INVITED PAPER *Special Section on Optical Access System for Social Life*

PON Convergence

Frank EFFENBERGER[†a)], ***Nonmember***

SUMMARY This paper discusses the concept of PON standards convergence. The history of PON standardization is reviewed in brief as a way to explain how the industry arrived at its current divergent form. The reasons why convergence is favorable are enumerated, with a focus on what has changed since the last round of standardization. Finally, some paths forward are proposed.

key words: *PON, access, convergence, standards*

1. Introduction to the History of PON Standards

The development of passive optical network (PON) systems has resulted in two families of standards being written by various groups [1]–[5]. The ITU-T PON systems are defined primarily in ITU-T recommendations, which cover the system, physical medium dependent (PMD), transmission convergence (TC), and management layers. The Broadband Forum (BBF) has developed documentation on how to build a total access system using the ITU-T technology, and also detailed test plans to confirm conformance with those recommendations and interoperability between optical network units (ONUs) and optical line terminations (OLTs). The IEEE PON systems are defined by a combination of PMD and TC layers defined by IEEE P802.3 and management/system layers defined in IEEE 1904.1. There are also conformance test plans developed in the IEEE P1904.2 project. In fact, the 1904.1 standard includes three "packages", labeled A, B, and C; which represent the common practices found in the US, Japan, and China, respectively.

As time has passed, the transmission technology has advanced in speed, and the two families have been in a semi-competitive race. The sequence of systems began with ITU-T A-PON and B-PON, then IEEE GE-PON, followed quickly by ITU-T G-PON (Fig. 1). After a short pause, there was IEEE 10GE-PON, following quickly by ITU-T XG-PON. Most recently, ITU-T has defined both NG-PON2 and XGS-PON, and IEEE has begun work on 100GE-PON. The simple-minded analysis that the system with the highest rate must be the better one has only helped to sustain this competition. This back-and-forth battle has sometimes left the industry wondering which system is next or which will be cost effective. This has worked in some cases to slow down the acceptance of PON technology, because every new entrant to the industry must first make the choice of which PON system to support.

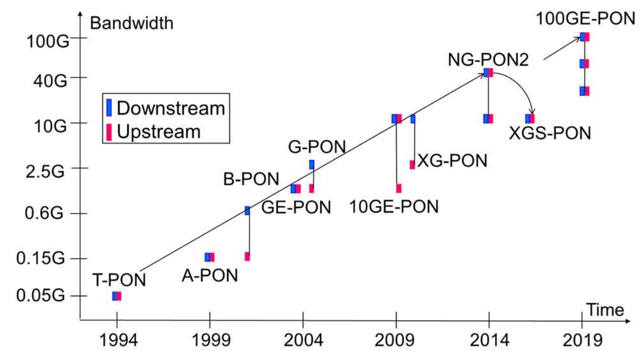

Fig. 1 The variety of PON systems over time.

The broadband world is fortunate that the truly cost sensitive aspects of the PON (the optical distribution network (ODN) and the PMD) are common between the two families (for the most part). The ODN's were driven to commonality because operators wanted to have a single design that could have a lifetime much longer (~30 years) than the technology life-cycle (~7 years). The PMD's were driven to commonality because they all operate on the same ODN (and the optical channel is what drives the PMD requirements), and they all want to leverage the same common pool of optical components (similar wavelengths, detector and laser types).

Unfortunately, the TC layers and some aspects of the PMD were constructed differently. These differences were due to the development forums in play, which had different membership and different design philosophies. ITU-T is a group that is heavily influenced by telephone network operators, and attended by their system vendors. Systems developed in ITU-T tend to be more conservatively specified, and work to develop a high level of efficiency and integration of features and services. The desire for efficiency led to decisions such as rather tight PMD timing parameters and data packet fragmentation. In time division multiple access (TDMA) PONs, the PMD requires time to turn on and stabilize, and to turn off, and these overhead times a directly subtract from the data rate delivered to the customer. The desire for full service integration drove solutions like asynchronous transfer mode (ATM) support, and the use of periodic framing (to provide timing references for services like voice).

