Links

In August 2011, the IEEE 802.3 Working Group authorized the formation of a Study Group to address "Next Generation 100G Optical Ethernet" with the goal of reducing the cost and power of, and improving density for 100G optical solutions. The IEEE P802.3ba project produced a number of 100G optical PHYs based on the 10x10G CAUI (100 Gigabit attachment unit interface) electrical interface standard, requiring a relatively large number of interface data paths, and for the longer-reach power management devices (PMDs), a gearbox to reduce the number of optical lanes from 10 to 4. This gearbox would normally be included in the optical module, creating a power constraint that inhibits higher port density; see Fig. 2 for differences in device architecture. Through the work being done in IEEE 802.3, this gearbox can be moved out of the module and eventually integrated directly into the host integrated circuit (IC). This allows smaller lower-power modules, driving density up and cost per port down.

While the 40 Gigabit Ethernet and 100 Gigabit Ethernet standards enabled Ethernet networking at 40G and 100G for the very first time, this next generation of 40G and 100G optical standards are expected to provide a substantial decrease in the cost and complexity required for broad deployment of 100G Ethernet at just the right time as network providers and customers see the demand curve rising. In addition, the 40 km specification for 40G will enable near-term deployment of long-reach optical channels.

The Study Group spent close to a year to define a set of technical objectives, including among others:

- Provide appropriate support for optical transport network (OTN)
- Define retimed 4-lane 100G electrical interfaces for chip-to-chip and chip-to-module applications
- Define a 40 Gb/s PHY for operation over at least 40 km of single-mode fiber (SMF)
- Define a 100 Gb/s PHY for operation up to at least 500 m of SMF, at least 100 m of multimode fiber (MMF), and at least 20 m of MMF in separate devices

The addition of the 40 Gb/s PHY objective was made subsequent to the March 2012 meeting, addressing the demonstrated demand for a longer-reach solution at this particular speed. Instead of creating a new project, it was decided that this objective would be added to the scope of the Next Generation 100G project.

The chip-to-chip interface objective calls for the definition of a CAUI-4 electrical interface, which is anticipated to leverage the OIF CEI-28G-VSR specification [12]. This interface will

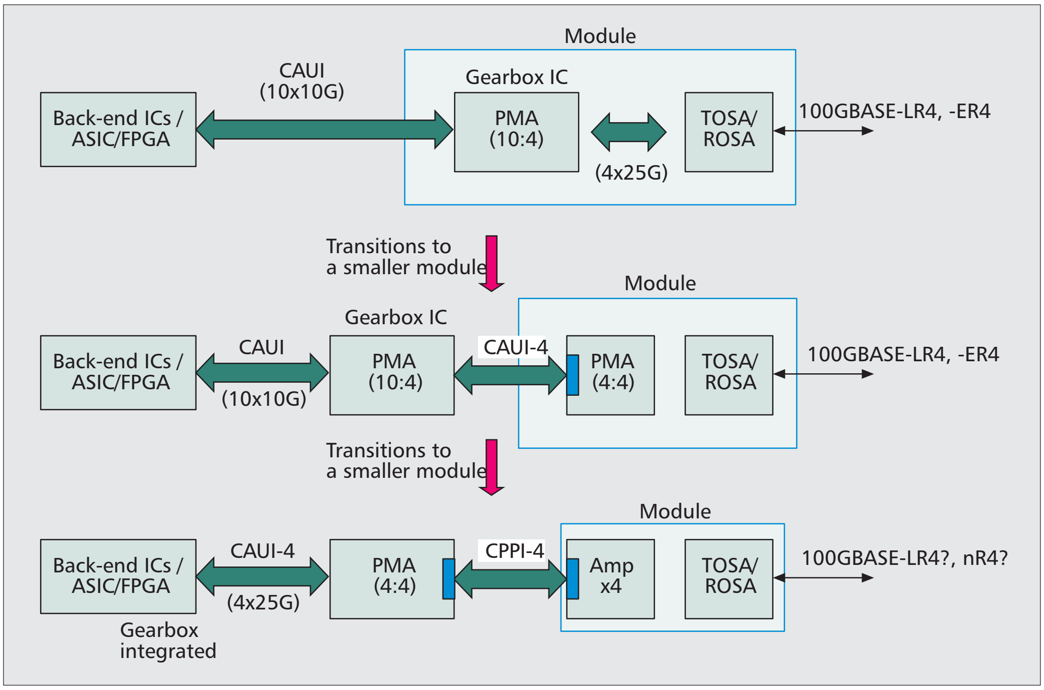

Figure 2. *Evolution of 100G module architectures (adapted from CFI materials for Next Generation 100 Gb/s Ethernet, http://www.ieee802.org/3/100GNGOPTX/public/jul11/CFI_01_0711.pdf).*

