Future Proof Next Generation Access Networks -- The Case for FTTH

Wolfgang Fischer

Cisco Europe
wfischer@cisco.com

Abstract. The insatiable demand for more and more bandwidth lets access bitrates in broadband networks grow exponentially at a rate of about 50% year-over-year. None of the copper-based wireline access technologies will be able to cope with this demand in the long term. Also wireless networks see similar growth and require higher and higher backhaul speeds. There is industry-wide consensus that only fiber-based access networks will be able to support this growth. It can be expected that the new generation of passive infrastructure will experience the same lifetime as today's copper-based infrastructure. Therefore, it is of vital importance that this infrastructure will be built in a way that it is as flexible as possible, supporting future applications and technologies without costly upgrades, while still being affordable at the time of deployment.

1 Introduction

Broadband access has many facets today with wireline (e.g., DSL, cable, Ethernet, PON, etc.) and wireless (e.g., WiFi, WiMaX, 3G, LTE etc.) technologies being abundant. Within the physical limits of the respective media all these technologies have evolved to higher and higher bitrates over the last two decades. Typical access bitrates (i.e., the physical bitrates in the access network for Internet subscribers) have consistently grown by about 50% year-over-year on average. Using available data and extrapolating them into the future, 100Mbit/s can be seen as a typical access bitrate around 2014, and 1Gbit/s would be typical in 2019 (see Fig. 1). Many stakeholders even believe that this projection is rather conservative. In a number of networks, e.g., in Sweden, Netherlands, or Hongkong, 1Gbit/s is already state of the art. This consideration is supported by a significant increase in video-related applications over broadband networks which go far beyond the usual TV-related applications, and which will be characterized by higher definition, symmetrical bitrate requirements, ultra-high-speed bursts for non-real time video applications, and multiple streams per household. High-Definition video cameras have become affordable for the masses, and the efficient transport of the content created by them is one of many factors which continue driving bitrate growth.

There is widespread agreement in the industry that only direct Fiber to the Home (FTTH) will allow all of these applications to be supported in the longer term, with wireless technologies being an indispensable complement providing mobility.

I. Tomkos et al. (Eds.): BROADNETS 2010, LNICST 66, pp. 302–313, 2012.

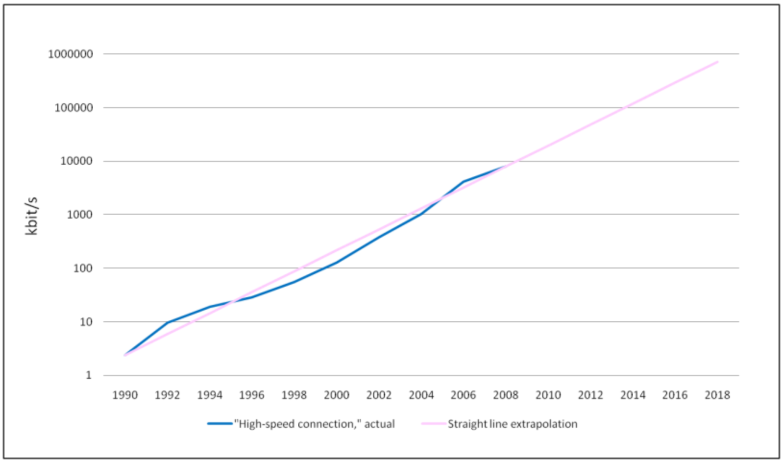

Fig. 1. Exponential growth of Broadband access bitrates (Source: Graham Finnie, HeavyReading)

2 Topology vs. Technology

While there is agreement about the need for fiber-based access networks, there is an ongoing debate about the best architecture for these networks, both in terms of the topologies of the fiber infrastructures as well as the technologies for transporting data over those fibers.

The two most fundamental types of topologies are Point-to-Multipoint and Point-to-Point.

Point-to-Multipoint topologies with multiplexing devices in the field are deployed in order to be used by one of the standardized Passive Optical Network (PON) technologies (EPON, GPON) and their evolutions, by Active Ethernet (Fiber to the Building, Fiber to the Curb), or by WDM-PON. The common fiber is shared and all subscribers connected to a tree use the same technology.

