Future bandwidth demand favors TDM PON, not WDM PON

Ed Harstead
Alcatel-Lucent,
Ed.Harstead@alcatel-lucent.com

Abstract: Claims to the contrary notwithstanding, quantitative trends indicate that TDM PON bandwidth supply is growing faster than subscriber bandwidth demand, and TDM PON will deliver future ultra-high speed services far more efficiently than WDM PON.

OCIS codes: 060.0060 Fiber optics and optical communications; 060.4258 Networks, network topology

1. Introduction

Since the 1990s, WDM PON has been studied and proposed as the successor to TDM PON for FTTH networks [e.g. 1]. The primary rationale is that to meet future subscriber bandwidth demand, the required aggregate TDM PON bit rate will become so high that lower speed dedicated wavelengths will become more cost-effective.

In this paper we quantitatively examine the historical trend of the bandwidth supplied by TDM PON systems, and forecast residential bandwidth demand, in order to answer several questions related to large scale residential FTTH deployment: are current TDM PON systems (GPON) at risk of running out of bandwidth in the near future? How long will 10G (TDM) PON systems satisfy bandwidth needs, and what technology should succeed it? We will also examine the nature of broadband demand to answer whether shared bandwidth systems like TDM PON or dedicated bandwidth systems like WDM PON are better suited for future ultra-high speed services.

2. TDM PON bandwidth doubles every two years.

For the purposes of consistently tracking the technology, we restrict our attention to TDM PONs[1] that have been engineered to meet the cost targets for FTTH[2] and were commercially deployed (beyond field trials). Table 1 is a list of first important commercial deployments for major phases in the evolution of TDM PON, and their downstream and upstream line rates, including projected commercialized deployments of 10G PON. The downstream data is plotted in Figure 1, and indicates a trend of approximately 40% year-over-year (YoY) growth. That is, for almost two decades the TDM PON downstream line rate is doubling every 2 years. Therefore a putative 40G (TDM) PON would be predicted in 2016. We shall see below if such a thing will be required.

PON type	Commercial Deployment	Year	Line rate (Mb/s)	
			Down	Up
Narrowband	Deutsche Telekom OPAL	1995	29	29
Narrowband	NTT Pi PON	1997	49	49
ATM PON	NTT	2001	155	155
BPON	NTT West	2003	622	155
EPON	NTT East	2004	1000	1000
GPON	Verizon FiOS	2007	2488	1244
10G PON (IEEE)	t.b.d.	2012	10000	1000
10G PON (ITU-T)	t.b.d.	2013	9953	2488

Table 1. Commercial deployments of TDM PON by generation.

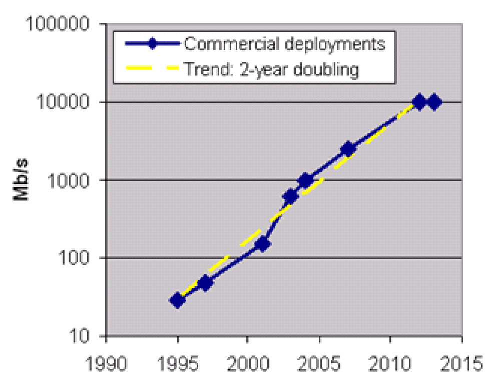

Figure 1. Evolution of TDM PON downstream rate.

3. Streaming video bandwidth demand and headroom for bursty services.

In a triple play access network, the source of largest bandwidth demand is streaming video, which could be delivered via a managed IPTV service and, increasingly, video delivered over-the-top via the Internet. While it would be difficult to find publicly available statistics to quantify and forecast this demand, we can create an upper bound. The first step is to estimate the current maximum bandwidth required to deliver streaming video by an IPTV operator. We'll take as our example AT&T, which offers "TV" service as part of its well-known 24 Mb/s

---

[1] We define a TDM PON as a point-to-multipoint passive fiber system achieving multiple access using TDM in the downstream direction and TDMA in the upstream direction. "TDM" does not refer to layer 2 transport; current TDM PON standards and products are packet-based.

[2] even if some actual deployment was in fiber-to-the-curb or fiber-to-the business applications.

downstream "U-Verse" offer. Versus incumbent cable and satellite providers, U-verse has been ranked high in customer satisfaction and has taken market share. It serves the U.S. market which is high among global leaders with respect to TVs per capita and HD TV penetration. In other words, 24 Mb/s is a viable bandwidth for providing successful digital video service in 2010.