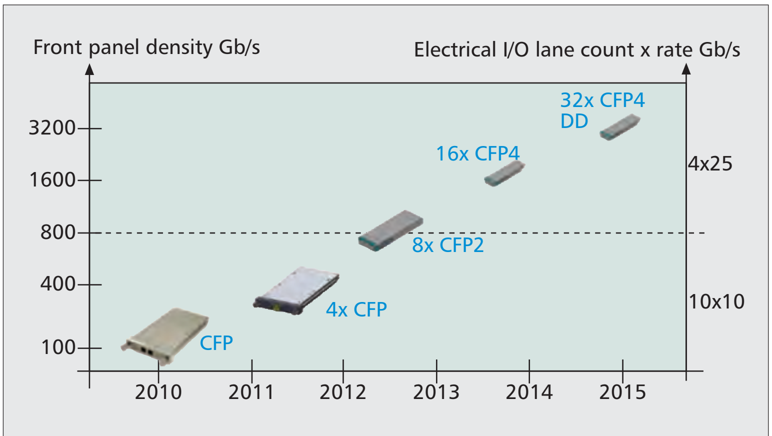

Figure 3. *Front panel density vs. I/O count × rate and time.*

foster higher-density implementations of the existing 100GBASE-LR4 and 100GBASE-ER4 PHYs simply by providing a standard 4 × 25 Gb/s electrical interface to enable development of optical modules using this particular interface. It will also enable an MMF solution that uses only 4 data lanes (4 fibers) in each direction vs. the current solutions requiring the use of 10 fibers in each direction, as required by the 100GBASE-SR10 standard. In September 2012, a baseline proposal for the chip-to-module interface was adopted. At the time of writing, work is ongoing for a chip-to-chip interface baseline proposal.

For MMF applications, two reach objectives were defined with the understanding that a significant cost difference must exist between them, or a single MMF solution that covers at least 100 m of MMF will address both objectives. Such a PHY is intended to address the demand for short-reach data center links running between server and top of rack (TOR), server to end of rack (EOR), or TOR to EOR switches. It is expected that continued work to identify the best way to achieve MMF operation over the defined objectives will consider the bandwidth of VCSEL sources, equalization methods, forward error correction (FEC), and the balance between these methods to achieve the best cost/reach solution.

For SMF applications, a reach objective of at least 500 m was defined to address the demand for a lower-cost alternative to 100GBASE-LR4, which was originally specified to support the reach of 10 km. A number of alternative technologies have been studied so far, including coarse wavelength-division multiplexing (WDM), LAN WDM, parallel SMF links, pulse-amplitude-modulated optical signaling using SiPhotonics, and mid-wave vertical cavity surface-emitting lasers (VCSELs). The project anticipates that at a minimum, 100GBASE-LR4 cost, power, and size requirements will be improved via removal of the gearbox, thus meeting the stated SMF objective. If a lower-power and cost-optimized solution appears that is capable of exceeding the 500 m reach objective, it must be distinct enough from LR4 to justify the development of a new PMD.

It is anticipated that a preliminary draft of the IEEE P802.3bm specification will be made available for Task Force review in May 2013 with the anticipated Working Group ballot starting by November 2013 and a final specification ready for ratification in March 2015.

For more information about the Task Force, please see http://www.ieee802.org/3/bm/index.html.

High-Speed Copper P2P Links and Backplane Technologies

In 2007, the family of PHYs for Ethernet operation over electrical backplanes for Gigabit Ethernet and 10 Gigabit Ethernet was first introduced. Two PHYs were introduced for 10 Gigabit Ethernet: 10GBASE-KX4 and 10GBASE-KR. The 10GBASE-KX4 PHY is a full duplex solution employing four data lanes in each direction, where each lane operates at 3.125 Gb/s and employs 8B/10B line encoding to support the effective data rate of 10 Gb/s. The 10GBASE-KR PHY is a serial lane solution operating at 10.3125 Gb/s and employing 64B/66B encoding to support the effective data rate of 10 Gb/s.

These two PHYs laid the groundwork for the 40 Gigabit Ethernet backplane PHY, which was developed during the IEEE P802.3ba project. Utilizing the 4-lane approach of 10GBASE-KX4 and the serial 10 Gb/s electrical signaling developed for 10GBASE-KR, the 40GBASE-KR4 PHY supports 40 Gigabit Ethernet operation across an electrical backplane. At the time of the IEEE P802.3ba project there was no 25 Gb/s per lane electrical signaling solution available; therefore, no 100 Gigabit Ethernet backplane PHY was developed.

The Call for Interest to develop the operation of Ethernet at 100 Gb/s across an electrical backplane, as well as across twin-axial cables, took place in November 2010. Fueled by the small form-factor pluggable plus (SFP+) supporting 10 Gigabit Ethernet, quad SFP (QSFP) supporting 40 Gigabit Ethernet, or the CXP or CFP supporting 100 Gigabit Ethernet, potential front panel capacities ranging anywhere from 480 Gb/s to 3.2 Tb/s were observed. These front panel capacities could create backplane requirements ranging anywhere from 3.2 to 44.8 Tb/s depending on the specific system configuration. A comparison was made between the then existing 10 Gb/s signaling technologies against potential 25 Gb/s signaling to understand the impact on the total number of copper differential pairs needed to support various backplane capacities. Figure 4 illustrates the impact of 10 Gb/s vs. 25 Gb/s on the total number of differential pairs needed when supporting various backplane capacities for various switch fabric configurations. Note that as the capacity requirement of the backplane increases, the ability to support an actual total capacity of the backplane with 10 Gb/s per lane signaling becomes questionable [10].

The challenge with electrical backplanes as compared to copper cabling is that they are essentially custom-designed. There are a multitude of factors that influence the electrical performance of the backplane channel: the FR4 board materials, trace geometries, surface roughness of the copper traces, and the actual system configurations, among other characteristics. This is further complicated by the cost sensitivity of channels, where material costs alone could increase the cost of a backplane by 500 percent depending on the materials compared [11].

The large variation in electrical performance and sensitivity to cost resulted in the development of 100 Gigabit Ethernet backplane objectives targeting different performance/cost targets:

- Define a 4-lane PHY for operation over a printed circuit board backplane with a total channel insertion loss of ≤ 35 dB at 12.9 GHz.
- Define a 4-lane PHY for operation over a printed circuit board backplane with a total channel insertion loss of ≤ 33 dB at 7.0 GHz.