Point-to-Point topologies provide dedicated fibers between the Central Office / POP and the subscriber. As there is a dedicated transparent medium a wide choice of transport technologies is available. All currently existing point-to-point FTTH deployments use Ethernet (100BASE-BX10 or 1000BASE-BX10) over single fibers, but this can be mixed with other transmission schemes for business applications (e.g., FiberChannel, SDH/SONET, …), and even with PON technologies by placing the passive splitters into the POPs. Upgrades and modifications of a transmission technology can be performed on a subscriber-by-subscriber basis.

		Technology		
		Ethernet	TDM-PON (GPON, EPON, DPON, XG-PON, 10GEPON)	WDM-PON
Topology	P2P	Ethernet P2P	PON P2P	N.A.
	P2MP	Active Ethernet	PON P2MP	WDM-PON

Dedicated medium

Shared medium

Fig. 2. Classification of FTTH architectures

Fig. 2 shows a classification of FTTH architectures in terms of fiber topologies and technologies.

Ethernet P2P and PON P2MP are the classical architectures in the context of FTTH. Active Ethernet is used mainly for Fiber-to-the-Building (FTTB) or Fiber-to-the-Curb (FTTC) deployment models.

PON P2P has become more popular recently as a way to combine some aspects of P2P topologies and PON technology, by placing the splitter in the central office. This, however, sacrifices the benefit of fiber savings in P2MP architectures. TDM-PON will always constitute a shared-medium architecture, even with a P2P topology as the port on the OLT is shared. Thus, most of the well-known characteristics of shared media remain.

WDM-PON is an emerging architecture which promises to provide dedicated connectivity over a P2MP infrastructure by providing pairs of wavelengths per customer on the fiber tree. WDM PON on a P2P fiber infrastructure makes no practical sense.

From a civil engineering perspective the structure of the fiber plant for Point-to-Point fiber deployments are identical to that of Point-to-Multipoint. From the POP location feeder fibers are deployed towards some distribution point in the field either in some underground enclosure or a street cabinet. From this distribution point individual drop fibers are laid towards each individual household. The difference between both structures is only the number of feeder fibers. As fiber densities in the feeder and drop part are very different, often different cabling techniques are employed, depending on the specific circumstances.

Cabling techniques include classical cables and blown fiber in microducts. Classical cables can either be directly buried or accommodated in ducts for larger flexibility. In the feeder part deployments can be greatly facilitated not only by existing conventional ducts, but by other rights of way, like sewers, tunnels, or other available tubes. Thus, in most cases the higher number of feeder fibers in point-to-point topologies do not pose any major obstacle for these installations.

3 The Case for Point-to-Point Topologies

Access to the telephone network has always been provided via dedicated copper lines into every household. Initially, these lines were dimensioned for voice communication with frequencies up to 3.4kHz. Technological progress during the past 30 years has enabled a rapid increase of transmission rates over this classical medium. With VDSL2 the physical limits of this medium have now been reached with frequencies on the medium reaching up to 30MHz over short distances. This evolution would not have been possible if the passive infrastructure had contained technology specific components (in fact, in the US the loading coils in the copper plant have caused challenges to the DSL introduction in many places).

The structure of the access network has also allowed to easily unbundle the local loop, and to create competition among access network providers. This competition has largely contributed to increased access bitrates at lower and lower prices.

For the deployment of new fiber-based access networks for consumers (Fiber to the Home – FTTH) during the past few years there has been an intense discussion whether they should be built in the same way as the copper-based networks, in a Point-to-Point topology, or whether the high capacity of fibers would rather imply a Point-to-Multipoint structure, which allows savings of fibers and optical interfaces. Nonetheless, they could provide relatively high bitrates for every customer. For this purpose optical splitters are introduced in the passive infrastructure, which restrict its use to a class of technologies – Passive Optical Networks (PON). These technologies were developed initially at a time when fibers and optical interfaces were still very expensive. These arguments have become largely irrelevant through mass production of fibers, cables with very high fiber counts (e.g., 720 fibers in a cable with 16mm diameter), and very cost effective optical interfaces.