To forecast demand, the next step is to pick an upper bound for the growth rate of that video consumption. The main drivers are the trend to video content with higher information rates, and possibly an increase in the number of video "tuners" (TVs, DVRs, and PCs) in the home. We propose 20% YoY, or a doubling every 4 years, as shown in Figure 2. An example of a stream mix evolution that could support that growth is shown in Table 2[3], where the subscriber starts with three HD streams and one SD stream, upgrades all his/her screens in the home to HD by 2011, and replaces all those with 3D screens by 2013. After that point, supporting 20% YoY would require adding more and more 3D screens and/or upgrading to a post-3D video standard, assuming one could be already available in consumer electronics by around 2015. This represents an early adopter making prodigiously aggressive investments in consumer electronics upgrades. Further, in the interest of setting an upper bound, in Table 2 we have ignored advances in video compression technology that continuously reduce encoded video bit rates, and we ignore the possible downward trend toward in video tuners, since the need for DVR recording of broadcast content will be eroded by the increase in viewing content that is cached in the operator's network or in the Internet. To sum, there is no present indication that the assumed 20% YoY growth will underestimate future demand for streaming video, and therefore the curve in Figure 2 should provide an upper bound.

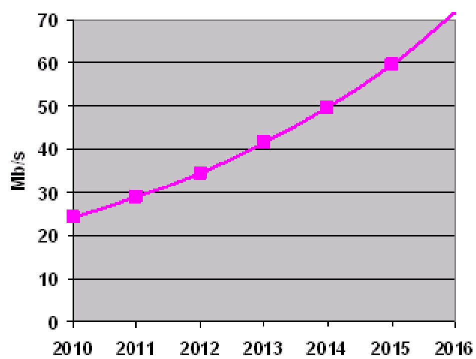

Figure 2. Upper bound of streaming video bandwidth demand.

Year	Video (Mb/s)	Number of TVs		
		SD	HD	3D
2010	24	1	3	
2011	29		4	
2012	35		2	2
2013	41			4
2014	50			5

Table 2. A stream combination to fill the upper bound

So far we have ignored bandwidth demand for bursty services, e.g. web browsing and file transfers. In an access system, the available bandwidth not consumed by streaming video would provide headroom for bursty services. The U.S. FCC has cited a goal of 100 Mb/s per home. Google has announced an initiative to deploy 1 Gb/s to the home. As we have seen, it will take a long time for streaming video to fill a 100 Mb/s pipe, let alone 1 Gb/s. The only way to justify this bandwidth would be for bursty traffic, e.g. extremely fast download/ upload of very large chunks of non-real-time data, e.g. music or video file transfers.

4. Can TDM PON meet current and future bandwidth demand?

Figure 3 shows GPON to 10G PON[4] aggregate downstream bandwidth supply versus the video demand in Figure 2 multiplied by 32 subscribers per PON[5]. Note that simply multiplying by 32 is a worst case, as it assumes all 32 subscribers on the PON are early adopters, and it prevents the TDM PON from taking advantage of usage statistics and IPTV multicasting gain. (This is borne out by proprietary real world traffic data on a live North American triple play network, which indicates much lower peak aggregated demand for 32 users, as one would expect).

First, it is clear from Figure 3 that even in this worst case, GPON will be able to meet streaming video bandwidth demand for years to come. Secondly, since TDM PON bandwidth supply is increasing at least twice as fast as streaming video bandwidth demand, headroom for bursty services will only continue to grow. If we extended the demand curve, we would see that 10G PON maintains 1 Gb/s bandwidth headroom for bursty services until the year 2023, even under the severe assumptions made here. So while the trend of TDM PON bandwidth doubling every

[3] This assumes SD, HD, and 3D streams consume 2, 7, and 10 Mb/s of bandwidth respectively.

[4] For simplicity, we are ignoring the effect of TDM PON-layer overhead that would decrease usable bandwidth to a small degree.

[5] 32 is a nominal number representing e.g. a 1:32 optical split with 100% video take rate, or a 1:64 optical split with 50% video take rate. The demand curve in Figure 3 would be lower or higher for e.g. 16 or 64 video subscribers per PON respectively.

two years predicts a commercial 40G PON in the year 2016, the industry actually has much more time at its disposal.