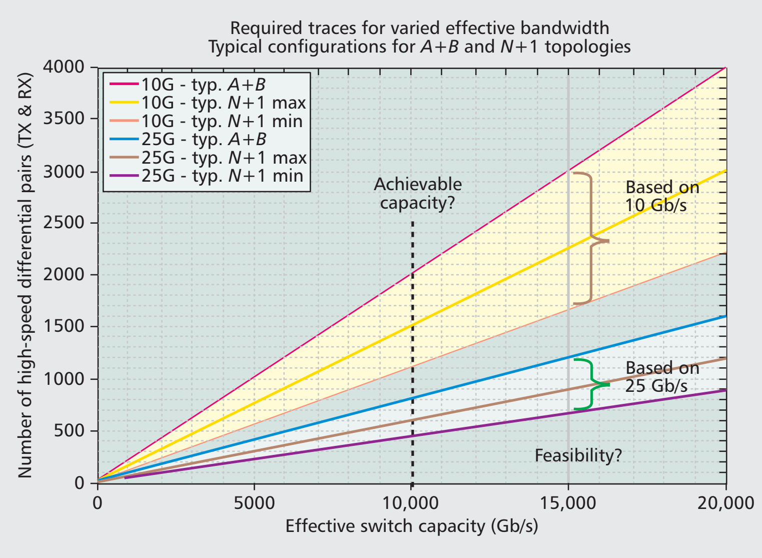

Figure 4. *Impact of signaling speed on switch capacities.*

The IEEE P802.3bj Task Force selected an non-return to zero (NRZ)-based solution for the 35 dB @ 12.9 GHz objective, while a PAM-4 based solution was selected for the 33 dB at 7.0 GHz objective.

In addition, the IEEE P802.3bj Task Force agreed on an objective to create a new 100 Gigabit Ethernet x4 twin-axial cable solution, to define a 4-lane 100 Gb/s PHY for operation over links consistent with copper twin-axial cables with lengths up to at least 5 m. An NRZ-based solution was selected to address this objective.

In addition to these objectives, the Task Force has agreed to define an optional EEE mode for PHYs described above.

At the time of writing, the IEEE P802.3bj 100 Gb/s Backplane and Copper Cable Task Force released D1.3 of the specification, and is working toward the next phase of the process (Working Group ballot). The anticipated ratification of the standard at the time of writing is expected in the first half of 2014.

For more information about the Task Force, please see http://www.ieee802.org/3/bj/index.html.

Optical P2MP Links and Evolution of EPON

The EPON is a relatively new addition to the family of Ethernet standards, with the first standard for this technology (1G-EPON, operating at the symmetric data rate of 1 Gb/s) published in 2004. In 2009, a higher-speed version of EPON was standardized, supporting the symmetric data rate of 10 Gb/s as well as an asymmetric data rate of 10 Gb downstream (toward the customer) and 1 Gb/s upstream (from the customer). Supporting nominal distances of 20 km (or more) and a nominal split of 1:32 (or

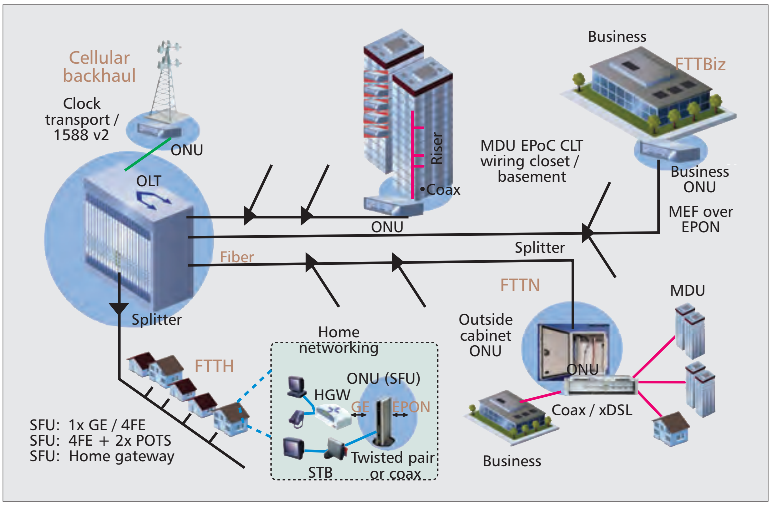

Figure 5. *Examples of various EPON deployment scenarios.*

more) with three available power budget classes, EPON is used in a variety of deployment scenarios, some of which are shown in Fig. 5. More details about EPON technology can be found in [2, 3], including definition of individual power budget classes.

However, along with the commercial success of this technology, the end users of EPON systems raised a number of questions on the deployment strategies for areas that have not been covered with EPON until now:

- Leveraging EPON architecture in a cost-effective way in rural areas with lower customer density
- Increasing subscriber density in the central office (CO) to get a better return on investment from active devices and improve the port occupation ratios
- Decreasing the cost of connection per subscriber (or per megabit per second), sharing the available links among a larger population of end subscribers
- Serving more people at longer distances from the CO using IEEE 802.3 EPON equipment and avoiding proprietary solutions offered by some system vendors.