The economic lifetime of a passive infrastructure will be more than 40 years. During this period novel applications will be developed, enabled by the high transmission bitrates, and we will see innovative fiber-based transmission technologies which will allow higher and higher speeds at lower and lower cost. Technology specific components in the passive infrastructure will restrict the degrees of freedom for these developments.

Dedicated fibers allow to provide individual, virtually unlimited speeds for each customer. Changing technologies for one customer does not cause technology changes for the other subscribers as is the case with P2MP architectures.

The P2P topology allows access to individual fibers to every household allowing potentially the same kind of physical unbundling as in copper networks. National regulators are split about the question whether a physical unbundling of the access network will be required to enable competing service providers access to the passive infrastructure. The joint usage of such an infrastructure, however, can be of high commercial interest and allow investment sharing into a common physical network by multiple providers. For example, in Switzerland bilateral agreements between the incumbent operator and individual utilities have been signed about co-investments into passive infrastructures in a P2P topology.

Furthermore, the European Commission has set the criteria for state aid in broadband networks within their Community Guidelines for the application of State aid rules in relation to rapid deployment of broadband networks:

"… whatever the type of the NGA network architecture that will benefit from State aid, it should support effective and full unbundling and satisfy all different types of network access that operators may seek (including but not limited to access to ducts, fibre and bitstream). …"

In the POP the fibers arriving from the outside plant are terminated on an Optical Distribution Frame (ODF) as the fiber management solution which allows to flexibly connect any customer to any port on switches or splitters in the POP. Due to the large number of fibers to be handled in a POP the density of such a fiber management solution has to be very high in order to minimize real estate requirements. Fig. 3 shows a high-density ODF that allows to handle more than 2300 fibers in a single rack. For illustration purposes it is positioned next to a rack with active equipment.

Take rates in FTTH projects typically take some time to ramp up and usually stay well below 100%. The fiber management allows a ramp up of the number of active ports in sync with the activation of customers. This minimizes the number of unused active network elements in the POP and enables a slow ramp up of the investments.

Fig. 3. High-density fiber management solution (source: Huber&Suhner)

The emergence of 4G radio access networks is another driver for fiber deployment in the access network.

FTTH deployment will create major synergies with the rollout of these 4G networks because subscriber densities and base station densities in an area are positively correlated. This means that in urban areas where FTTH is deployed with priority also sufficient fiber can be made available for the large number of base stations.

As base stations have different connectivity requirements from residential customers – primarily clock synchronization – their connectivity should not be merged with that for residentials in a P2MP infrastructure.

4 The Case for Ethernet Technology

4.1 Transmission Technologies

After many years of technological debate Ethernet has emerged as the most prevalent transmission technology for virtually every application, from home networking to backbone networks. This is due to its conceptual simplicity, its support of a variety of different media – coaxial cable, shielded / unshielded twisted pair, single / dual multi-mode / single-mode fiber – and its scalability. There are Ethernet versions available, starting at 2Mbit/s on a shared coaxial cable up to 100Gbit/s on a fiber pair. Backwards compatibility has been largely maintained, leading to the ubiquitous 10/100/1000Mbit/s auto-sensing interfaces in homes and offices.

This ubiquity has led to the emergence of very cost-effective components for transmission and switching of Ethernet frames, and to a huge market for Ethernet equipment.

Recognizing the need for Ethernet in Service Provider access networks IEEE had established the IEEE 802.3ah working group already in 2001, creating a standard for "Ethernet in the First Mile (EFM)". Besides standards for OAM, Ethernet over copper and EPON, two standards for Fast Ethernet and Gigabit Ethernet over Single-Mode Single-Fiber were created.

The EFM standard was approved and published in 2004, and was included into the base IEEE 802.3 standard in 2005.

The specifications for the transmission over Single-Mode Single-Fiber are called 100BASE-BX10 for Fast Ethernet and 1000BASE-BX10 for Gigabit Ethernet. Both specifications are defined for a nominal maximum reach of 10km.

For the separation of the directions on the same fiber wavelength-division duplexing is employed, such that for each of the bitrate classes two specifications for transceivers are defined, one for "Upstream", i.e. from the CPE towards the POP and one for "Downstream", i.e., from the POP towards the CPE.

The following table provides the fundamental optical parameters of these specifications.