Another way to look at the forecast is that in 2016 the early adopter will require 72 Mb/s of streaming video, Therefore, on a 10G PON with 32 early adopter subscribers at simultaneous peak consumption, there would still be about 7.5 Gb/s of downstream bandwidth headroom available for bursty high speed data transfer, which could be grabbed by any subscriber. This is about 8 times more headroom than provided by a dedicated 1 Gb/s wavelength. Simply put, to beat 10G PON, WDM PON requires 10 Gb/s per wavelength. But such a massive pipe would be virtually empty virtually all of the time. Dedicated bandwidth is highly inefficient for the delivery of bursty services. As an aggregated bandwidth medium, TDM PON is favored for handling bursty traffic, which is the only kind of traffic capable of filling such large pipes in this decade.

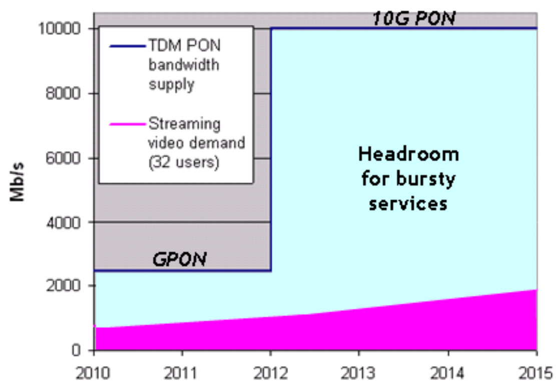

Figure 3. TDM PON downstream bandwidth vs. worst case demanded.

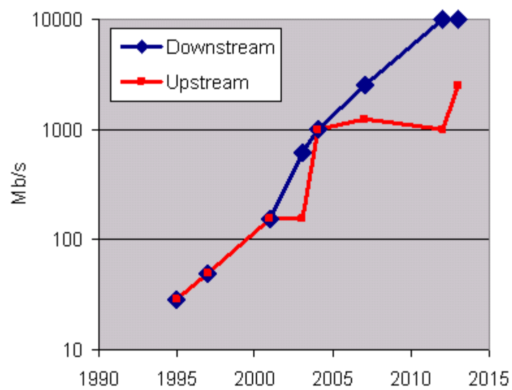

Figure 4. Evolution of TDM PON upstream bit rates.

4. Upstream bandwidth supply and demand

So far only downstream bandwidth supply and demand has been considered. The evolution of TDM PON upstream bit rates is presented alongside downstream bit rates in Figure 4. The upstream bit rates exhibit a random-walk-like trend due to two factors. Early proprietary systems delivered symmetrical narrowband (POTS and ISDN) services, and PON systems standardized by the ITU-T and the IEEE are (usually) constrained by their traditions of bit rates advancing by factors of 4 and 10 respectively. At any rate, the industry consensus seems to be that 10:1 is too asymmetrical and 1:1 symmetry is not really required.

How much upstream bandwidth is needed? Consider a worst case household engaging in several simultaneous video phone calls and running some surveillance cameras. It is probably safe to assume that these, especially the latter, will continue to be encoded in lower resolution and with higher compression than downstream video entertainment. An aggregate of 32 such households on a PON is unlikely to consume more than a few hundred Mb/s. TDM PONs with 1 Gb/s or more of upstream bandwidth will have a large headroom for very high speed upstream data transfers. It is hard to imagine how this might be exhausted.

4. Conclusion

In this paper we examined the trend in TDM PON bandwidth supply and forecasted an upper bound to residential video bandwidth demand. We concluded that current GPONs have ample bandwidth for years to come, and that 10G PONs will satisfy bandwidth demands beyond the foreseeable future. We also concluded that the only kinds of services that are able to fill huge FTTH pipes are future ultra-high speed file transfers, and that a shared bandwidth TDM PON more efficiently handles this bursty traffic than a dedicated bandwidth WDM PON. Accordingly, the WDM PON per-wavelength bit rate needs to be essentially equal to the aggregate TDM PON bit rate, otherwise the TDM PON solution will always be able to offer higher burst rates. For example a WDM PON competing with a 10G PON should provide 10G per wavelength. Given the ability of TDM PON to exceed bandwidth requirements, by far, the target for WDM PON should be cost parity with TDM PON. In other words, from these findings we believe that there is no bandwidth-based argument for an operator to pay a premium for WDM PON.

4. References

[1] R. Feldman, et. al., “An Evaluation of Architectures Incorporating Wavelength Division Multiplexing for Broad-Band Fiber Access”, J. of Lightwave Technology, v. 16, no. 9, Sept. 1998, pp. 1546-1559.