All these questions led to the creation of a new project (IEEE P802.3bk, Extended EPON Task Force), chartered with extending the coverage of EPON to higher split and/or longer distance, addressing the concerns brought forth by operators already using EPON or planning to use EPON in the near future. This future amendment will add support for new power budget classes with the option of a single-ended upgrade at the optical line terminal (OLT) through the replacement of OLT line cards. Simultaneously, it will also create a framework for the use of mid-network power budget extender (PBEx) devices, increasing the available link budget beyond what can be achieved with simple replacement of optical network unit (ONUs) and OLT transceivers. There are unique challenges specific to this approach, including preservation of existing deployed ONU and OLT devices, and supporting coexistence of 1G-EPON and 10G-EPON on the same distribution network. These challenges will be addressed by the IEEE P802.3bk Task Force, leading to the planned publication of the completed standard late in 2013.

From the technology point of view, this project is evolutionary in that it does not modify the existing EPON architecture, increase the data rate supported by the end stations, or add new types of devices. However, it does provide the ability for the operator to reach larger subscriber populations or remotely located subscribers, providing savings on capital and operational investment, as well as allowing customers to be served where previously no fiber-based services were available.

It is also expected that in the future a more revolutionary shift in the EPON architecture will take place, primarily to increase the channel capacity beyond 10 Gb/s downstream and perhaps also upstream. Before technical choices are made, though, a detailed study of the market drivers for such higher capacity systems is needed to make sure that the market requirements are addressed appropriately, and a cost-effective as well as backward-compatible solution is selected. This effort is expected to start within the EPON community in the near future, especially now that 10G-EPON deployments are on the rise and operators are beginning to consider the path for future network upgrades.

For more information about the Task Force, please see http://www.ieee802.org/3/bk/index.html.

EPON Protocol over Coax: Bringing the Copper and Optical Worlds Together

Deployment of gigabit-capable EPON (based on 1G-EPON and 10G-EPON) services by cable operators in both China and the United States has been increasing over the past several years. The market drivers in both markets are slightly different. Hybrid fiber-coaxial (HFC) deployments as well as Data over Cable Service Interface Specification (DOCSIS®) [4] are not as widely deployed in China as in other parts of the world. Cable operators there are looking at the opportunity to transparently extend EPON services over legacy coaxial cabling in multitenant dwelling units (MxU) and businesses. In North America, high-speed data (HSD) residential services are currently provided using DOCSIS technology. However, for competitive multi-gigabit business class services, cable operators are increasingly deploying EPON to capture market share leveraging Metro Ethernet Forum (MEF) [5] service performance and competitive service level agreements (SLAs), all managed by DOCSIS Provisioning of EPON (DPoE™ [4]) technology, developed jointly by operators, vendors, and CableLabs.

Another trend in the worldwide cable network industry is the step-wise migration from back-end legacy MPEG-2 transport to MPEG-4 video distribution via IP over Ethernet. In the future, a large Ethernet-based gigabit pipe to the home and business will be fundamental for cost-effective growth and evolution. Figure 6 shows examples of some of the target applications of the mix of EPON and EPoC technologies, leveraging fiber-deep access architectures of current networks and also reusing existing coaxial distribution infrastructure to the greatest extent possible.

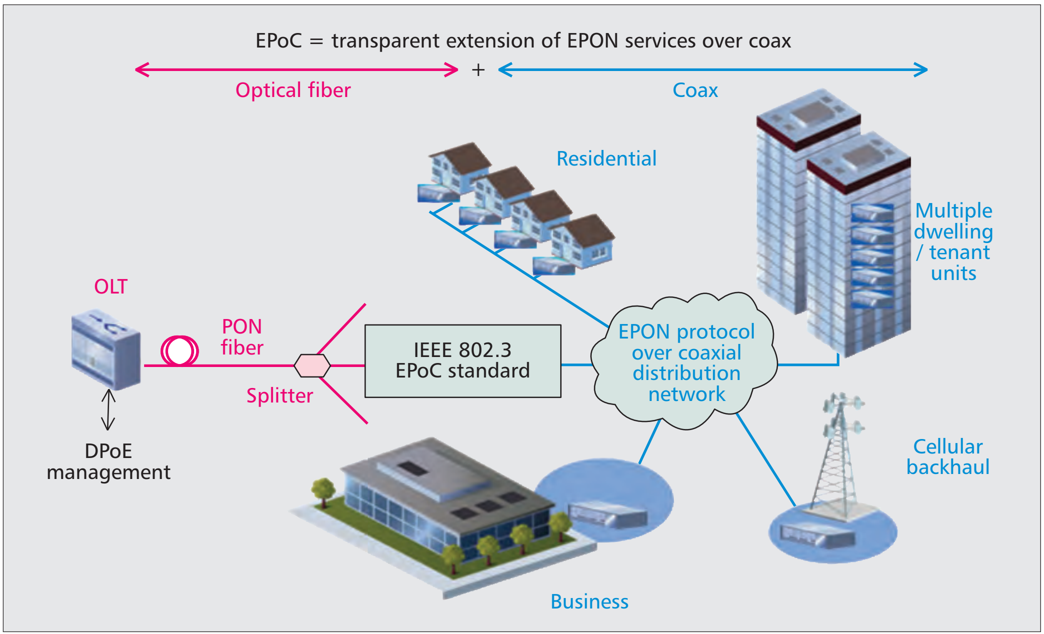

Figure 6. *EPoC applications for extending EPON services over coax.*

It is expected that in the future a more revolutionary shift in the EPON architecture will take place, primarily to increase the channel capacity beyond 10 Gbit/s in downstream and perhaps also in the upstream. Before technical choices are made, though, a detailed study of the market drivers for such higher capacity systems is needed

Both China and North American operators agree in the desire for the simplicity of Ethernet at gigabit speeds and collectively asked that the IEEE 802.3 Ethernet Working Group [6] create a standards effort for extending the operation of EPON protocols over coaxial distribution networks, a project that was called EPON Protocol over Coax, or EPoC for short. There are many opportunities where EPON has been deployed adjacent to or alongside of existing coaxial networks, and some customers are more opportunistically reached by simply extending EPON over coax. Key to this transparent extension is unified management, service, and quality of service.