Manuscript received May 14, 2017.
Manuscript revised August 11, 2017.
Manuscript publicized October 18, 2017.
[†]The author is with Futurewei Technologies, Plano, TX 75035, USA.
a) E-mail: frank.effenberger@huawei.com
DOI: 10.1587/transcom.2017OAI0001

In contrast, the IEEE 802.3 working group tends to have a different mix of members, primarily from the data communications industry, with more participation from the semiconductor device sector. This, and the relative absence of network operators, has led to a design approach that is more idealistic, and that is less focused on transmission efficiency in favor of simplicity. What's more, IEEE 802.3 maintains a tight control over the scope of problems they will address, so as to focus everyone's effort on the basic Ethernet transmission technology. This basically rules out the development of an integrated full service system in IEEE 802.3. To fill this gap in the specifications, the IEEE 1904.1 project was begun. Between the development of the original GE-PON standard and the 1904.1 project, different GE-PON ONU management systems were developed by several regional operators. The more important of these schemes were standardized after the fact (the three packages mentioned above), but they remain fundamentally different.

2. The Convergence Movement

The desire to pull these two families together has existed for a long time. In 2001, there were efforts made to align the ITU-T and IEEE PON efforts, but ultimately these failed due to the fundamental differences of design philosophy. What was worse, many of the vendor participants' own self interest was in developing their own new system, or developing a system that required technologies that they alone had. It is fair to say that both sides made design mistakes. For instance, ITU-T G-PON originally supported ATM transport, but this was removed some years later due to lack of use. Also, G-PON's use of tight timing parameters on the OLT receiver delayed introduction of the system. On the other side, IEEE GE-PON's lack of packet fragmentation (and its necessary complications on dynamic bandwidth allocation (DBA)) was a fault, as was the insistence on defining low loss budget values (PX10 and PX20) that were impractical from an operators' point of view.

Over time, the groups have learned and improved, largely due to real world application experience, and sometimes through the observation of the other. Already mentioned was the deprecation of the ATM features of G-PON. In the later XG-PON rendition of that system, the OLT receiver timing was made adjustable to meet whatever timing was most economical. 10GE-PON adopted the now standard 29 dB loss budget, and the 1904.1 and .2 projects filled the system engineering void left bare by the 802.3 project. The result is that XG-PON and 10GE-PON are two systems that are basically the same in capability and function, but different in their detailed implementation.

Most implementers agree that this situation is a waste of effort, because each system requires the same development effort to create, verify, deploy, and support, while having two systems does not increase our broadband capability at all. Currently, there is a growing movement to converge the PON technologies in the coming 25G generation of systems. Beyond the natural desire to have one set of standards to work from, there are a few new aspects this time.

The first is that the PMD is increasingly challenging, and the technical limitations of the system leave far fewer options open to decision. It is a fact that one of the biggest decision spaces to explore in any PON system is the optical spectrum plan. In previous generations, relatively little thought went into the choice of operating wavelengths, typically choosing rather wide spectrum bands at relatively arbitrary assignments. Over the generations of PONs, more and more spectrum has been occupied (simply taking it out contention for any future system). At 25 Gb/s operation, the fiber dispersion becomes a major impairment, so much so that operation in the low end of the O-band is desired to enable the avoidance of dispersion compensation technologies. While these technologies are available, they will add costs to the system, and the number one requirement for any PON system is low cost.

The second is the allowance of the frame pre-emption design in the 802.3 system. The biggest architectural disconnect between GE-PON and G-PON was the support of fragmentation or not. Recently, IEEE 802.3 has standardized a method of pre-empting a data frame to allow the transmission of a higher priority frame. The essential function that was added was the addition of a signal that stops one medium access control (MAC) from transmitting further data, so that the other MAC can begin its transmission. This very fine-grained control of the MAC operation also happens to be perfect to allow frame fragmentation for PON purposes. This design concept has already been accepted in the currently evolving 100GE-PON standard. So, in this generation frame fragmentation is no longer an issue.