Table 1. Ethernet transmission technologies for FTTH

Description	100BASE-BX10-D	100BASE-BX10-U	1000BASE-BX10-D	1000BASE-BX10-U
Transmit direction	Downstream	Upstream	Downstream	Upstream
Nominal transmit wavelength	1550nm	1310nm	1490nm	1310nm
Minimum range	0.5m to 10km			
Maximum channel insertion loss	5.5dB	6.0dB	5.5dB	6.0dB

In order to cope with requirements not considered in the standard the market also offers optical transceivers with non-standard characteristics.

Some types can bridge significantly longer distances, e.g., for deployments in rural areas.

As the nominal transmit wavelength of 100BASE-BX-D (1550nm) is the same as the standard wavelength for video overlays in PON systems, transceivers exist which can transmit at 1490nm which allows to insert an additional signal at 1550nm carrying an RF-modulated video overlay signal on the same fiber using off-the-shelf video transmission equipment (see Fig. 5).

4.2 Operational Considerations

Traffic Management

For obvious reasons all communication networks are oversubscribed, starting at the uplink interface of the access network element. In shared media architectures oversubscription already occurs in the access network based on appropriate Media Access Control (MAC) protocols. Management of oversubscription in packet networks requires queueing and priority mechanisms. As a general rule, such queueing and priority implementations in switches and routers can be made significantly more sophisticated and effective than those based on MAC protocols. Using a dedicated port per fiber, therefore, removes the need for a MAC protocol in the access network, and it removes also an additional point of congestion to be managed within the entire context of traffic management.

Access network elements follow fast innovation life cycles and are replaced or upgraded in regular intervals which are much shorter than the lifetime of the fiber. Therefore, any bottlenecks in the active infrastructure can be removed on a regular basis, but this should not have any impact on the passive access infrastructure.

Security

A dedicated port on a dedicated fiber is inherently secure because the information transmitted over it can only be received at either end of the fiber, and it will not be shared, by default, with any other subscriber in the same access network domain. This inherent security, therefore, obviates the need for encryption on the fiber which is a necessary function in any shared medium architecture.

A failed ONT / CPE device connected to a dedicated fiber cannot impact the traffic of any other subscriber because any unplanned behaviour can be detected by the associated switch port in the POP, and the port can be deactivated until the defect is repaired. This property also eliminates denial of service attacks which can be launched by malicious users in shared medium architectures by deliberately jamming the signals.

CPE Deployments

For the deployment of the CPEs – either simple ONTs with integrated Ethernet switches, or more sophisticated routed Home Access Gateways – the service providers have the choice of two different scenarios.

- They own and install the CPE by themselves and also test the integrity of the transmission. In this case the subscriber has no need to touch the fiber in any way but only connects his home network to the subscriber-facing interfaces of the CPE.
- They can drop-ship the CPEs to the subscribers and have them connect their CPEs via optical patch cables to optical wall outlets. This requires more confidence in the subscribers' capabilities to handle optical fibers. Eventually, this will also allow CPEs to be distributed over retail channels.

While there are operational pros and cons for either solution, with Ethernet over dedicated fibers in no case is there any security risk involved in having the subscriber handle the fiber. Any potential problem is strictly confined to this particular subscriber's access line. One of the main questions in this context is rather whether the savings from self installation can compensate the higher cost of support calls.

Trouble-Shooting

Optical Time Domain Reflectometry (OTDR) mechanisms are used to determine any discontinuities or reflections in the fiber plant. OTDR transmits light pulses into the fiber, and the timing and intensity of the reflections indicates the location and the nature of any problem. In a point-to-point deployment the fiber is visible on its entire length from the POP to the subscriber, which greatly facilitates trouble-shooting of the fiber plant, compared to point-to-multipoint architectures which create ambiguous results for the reflections coming from the drop fibers.

Power Budget Planning

The Ethernet transmission technologies typically employed for point-to-point connectivity provide sufficient margin, even at the nominal maximum distances, such that a substantial number of splices and connectors can be accommodated without exceeding the maximum channel insertion loss. This facilitates the planning of the outside plant and allows for some aging of the passive components.