IEEE Standards Status — In November 2011, following a successful Call for Interest, the IEEE 802.3 Working Group approved the creation of the EPoC Study Group [7]. In May 2012, the EPoC Study Group completed its draft Project Authorization Request (PAR), Responses to 5 Criteria, and a set of Objectives. Following the approval of these documents, the IEEE P802.3bn Task Force was created in August 2012, and after meetings in September and November of 2012, a timeline was adopted, estimating a draft submission to the IEEE 802.3 Working Group in the latter half of 2013 with the goal of having a published standard by the end of 2014.

In addition, a sub-Task Force was created to develop a time-division duplex (TDD) alternative mode of operation that would augment the primary frequency-division duplex (FDD) focus of the specification development effort.

For more information about the Task Force, please see http://www.ieee802.org/3/bn/index.html.

EPoC Architecture — For EPON, IEEE 802.3 standards define the MAC and PHY sublayers for a service provider OLT and a subscriber ONU. The fiber optical interconnecting medium uses two wavelengths for full-duplex operation, one for continuous downstream channel operation and another for upstream burst mode operation. The OLT MAC controls time-division sharing of the upstream channel for all ONUs.

Similarly, the EPoC architecture consists of a service provider coax line terminal (CLT) and a subscriber coax network unit (CNU). The EPoC CLT and CNU MAC sublayers will be substantially similar (if not the same) to respective layers found in the OLT and ONU, respectively. A new PHY will be specified for operation over coaxial distribution network (CxDN) media. Downstream and upstream communication channels will utilize radio frequency (RF) spectrum as assigned by a cable operator for their coax network.

System Models — Two system models have been under discussion in the EPoC Task Force, as shown in Fig. 7. The first is a CLT with one or more CNUs interconnected by a coaxial distribution network. The second is enabled by the future EPoC standard, but is outside the scope of the IEEE 802.3 Working Group. That is a traditional EPON with an OLT and multiple ONU devices together with one or more optical-to-coax media converters devices that attach between the PON and a CxDN using EPoC, permitting CNUs to appear as ONUs to the OLT. The industry will likely create products using this second model.

TIMING AND SYNCHRONIZATION IN ETHERNET (TIMESYNC)

Support for synchronization in Ethernet is rapidly becoming a critical feature, especially due to requirements of digital content distribution, video and audio systems with remote streaming, or even mobile backhauling. All these application areas require not only delay-guaranteed, engineered and strictly controlled links (in terms of QoS, bandwidth, and jitter), but also the abili-

The potential use of IEEE Std 802.3bf to support IEEE Std 1588v2 resolves one of the long-standing problems of this specific synchronization protocol, namely the lack of a standardized way to retrieve correlated information between the frame transmission time and synchronized time.

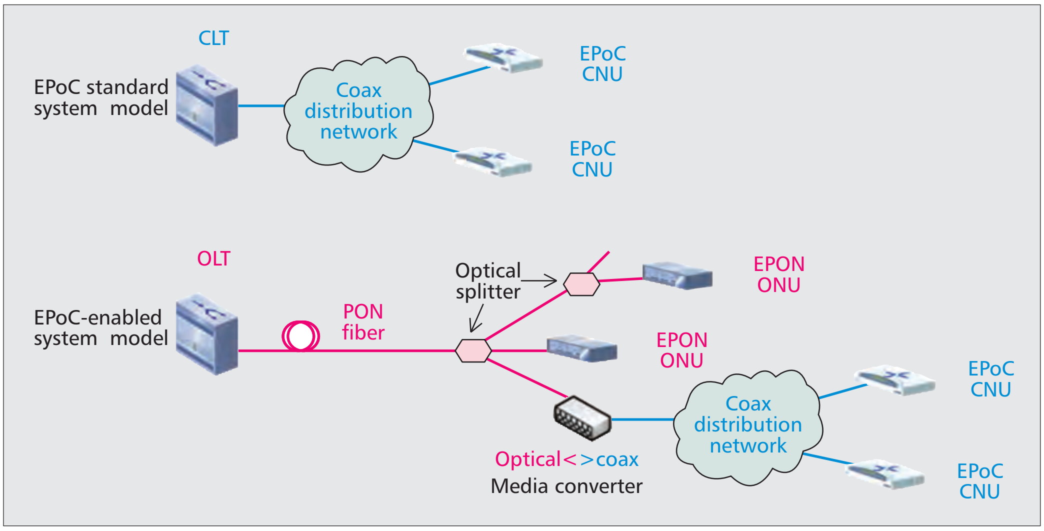

Figure 7. *EPoC standard and EPoC-enabled system models.*

ty to synchronize with a common reference clock to ensure proper operation of specific features of the given application. To address these requirements, the IEEE P802.3bf Task Force was created in 2009 to develop a method for "an accurate indication of the transmission and reception initiation times of certain packets, as required to support IEEE P802.1AS" [8].

The resulting architecture is presented in Fig. 8. This project added the following new features to the IEEE 802.3 architecture:

- Rx SFD Detect and Tx SFD Detect functions, responsible for detecting the reception and transmission of an Ethernet frame, respectively, and indicating this information to upper layers (time sync client) via the time sync service interface (TSSI).
- Set of managed objects and registers, providing the time sync client the ability to read ingress and egress latency information characteristic of the given PHY. This provides the time sync client with the ability to perform necessary synchronization calculations relative to the reference plane located at the bottom of the 802.3 stack (at the medium dependent interface).