The third reason why convergence is more likely is that all parties realize the importance of channel bonding. One of the key advantages of TDMA-PON is that any ONU on the PON can take advantage of the full bandwidth of the PON (through DBA, and subject to the traffic contract involved). This flexibility is very useful, and it is a commercial fact that many users rate the quality of their broadband service not on its average throughput, but on its peak bandwidth. The demand for more peak bandwidth has been recognized in both ITU-T and IEEE, and both groups have been working on various concepts on exactly how to distribute user data over multiple channels. For the next generation of PON, both groups are coming to the same conclusion, which is to distribute frames in manageable units (in 100GE-PON, they are called Envelope Quanta) over the available channels in a simple round-robin fashion. In this way, each packet can achieve the peak rate, and packet order is maintained very naturally.

The fourth is the increasing cost of integrated circuit development. Each process node of complementary metal oxide semiconductor (CMOS) is more expensive than the last, particularly considering the non-recurring engineering costs (design, verification, and mask sets). In previous generations of PON, chip designers handle multiple standards by implementing both on a single chip, and then just switch one off. This becomes wasteful for the 25G generation since it will be using the more advanced CMOS process where even idle circuits consume significant power due to their leakage current, and because the amount of circuitry (gate count) is certainly going to scale at least with the data rate, and probably even faster than that.

The fifth is the recent ascendancy of NetConf / YANG as the new network management scheme for communications networks. This reconstruction is driven by many diverse business needs, and it seems to be irresistible. One of the classic problems of management systems is that the current users grow very attached to their own system, and do not want to change. Politically, it is difficult to pick one winner, since that would unduly advantage one user group over the others. NetConf can be the common language to which all the previous management systems can evolve towards. To say it another way, everybody will be equally challenged with the new system.

For all these reasons, PON convergence may succeed in the near future. A strong indication of that is a recent public statement made by the leaders of all the major PON standard developing organizations (SDO): BBF-fiber access network (FAN), full services access network (FSAN), IEEE 802.3ca, IEEE 1904, ITU-T Q2/15 [6]. This statement reviews the ad-hoc sessions that have been held in the various groups, and the general support for the concept of convergence. It also puts forward the fact that while direct collaboration agreements between the groups are unlikely, convergence can still be achieved through a ‘grass-roots’ approach. That is, if a substantial majority of all the parties agree that convergence is in their interest, then their contributions and comments will support convergence. Everyone will do what is right, which is convergence.

3. What is Convergence, Really?

This is indeed a good question. A great deal of focus has been placed on the 100GE-PON system, as this system is currently under development. Figure 2 shows one proposed work arrangement of the three standards development organizations for the 100G-PON system (note, we remove the “E” from the acronym, because this is a generalized PON). The acronyms in Fig. 2 are physical coding sublayer (PCS), multi-point reconciliation sublayer (MPRS), and multi-point control protocol (MPCP). The IEEE 802.3 group could develop the physical layers of the system (PMD and PCS), as well as the “lower MAC” part (MPRS and MPCP). The ITU-T could develop the more system-level aspects of the system, including the “upper MAC” functions such as DBA, PON protection, and power saving. The Broadband Forum could develop the NetConf and YANG management schema for all the systems. Importantly, all of this work will be done to ensure backward compatibility with existing systems to the largest degree possible, and to develop evolution paths that minimize any discontinuity to the PON community.

It should be noted that this work arrangement would pertain only to the 100G-PON system as currently envisioned in IEEE 802.3ca. Other systems are conceivable, such as point-to-point and wavelength division multiplex (WDM)-PON systems. Consider a hypothetical system which uses sub-band coarse WDM transmitters to carry Ethernet formatted data. Such a system would likely have a SDO task distribution as shown in Fig. 3. The optical channel and PMD functions would be specified by ITU-T, while the PCS and MAC functions would be specified by IEEE 802.3 (likely just a reuse of already existing specifications). The higher layer functions would again be specified by ITU-T, and perhaps this system is so simple that it doesn’t require BBF

