1

Introduction to Broadband Access Networks and Technologies

1.1 Introduction

In the mid-1990s, there were many doubts about the future of broadband access. No one was sure if the mass market needed or wanted more than 100 kbit/s; what applications would drive that need; what broadband access would cost to deploy and operate; what customers were willing to pay; whether the technology could provide reliable service in the real world; or which access technology would "win." Government regulation in many countries made it unclear if investment in broadband would yield profits. It seemed that broadband access would be available only to wealthy businesses. Fortunately, there were some people who had a vision of a broadband world and who also had the faith to carry on despite the doubts.

We now live in a world where broadband access is the norm and households without it are the exception. No one asks today why the average household would need broadband access. The answer is obvious: we need internet access, with its ever-growing number of applications, and VOD (video on demand). With more than 600 million customers connected to broadband networks, no one asks if the technology works or whether it can meet the customer's willingness to pay.

Furthermore, a growing application of broadband access is the support of femtocells, and small cells in general. Resorting to small cells has today become the most promising trend pursued for increasing wireless spectral efficiency, and the key to its success is the availability of a high capacity wired line to the home. Also, a growing fraction of cellular data is today generated indoors. In addition, it has become clear that no single broadband access technology will win the entire market, and that the market shares of the different technologies will change over time.

Each access technology has its strengths and weaknesses. A common constraint is that we can have it fast, low cost, and everywhere – but not all at the same time. In many cases, the choice of broadband access technology is driven by the legacy network infrastructure of the network provider. In other cases, national regulatory considerations are a significant factor. As a result, each access technology has its areas of dominance in terms of geography, applications, and political domains.

The book is divided into three sections:

- The chapters in the first section of the book cover technologies and standard protocols for broadband access over fiber-based access networks.

Broadband Access: Wireline and Wireless – Alternatives for Internet Services, First Edition.
Steven Gorshe, Arvind Raghavan, Thomas Starr, and Stefano Galli.

- The chapters in the second section cover technologies and standards associated with non-fiber, non-wireless broadband access.
- The chapters in the final section of the book address wireless broadband access technology and standards. Some of these technologies have been widely deployed, while others are anticipated to see deployment soon.

1.2 A Brief History of the Access Network

The traditional access network consisted of point-to-point wireline connections between telephone subscribers and an electronic multiplexing or switching system. The early access network used a dedicated pair of wires (referred to as a copper line or "loop") between the subscriber and the central office (CO) switch.[1] As the cost of multiplexing technology decreased, it became more economical in many cases to connect subscribers to a remotely located terminal. This remote terminal (RT) would multiplex calls from multiple subscribers onto a smaller number of wires for the connection to the CO. Network cost was reduced by having far fewer pairs of wires from the CO to the remote areas. As the technology evolved from analog frequency domain multiplexing (FDM) to digital time domain multiplexing (TDM), the RT systems became known as digital loop carrier (DLC) systems.

Data access to the telephone network began with the introduction of voiceband modems that could transmit the data as a modulated signal within the nominally 4 kHz voiceband pass-band frequency. The shorter lines (loops) allowed by DLCs made increasingly efficient modulation technologies practical. However, as explained in Appendix 1.A, the maximum data capacity of voiceband modems was limited to 33.6 kbit/s, or 56 kbit/s under special circumstances. Modems and their evolution are also discussed briefly in Appendix 1.A.

As a result, out-of-band technologies were introduced that transmitted signals over the copper line at frequencies outside the voiceband. Since these technologies sent digital information in the out-of-band signals, they became known collectively as digital subscriber line (DSL) technology. DSL is discussed further in Section 1.3 and Chapters 7–10.

Since the subscriber lines are implemented with twisted wire pairs, with multiple lines sharing the same cable without being shielded from each other, there are limits on the bandwidth that is achievable with DSL. For this reason, network providers became interested in alternatives to the subscriber line for providing broadband access. The three main contending technologies are coaxial cable, fiber optic cable, and wireless radio frequency connections. Each of these technologies is reviewed in later chapters of this book.