The IEEE Std 802.3bf architecture was designed to provide direct support for the IEEE Std 802.1AS time sync client, operating on top of IEEE 802.3 PHYs. However, it was quickly identified that potential applications of the newly specified TSSI could also cover other synchronization protocols (e.g., IEEE Std 1588v2) and other proprietary use cases, which can benefit from information about transmit and receive path latencies as well as identification of the frame transition event through the RS sublayer.

The potential use of IEEE Std 802.3bf to support IEEE Std 1588v2 resolves one of the long-standing problems of this specific synchronization protocol: the lack of a standardized way to retrieve correlated information between the frame transmission time and synchronized time. Various proprietary mechanisms have been developed over the course of the last few years, some of them quite similar to the solution proposed in IEEE Std 802.3bf. It is expected that the TSSI will become a de facto standard for future implementations of IEEE Std 1588v2 protocol operating on top of Ethernet PHYs.

For more information about the Task Force (currently achieved due to completion of its charter), please see http://www.ieee802.org/3/bf/index.html.

Ethernet in New Markets: Applications in the Automotive Industry

The IEEE 802.3 family of standards provides a wide variety of solutions for data networks with many different operating speeds over copper wire, electrical backplanes, and various optical media. Recently, the global automotive industry has decided to deploy Gigabit Ethernet as a network backbone in automobiles and light trucks by 2020 [9]. 1000BASE-T as defined in IEEE Std 802.3 utilizes four twisted copper wire pairs. While this is not an issue for structured wiring plants, it results in a cable that is too heavy, costly, and cumbersome for vehicular use; see [9, slide 15] for an example of a typical passenger vehicle harness.

The Reduced Twisted Pair Gigabit Ethernet (RTPGE) Task Force (IEEE P802.3bp) is working on developing a robust 1 Gb/s copper PHY for this new market space. Estimates place the number of Ethernet ports in cars at around 300 million ports/year by 2019. The RTPGE PHY will allow for smaller and lighter cabling, and the use of a network backbone architecture that will make the in-car wiring harness easier to manufacture and lower in cost. In fact, the wiring harness in a car is the third most expensive component in the car (behind the engine and chassis), and also the third heaviest.

In 2012, modern cars have between 40 and 60 microcontrollers; high-end cars have over 120. Future sophisticated camera and control systems, vehicle safety devices (automatic braking, collision avoidance, etc.), infotainment, and GPS systems create large traffic volumes for the in-car network that previous automotive networks cannot handle.

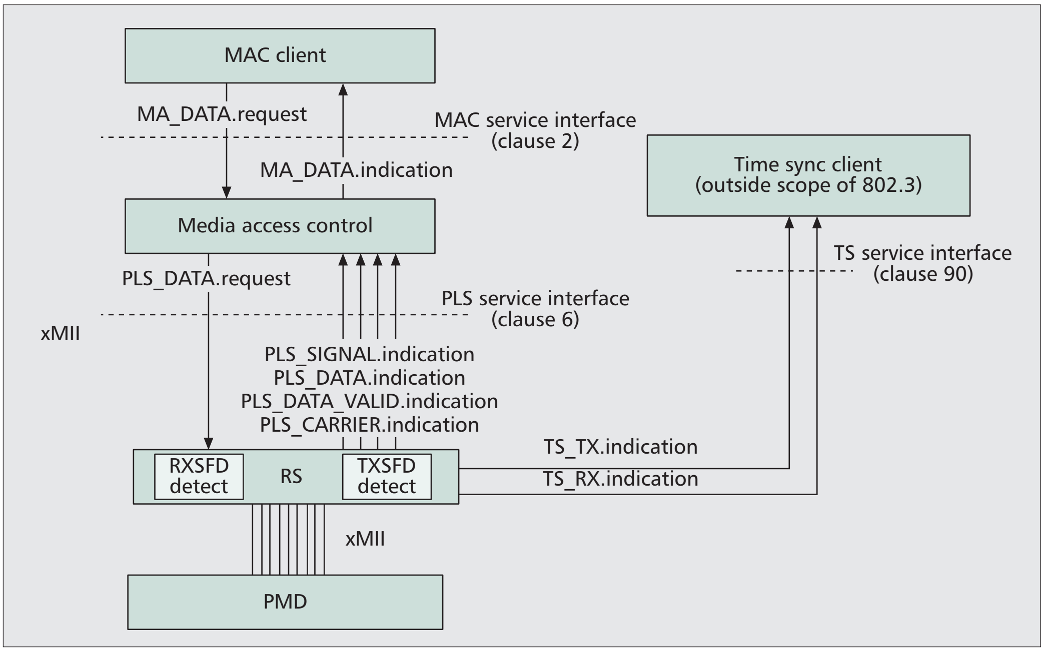

Figure 8. *Relationship between IEEE Std 802.3bf functions, TSSI, and remaining IEEE 802.3 layers. All clause numbers relative to IEEE Std 802.3.*

Ethernet as defined by the IEEE 802.3 Working Group, continues to evolve by adding support for higher data rates, new media types, and new features, extending into new markets and addressing new application spaces.

The number of microcontrollers or electronic control units (ECUs) in cars is expected to rise dramatically over the next decade as these functions become standard features in new cars. Automotive networks will also leverage other IEEE 802.3 technologies such as IEEE Std 802.3az EEE and IEEE Std 802.3bf time synchronization (discussed in the previous section) to provide an optimized network solution.