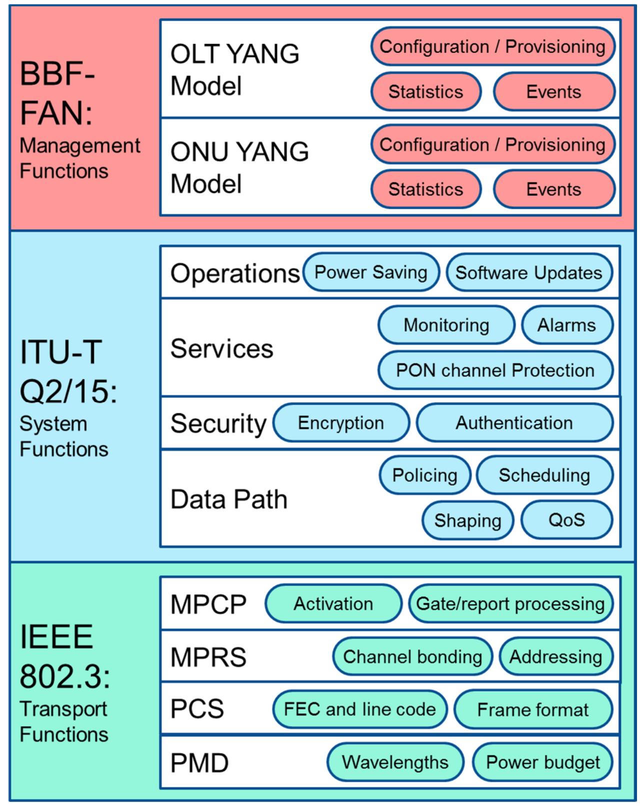

Fig. 2 Proposed 100G-PON standards layout.

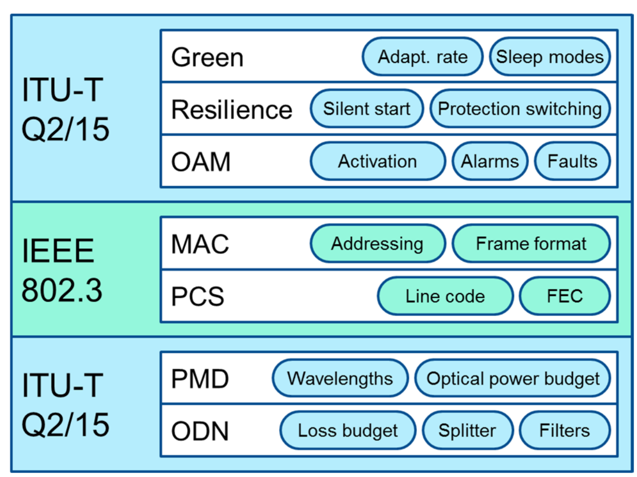

Fig. 3 Possible sub-band CWDM-PON standards layout.