Coaxial cable networks were deployed by community access cable television (CATV) companies to provide broadcast video distribution. Due to the high bandwidth capabilities of coaxial cables, they had the potential for offering broadband services to their subscribers. In order to offer broadband data services, CATV companies evolved their networks to support upstream data transmission, and introduced fiber optic cables for higher performance in the feeder portion of their networks. As discussed below and in Chapter 11, coaxial networks have their own challenges as well as advantages.

Telephone network providers responded to the potential broadband advantages of the CATV companies by deploying additional fiber in their access networks. Telephone companies have deployed fiber directly to each subscriber's premises in some areas. Others are deploying fiber to terminals near enough to the subscribers' premises that broadband services can be provided by the latest very high-speed DSL technologies. The most attractive aspect to fiber is its virtually unlimited bandwidth capability. The primary drawback has been the relatively high cost of the network and its associated optical components.

Wireless access had not originally been a significant contending technology for residential broadband access. However, as wireless mobile networks have become widely deployed, and new technologies and

[1] Of course, some of the earliest access lines were "party lines" where several subscribers were connected in parallel to the same loop. Due the inherent lack of privacy and decreased cost of providing access lines, party-lines have become a historical footnote.

protocols have been developed, wireless broadband access has become increasingly important. It is especially attractive in regions that lacked a legacy wireline infrastructure capable of evolution to broadband services. Examples of such regions include developing nations and rural areas. It also offers the very significant advantage of allowing mobile, ubiquitous service rather than being restricted to service at the subscriber's premises.

Since a limited amount of spectrum is available for use in broadband services, the networks to support it have become increasingly complex. For spectrum efficiency, wireless networks use grids of antenna, where each subscriber only needs enough signal power to reach the nearest antenna. The region covered by each antenna is referred to as a cell. The result is that the same frequencies can be used by subscribers in non-adjacent cells, since their signals should not propagate far enough to interfere with each other. The signal formats have been optimized in the latest protocols to approach the Shannon limit for data bits transmitted per Hertz of transmission channel bandwidth. Capacity is further increased by re-use of the spectrum through smaller cells and smart antenna technologies. Both add cost, and radio signals are always more vulnerable to various types of interference than wireline technologies. Wireless technologies are discussed further in Section 1.6 below, and in detail in Chapters 14–17.

1.3 Digital Subscriber Lines (DSL)

1.3.1 DSL Technologies and Their Evolution

DSL operates over a copper line at frequencies outside the voiceband, sending digital data directly from the subscriber, and thus avoiding the need for an analog to digital conversion. Since the telephone lines were designed to provide good quality for voiceband signals, they are often not particularly well suited for higher rate data signals. Reflections become a significant problem in the electrical domain at rates beyond the voiceband. One of the worst sources of reflections in North American networks is bridge taps. When the feeder cables are installed from the CO into the loop area, they go to splice boxes where the wires going to the subscribers are connected. When service is disconnected to a subscriber (e.g., due to the homeowner moving), a second pair of wires may be connected to the feeder cable to serve a different subscriber without removing the other line. The result is a bridge tap, and it is possible to have bridge taps at more than one location along the connection to a subscriber. The unterminated end of the unused line(s) causes electrical reflections of the DSL signals, and these reflections can cause destructive interference for certain frequencies (any impedance mismatch along the copper connection to the CO to the subscriber can cause harmful reflections, but the bridge taps are especially bad).

The first widely deployed services using a digital subscriber line were the Digital Data Service (DDS) from AT&T. DDS used baseband signals over the line and offered data rates including 2400, 4800 and 9600 bits/s, and 56 kbit/s. The lower rate signals were sometimes converted to analog signals at the CO and then mapped into a voiceband channel, thus avoiding any noise or distortion from the subscriber line. DDS required the end-to-end service be synchronized to a common atomic clock. DDS circuits also usually required that the line be groomed to remove impairments such as bridge taps. While DDS circuits were very valuable for some customers (e.g., banks using them for connections to ATM machines), they were too expensive to deploy to residential subscribers or even to many business subscribers.