Ethernet is already used in aircraft for in-flight entertainment systems. The RTPGE PHY will allow for weight reduction as seen in ground-based vehicles. Weight savings in all modes of transportation result in lower fuel costs and effectively reducing the environmental impact of travel.

There is also growing interest in using the RTPGE PHY in more traditional IEEE 802.3 Ethernet applications. In data centers, RTPGE would allow for increased port density and higher throughput without the need for new cables. The work in RTPGE is generating a great deal of exciting discussion about new markets and applications that might be served in the 2015 and beyond timeframe.

For more information about the Task Force, please see http://www.ieee802.org/3/bp/index.html.

Management for Ethernet Networks

Ethernet, as defined by the IEEE 802.3 Working Group, continues to evolve by adding support for higher data rates, new media types, and new features, extending into new markets, and addressing new application spaces, as discussed in the previous sections. This evolution, though, may require changes in the managed objects stored in management information base (MIBs), allowing management systems to take full advantage of the newly added Ethernet features.

In order to provide a consistent up-to-date version of MIB definitions and eliminate dependence on external MIB definitions produced outside of the IEEE 802.3 Working Group, a project was started at the end of 2008 (IEEE P802.3be), targeting organization, update, and consolidation of managed object definitions provided in IEEE Std 802.3-2008, including as well the Logical Link Discovery Protocol Ethernet extensions provided in IEEE Std 802.1AB-2009 Annex F. In addition, the initial version of this standard incorporated and updated the MIB module definitions formerly defined in a series of RFC documents: RFC 2108, RFC 3621, RFC 3635, RFC 3637, RFC 4836, RFC 4837, RFC 4878, and RFC 5066. The final version of IEEE Std 802.3.1 was published in July 2011, containing both the definitions of individual MIBs, associated descriptions, and the machine-readable MIB files, available from the website of the IEEE P802.3be project.

IEEE Std 802.3.1 did not account for any of the published amendments to IEEE Std 802.3-2008, including IEEE Std 802.3at, IEEE Std 802.3av, IEEE Std 802.3az, IEEE Std 802.3ba, IEEE Std 802.3bc, IEEE Std 802.3bd, IEEE Std 802.3bf, and IEEE Std 802.3bg. A new project, IEEE P802.3.1a, was started in May 2011, after the publication of the base IEEE Std 802.3.1, with the intent of bringing Ethernet MIB definitions aligned with the status of the published base Ethernet standard (IEEE Std 802.3-2012). This project completed the final stage of the revision process (Sponsor Ballot) early in 2013 and became a published standard soon after.

For more information about both IEEE P802.3.1 and IEEE P802.3.1a Task Forces, please see http://www.ieee802.org/3/be/index.html and http://www.ieee802.org/3/1/index.html.

Summary

The work within the IEEE 802.3 Working Group is far from done, with the next generations of 40G and 100G links aiming for broader market

It can be expected that the innovation in the area of wired Ethernet continues in the years to come, bringing the same highly reliable and well understood networking philosophy to new markets, enabling new applications and making networking in general more ubiquitous.

adoption through increasing the cost effectiveness of solutions, decreasing the power consumption and complexity of compatible products. However, projects such as IEEE 802.3 Industry Connections Ethernet Bandwidth Assessment Ad Hoc (http://www.ieee802.org/3/ad_hoc/bwa/index.html) aim at establishing the market demand for the following generation of high-speed links, achieving 400G or even 1000G capacities, collecting requirements from the Ethernet ecosystem in terms of their technical feasibility and target application spaces. At the time of writing, a new Study Group for 40GBASE-T (http://www.ieee802.org/3/NGBASET/index.html) was created out of the July 2012 plenary meeting, focusing at this time on development of future project documentation for 40 Gb/s operation over twisted pair media. The IEEE 802.3 Working Group is thus looking for ways to expand the Ethernet market coverage high and wide, to support higher data rates, while also providing coverage for emerging markets, such as the automotive industry.

It can be expected that innovation in the area of wired Ethernet will continue in the years to come, bringing the same highly reliable and well understood networking philosophy to new markets, enabling new applications and making networking in general more ubiquitous.

References

[1] S. Kipp, "Storage Growth and Ethernet," http://www.ieee802.org/3/ad_hoc/bwa/public/sep11/kipp_01a_0911.pdf

[2] G. Kramer, *Ethernet Passive Optical Networks*, McGraw-Hill Professional, 2005.

[3] M. Hajduczenia and H. J. A. da Silva: "Next Generation Ethernet Passive Optical Networks: 10G-EPON," in *Wireless Edge, Optical Core — Next Generation Networks*, K. Iniewski, Ed., Wiley, 2010.

[4] DOCSIS® and DPoE™ are trademarks of CableLabs, Inc. More information on DOCSIS and DPoE can be found at the CableLabs specifications web site, http://www.cablelabs.com/specifications/.

[5] MEF home page: http://metroethernetforum.org/; for more information on MEF Carrier Class Ethernet "Class of Service" requirements, look under Information Center and then MEF Technical Specifications for document MEF 23.1.

[6] IEEE 802.3 Working Group home page: http://www.ieee802.org/3/.

[7] IEEE 802.3 EPoC Study Group home page: http://www.ieee802.org/3/epoc.

[8] Objectives for the IEEE P802.3bf "Ethernet Support for the IEEE P802.1AS Time Synchronization Protocol," Task Force, http://www.ieee802.org/3/time_adhoc/P802_3bf_objective_802_3_approved_1109.pdf.