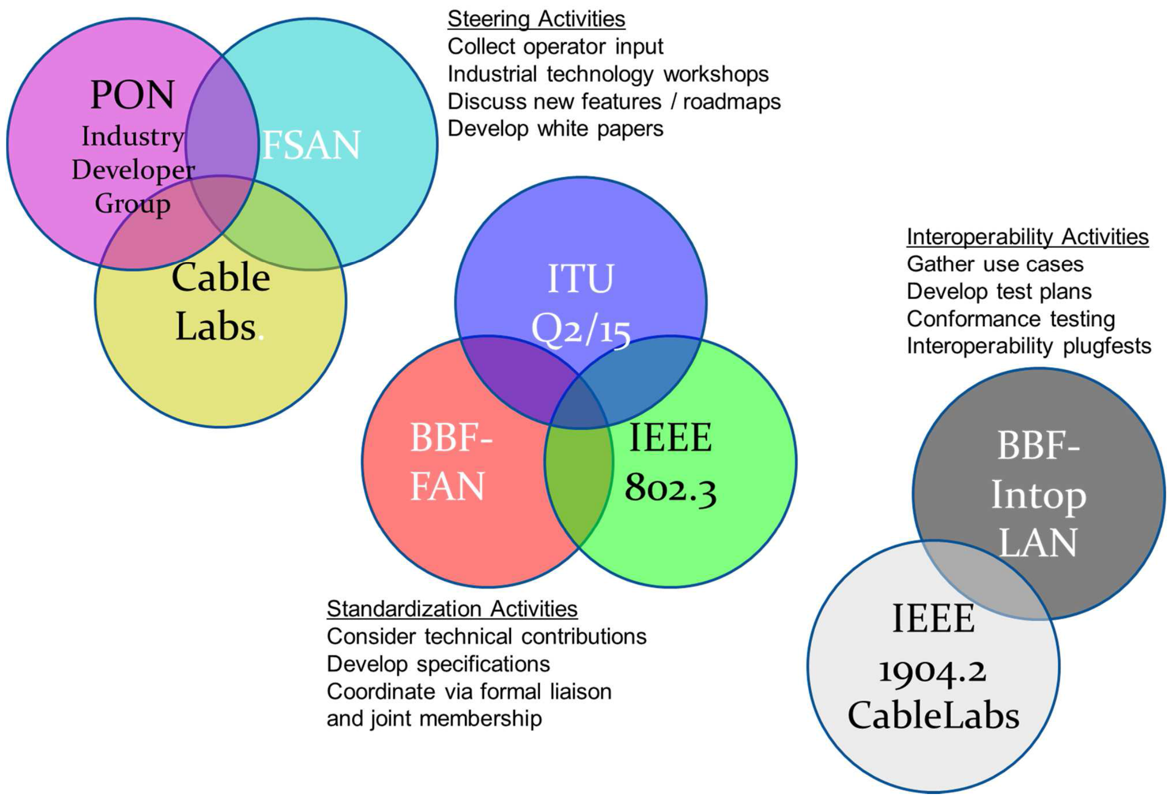

Fig. 4 All PON groups working together.

management support. It should be observed that this kind of arrangement was already used in the G.9801 (G.epon) standard system - proving that it can be done.

Of course, we don't know what the future may hold, and so the arrangement of SDOs must remain flexible. However, the general idea of convergence stays the same, and that is the industry will work as a whole to identify new requirement sets (problems) and new systems (solutions), and then to decide as a group what is the best way to standardize them. The structure for each project could be different, or evolve over time. The basic goal would be to avoid having two solutions for the same problem.

In the future, all the PON SDOs should work together as a team. This is diagramed in Fig. 4. There are three general categories of organizations: steering committees, formal SDOs, and interoperability efforts. The pre-standard steering groups include FSAN, CableLabs, and a new group named the PON Industry Development Group (PIDG), which has been proposed to the BBF. These groups work to uncover new problems in access and the new technologies that might be useful to solve them. The output would be a fully described set of system requirements that could then drive the SDOs to develop the specifications of the presumed system. This they would do collaboratively, finding a work arrangement that is maximally efficient. Lastly, the resulting standard would be implemented, and the resulting systems would be tested for conformance and interoperability by the testing organizations.

4. Conclusions

The development of PON standards has had many ups and downs, with some false starts and reckless competition. The result has been multiple systems that are remarkably similar in function but different in form. We now have another chance to converge the development of future PON systems. There are many reasons why this may be successful at this juncture of time.

References

[1] F. Effenberger, H. Ichibangase, and H. Yamashita, "Advances in broadband passive optical networking technologies," IEEE Commun. Mag., vol.39, no.12, pp.118–124, 2001.
[2] F. Effenberger, D. Cleary, O. Haran, G. Kramer, R.D. Li, M. Oron, and T. Pfeiffer "An introduction to PON technologies [Topics in Optical Communications]," IEEE Commun. Mag., vol.45, no.3, pp.S17–S25, 2007.
[3] J.I. Kani, F. Bourgart, A. Cui, A. Rafel, M. Campbell, R. Davey, and S. Rodrigues, "Next-generation PON-part I: Technology roadmap and general requirements," IEEE Commun. Mag., vol.47, no.11, pp.43–49, 2009.
[4] F. Effenberger, H. Mukai, S. Park, and T. Pfeiffer, "Next-generation PON-part II: Candidate systems for next-generation PON," IEEE Commun. Mag., vol.47, no.11, pp.50–57, 2009.
[5] F. Effenberger, H. Mukai, J. I Kani, and M. Rasztovits-Wiech, "Next-generation PON-part III: System specifications for XG-PON," IEEE Commun. Mag., vol.47, no.11, pp.58–64, 2009.
[6] http://www.lightreading.com/gigabit/fttx/sdos-team-up-on-pon-convergence/d/d-id/731234

Frank Effenberger has worked in the optical access field at Bellcore, Quantum Bridge Communications (Motorola), and Futurewei Technologies, where he is now the Vice President of the fixed access network lab. Frank is a Fellow of the OSA, the IEEE, and Huawei.