The first serious attempt to provide higher data rates to subscribers was the basic rate interface of the Integrated Services Digital Network (ISDN-BRI). ISDN-BRI used baseband signals[2] over the subscriber line to offer bidirectional data rates of 144 kbit/s. ISDN-BRI was designed to operate over most subscriber lines of up to 18 000 feet without having to remove impairments such as bridge taps from the lines. The 144 kbit/s signal was typically divided into two 64 kbit/s bearer (B) channels and a 16 kbit/s data (D) channel. The B channels could be used for voice or data, while the D channel carried the connection signaling information, with its leftover bandwidth available to carry subscriber data packets. It was also

[2] Specifically ISDN BRI used the 2B1Q line code, which mapped two input data bits to a quaternary symbol (i.e., a symbol that has four possible amplitude values).

possible to merge the two B channels or merge the Bs and D channel into a single 144 kbit/s channel. The cost of ISDN-BRI was relatively expensive, however, and there were no driving subscriber applications to generate high demand. ISDN also required that the connection signaling protocol be processed by the CO switch, which meant a major upgrade to the switches. By the time that Internet connectivity became a driving application, much higher rates were practical for DSL.[3] In effect, ISDN BRI provided too little bandwidth, too late, with too much network complexity.[4]

DSL modems that were dedicated to data services began to be widely deployed instead of ISDN-BRI. Initially, there were two broad categories of DSL. The first was a high speed DSL (HDSL) that provided bidirectional symmetric service at half the DS1 rate over a single pair,[5] or symmetric full DS1 rate over two pairs (half on each pair). Although it would seem that HDSL had no advantage over T1[6] service, which also uses two pairs, HDSL was capable of operating over much longer line lengths than T1, and it could do so without requiring repeaters. The total cost of HDSL was less than half of T1 lines, mainly due to eliminating most of the labor needed to install repeaters and remove bridged taps. It became common for carriers to use HDSL as the primary technology for providing DS1 connections to business customers. The current generation of HDSL is HDSL2, which allows bidirectional symmetric transmission of up to 2.048 Mbit/s payloads over a single wire pair.

The second category is the DSL lines optimized for residential subscriber access. The first generation was called ADSL (asymmetric DSL) due to its use of asymmetric data rates in the upstream and downstream directions. Since residential subscriber are typically downloading more information than they are providing to the network, they typically require much higher data rates from the network (downstream) than they do for upstream. This asymmetry in the desired data rates per direction was exploited to achieve the higher downstream rates. The service rate for ADSL is affected by several factors, but line length is the primary one. Over the past 25 years, telephone companies have tried to limit the line lengths to 12 000 feet.[7] Rates of 768 kbit/s downstream with 384 kbit/s upstream are possible over most of these lines. The actual rate is often determined adaptively as the system uses feedback to determine the frequency response of the line. In addition to higher data rates, another advantage of these DSL systems over ISDN-BRI was that they left the voiceband frequencies available for voice signals. This allowed analog POTS (Plain Old Telephone Service) signals to “ride underneath” the DSL data in its native format, which kept the voice and data signals separate within the network and allowed subscribers to use their existing telephone sets without conversion to digital signals at the subscriber premises.

ADSL rates and signal formats have been standardized by the ADSL Forum (now the Broadband Forum), by T1E1 (now ATIS COAST-NAI) and by the ITU-T SG15. SG15 is the primary body developing the current generation of DSL standards. The latest generation of ADSL is specified in the ITU-T G.992.5 standard for ADSL2plus which enables up to 20 Mb/s, with 12 Mb/s possible at 3000 feet.