4.3 Support for Video Broadcast

IPTV-based video solutions provide superior features over simple broadcast solutions and have, therefore, become an indespensable part of any triple-play offering. Quite frequently, however, there is a need to provide RF video broadcast overlays in order to support existing TV sets in the subscribers' households. In splitter-based P2MP architectures this is typically accomplished by providing an RF video signal, compatible with cable TV solutions, over an additional wavelength at 1550nm. In Ethernet P2P installations this can be achieved by two different approaches, depending on the possibilities for fiber installation:

- In the first approach an additional fiber per customer is deployed in a tree structure (see Fig. 4) and carries only an RF video signal that can be fed into the inhouse coax distribution network. In this case the split factors (e.g. ≥ 128) exceed those typically used for PONs so that the number of additional feeder fibers is minimized.

- In the second approach a video signal is inserted into every point-to-point fiber at 1550nm. Fig. 5 shows how the RF video signal carried by a dedicated wavelength from a Video-OLT is first split into multiple identical streams by an optical splitter and then fed into each point-to-point fiber by means of diplexers. Mechanically, this solution is implemented in structures similar to Optical Distribution Frames. On the customer side the wavelengths are separated, the 1550nm signal converted into an RF signal for coax distribution, and the 1490nm signal made available on an Ethernet port.

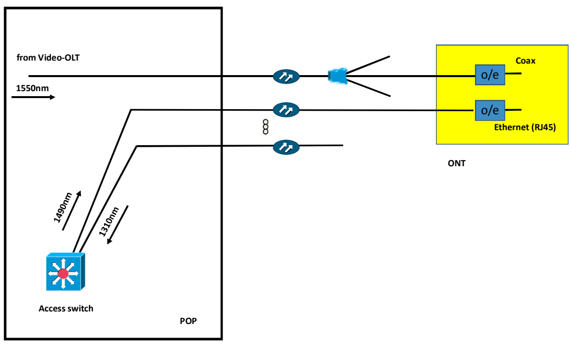

Fig. 4. RF video overlay using a second fiber per subscriber, deployed in a P2MP structure

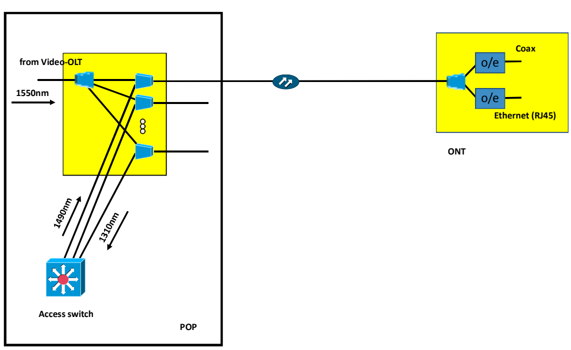

Fig. 5. Insertion of RF video signal into P2P fibers

In both cases the CPE/ONT devices comprise two distinct parts:

- a media converter which takes the RF signal on 1550nm and converts it into an electrical signal that drives a coax interface,
- and an optical Ethernet interface into an Ethernet switch or router.

In the single-fiber case the signals are separated by a diplexer built into the CPE, while in the dual fiber case there are individual optical interfaces for each fiber.

4.4 Support for Wireless Base Stations

To support the backhaul of traffic from wireless base stations tight clock synchronization is required. This can be achieved by special Ethernet interfaces with capabilities like Synchronous Ethernet (SyncE) or Precision Time Protocol (PTP) according to IEEE1588. As these capabilities incur cost which is not required to connect residential or business customers they are provided by dedicated linecards or systems.

5 Cost Considerations

5.1 Capital Cost

FTTH deployment involves a number of different cost components that can each be individually optimized. However, it is important to understand the relative contribution of each component and, thus, the relative saving potential. Fig. 6 shows a typical cost distribution for greenfield FTTH deployments.

This graph confirms what intuitively could be expected: civil works comprise almost 70% of the total initial capex. Obviously, this is the cost component where saving efforts have the largest effect. Therefore, usage of every potential right-of-way solution should be considered in order to reduce this component. As already mentioned earlier this comprises existing ducts, sewers, tunnels, …

With 6%, respectively 2% the fiber and other passive optical elements only contribute a very small part to the capex. Therefore, the saving potential from these components is very limited.