[9] Reduced Twisted Pair Gigabit Ethernet PHY, Call for Interest, http://www.ieee802.org/3/RTPGE/public/mar12/CFI_01_0312.pdf.

[10] "100GbE Electrical Backplane / Cu Cabling Call-For-Interest," http://www.ieee802.org/3/100GCU/public/nov10/CFI_01_1110.pdf.

[11] Patel and D'Ambrosia, "Blade Servers and the Relative Cost Impact of Materials," http://www.ieee802.org/3/bj/public/nov11/dambrosia_02a_1111.pdf.

[12] Optical Interconnect Forum CEI-28G-VSR, http://www.oiforum.com/public/currentprojects.html.

Biographies

DAVID LAW is a Distinguished Engineer at Hewlett-Packard Networking, and has worked on the specification and development of Ethernet products since 1989. Throughout that time he has been a member of the IEEE 802.3 Ethernet Working Group, where he has held a number of leadership positions. He has served as Chair of IEEE 802.3 since 2008 and before that served as Vice-Chair from 1996. He has been a member of the IEEE-SA Standards Board since 2005 and is currently serving as the Vice-Chair of the IEEE-SA Standards Board and Chair of the IEEE-SA Standards Board Patent Committee (PatCom). In 2000 he received the IEEE-SA Standards Medallion, in 2009 he received the IEEE-SA Standards Board Distinguished Service award, and in 2012 he received the IEEE-SA International Award. He has a B.Eng. (Hons) in electrical and electronic engineering from Strathclyde University, Glasgow, Scotland.

DAN DOVE is Chair of the IEEE 802.3bm Task Force, which is working to define the next generation of 40G and 100G Ethernet optical technologies. He has been working at Applied Micro for a year, helping them to establish a technology strategy that will position them for success in data center networking. In addition to his responsibility to identify new technology opportunities, he represents them in key networking consortiums and standards bodies including IEEE 802.3 and OIF, and as a member of the Board for the Ethernet Alliance. Prior to his employment at Applied Micro, he was principal engineer for Physical Layer Technologies at Hewlett Packard's Networking Business unit. He worked at HP for 31 years, spanning a career that started in production assembly and worked through a technical ladder to his ultimate role as a master-level engineer. He has been working in the IEEE 802.3 Working Group since 1988, participated in many projects, and led the 802.3ak 10GBASE-CX4 project to successful completion. He holds 21 patents including the broadly used Auto-MDIX patent used for twisted pair physical layers.

JOHN D'AMBROSIA is the Chief Ethernet Evangelist in the CTO Office of Dell. In this capacity he has been an industry leader in the development of Ethernet-related technologies since 1999. Currently, he is chairing the IEEE 802.3 400 Gb/s Ethernet Study Group. In addition, he is a founder of the Ethernet Alliance, and is currently serving as Chairman of the Board of Directors. Prior to these efforts, he served as Chair of the IEEE P802.3ba Task Force, which developed the specifications for 40 Gb/s and 100 Gb/s Ethernet, the IEEE P802.3bj 100 Gb/s Backplane and Copper Cable Task Force, the IEEE 802.3 Ethernet Bandwidth Assessment Ad Hoc, as well as chair of the OIF's Market Awareness & Education Committee. He writes a blog for *EE Times*, called Ethernet Watch, which may be found at http://www.eetimes.com/electronics-blogs/4205110/Ethernet-Watch

MAREK HAJDUCZENIA (marek.hajduczenia@zte.pt) is currently employed by ZTE Corporation and focuses on R&D for EPONs. He is involved in the IEEE 802.3 WG, where he was an Associate Editor for IEEE Std 802.3av-2009 (10G-EPON) and Chief Editor for IEEE 802.3bf-2011, and is currently acting Task Force Chair for IEEE P802.3bk (Extended EPON). He supports IEEE P1904.1 SIEPON with technical contributions and acts as Chief Editor. He also participates in CableLabs in development of the DPoE™ series of specifications. He received his Ph.D. degree in the area of electronics and telecommunication from the University of Coimbra, Portugal. Currently, he holds more than 35 international and European patents.

MARK LAUBACH [SM] is a technical director in the Broadband Communications Group of Broadcom Corporation. He is also the chair of the IEEE P802.3bn EPON Protocol over Coax Task Force. He has more than 20 years of experience spanning broadband, gigabit, and cable access network technologies including being CEO and president at Broadband Physics, vice president and CTO at Com21, and a senior engineer at HP Labs. He is an inventor of several U.S. patents on cable modem technology. His past and current activities in standards include IEEE, IETF, SCTE, and the ATM Forum. He is an IEEE 802.3 voting member and an SCTE Senior Member. He is the principal co-author of a cable modem technology book, *Breaking the Access Barrier: Delivering Internet Connections over Cable* (Wiley, 2000). He holds a B.S.E.E. and an M.S.C.S. from the University of Delaware under Professor David J. Farber.

STEVEN B. CARLSON is the president of High Speed Design, Inc., a Portland, Oregon-based consulting company. He has over 38 years' experience in embedded control systems and networking for the entertainment and energy management industries. He currently serves as Chair of the IEEE P802.3bp Reduced Twisted Pair Gigabit Ethernet Task Force and is Executive Secretary of the IEEE 802.3 Ethernet Working Group. He served as Chair of the IEEE 802.3af-2003 DTE Power via MDI project, usually referred to as "Power over Ethernet," and the IEEE 802.3bf-2011 Time Sync Task Force, and was a founder of the Professional Light and Sound Association's Technical Standards Program, ANSI E1-Entertainment Technology.