Active network elements are the second largest component with 12% contribution. Independent of the particular technology employed, this is a component where technological progress will continue to bring down per-port cost. In the case of Ethernet point-to-point architectures the access switches in the POPs are usually derived from systems that are deployed in very large quantities in Enterprise networks. Therefore, their cost benefit very strongly from manufacturing volumes which are significantly larger than they would be in the case of Service-Provider-only equipment. Furthermore, their technological evolution can be more broadly funded which leads to rapid innovation cycles and cost reduction.

It is also interesting to consider the distribution of cost for the infrastructure deployment. Swisscom have determined that 85% of the cost is in the drop segment and the inhouse segment of the network. Only 15% account for the feeder and the central office segment. As the former two segments are the same for each topology the saving potential from a P2MP topology, therefore, is very limited.

Based on detailed business case analyses with European Service Providers some typical values for the cost differences between point-to-point and point-to-multipoint deployments have been derived. In those cases where just sufficient duct space in the feeder plant was available to deploy the smaller cables for point-to-multipoint architectures, but not enough for point-to-point, the cost premium for Ethernet point-to-point could run up to 25% because of the need for additional civil works.

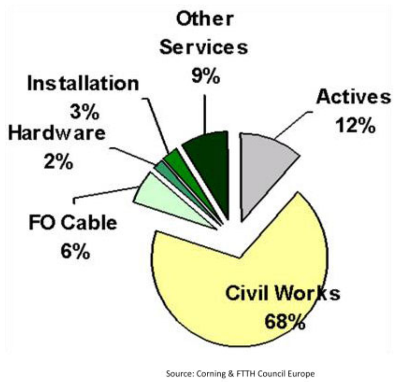

Fig. 6. Typical initial capex distribution for green-field FTTH deployments

Experience with real deployments, however, implies that this situation occurs far less frequently than situations where there is either no duct shortage, or where civil works have to be carried out because either all the ducts are occupied by other cables, or in greenfield scenarios where no infrastructure is available in the first place. In those cases the cost premium for point-to-point deployments usually stays below 5%. This difference can be mainly attributed to the more extensive fiber management in the POPs.

It can be expected that these initial small project cost differences will be overcompensated, over the lifetime of the fiber plant, by the inherent benefits of point-to-point deployments (flexibility, upgradability, fiber plant maintenance, …).

5.2 Operating Cost

Operating cost are a multi-faceted subject. Most of these cost items are not specific to any particular access technology, like marketing, subscriber acquisition, subscriber management, …

It is very difficult, though, to quantify operational cost differences. Therefore, we try to identify qualitatively the technology and architecture dependent aspects that can impact the operating cost of an FTTH access network.

Aspects which are favourable for Ethernet point-to-point deployments are mainly due to the relative simplicity of the architecture, as there are

- Ease of traffic management
- No encryption
- Ease of troubleshooting the physical layer
- Easy upgrade to higher speeds / new technologies on a per-customer basis.

Certainly, there are also some operational disadvantages of Ethernet point-to-point deployments:

- More real estate in the POP location due to the more extensive fiber management
- Slightly higher space requirements for the active equipment, although for typical penetration rates the difference is rather small as only active subscribers require

switch ports, in contrast to point-to-multipoint architectures where the first subscriber on a tree requires the allocation of an OLT port
- Higher power consumption in the POP as every active subscriber is connected to a dedicated port, although also this aspect is mitigated by the penetration rate, and technological advances will continue to reduce power consumption.

6 Summary

The need for Fiber to the Home is largely undisputed. The discussion is focused now on the timing, the business models, and the right architectures.

Synergies can be created with the backhaul connectivity of wireless base stations.

A variety of architectures for FTTH are available, each with its particular pros and cons. Access fiber deployment will form the basis of a new infrastructure which holds the promise to last for the next 40…50 years. This longevity should lead to deployment models which are technology agnostic and allow the maximum flexibility for the introduction of novel technologies and the support of new business models.

Point-to-Point architectures meet all these requirements at a very small initial capex premium.

On the technology side Ethernet is the most mature approach with its scalability, security and ubiquity. Recent developments provide also support for clock synchronization required by wireless base stations.