Video delivery will require rates of 10–50 Mbit/s, depending on the service. For these rates, very high-speed DSL (VDSL) is required. VDSL requires lines lengths limited to 5000 feet. ITU-T SG15 has developed the VDSL2 standard, whose specifications are provided in ITU-T G.993.2 which enables rates up to 100 Mb/s upstream and downstream with 25 Mb/s downstream possible at 3000 feet. The ITU-T is developing the G.fast standard which promises to achieve bit rates up to 1 Gb/s over short copper lines.

[3] While DDS and ISDN BRI are digital subscriber loop technologies, it is most common to use the term DSL to refer to their successors that operated at higher data rates.

[4] Note that another early application of digital subscriber loops was to provide a second voice line over the same loop. This application is known as “pair gain” since it provides multiple voice channels over the same pair. Some carriers used ISDN technology to provide the simple pair gain service.

[5] When discussing subscriber loops, the term “pair” means a twisted pair of wires used for differential signal transmission.

[6] The term “T1” is commonly misused as being equivalent to a DS1. Strictly speaking, DS1 refers to the 1.544 Mbit/s signal and frame format, while T1 refers to a specific AT&T carrier system that transmits DS1 signals over 4-wire repeatered copper pairs.

[7] In the Bell System, the 12 000 ft. range was known as the Carrier Serving Area (CSA). Independent companies like GTE, who served more rural areas, specified their loop limits at 18 000 feet. Beyond 18 000 feet, inductive load coils need to be added to the loops to compensate their frequency response. DSL technology typically cannot operate through these load coils.

Both ADSL2plus and VDSL2 support transmission of packet transport mode (PTM), asynchronous transport mode, and synchronous transport mode (STM). ITU-T G.997.1 specifies management parameters for ADSL2plus and VDSL2.

1.3.2 DSL System Technologies

The first generation of DSL equipment connected DSL modems at the subscriber premises to DSL access multiplexers (DSLAMs) located in the central office. DSLAMs were next deployed in remote locations that were often co-located with DLC RTs. If the DLC RT was served by a SONET fiber connection, the DSLAM traffic would be multiplexed onto the same SONET signal as the DLC voice traffic. One of the challenges of co-locating the DSLAM and RT is that the DSLAMs require much more power per line than DLC equipment. This leads to heat dissipation issues when they shared the same cabinet, which can restrict the number of DSL lines that can be served. The DSLAM is not connected to the RTs backup batteries, however, since there is no requirement to maintain DSL service during a power outage.[8]

DSL was developed at a time when Asynchronous Transfer Mode (ATM) appeared to the preferred multiplexing technology for next generation networks. ATM provided adaptation techniques to carry a wide variety of packet-oriented data and constant bit rate (CBR) traffic such as voice signals. Hence, ATM was a natural choice for the encapsulation technology over the DSL line and for the multiplexing technology within the DSLAM. ATM allowed some statistical multiplexing for more efficient bandwidth utilization on the trunk from the remote DSLAM to the CO, or within the network.

There are two drawbacks to ATM, however. The first is that it adds at least five bytes of overhead to each 53-byte cell, causing a roughly 10% bandwidth overhead penalty. The bandwidth penalty is sometimes referred to as the ATM "cell tax." The other drawback to ATM is that it typically uses a rather complex signaling protocol that is overkill for purposes such as carrying connections to the Internet. Since most of the data going over DSL systems uses the Internet Protocol (IP) for Layer 3, it makes sense to use lower layer protocols that are more efficient with IP packets. Consequently, the emerging generation of DSLAMs is IP-based and uses Ethernet for the Layer 2 protocol instead ATM. These are commonly referred to as IP-DSLAMs.

1.4 Hybrid Fiber-Coaxial Cable (HFC)

While the telephone companies have focused on DSL, the CATV companies have deployed a network that is optimized for broadband broadcast traffic. As the demand for internet connectivity increased and the regulations allowed competition for providing telephone service, CATV companies have upgraded their networks to allow upstream data transmission from their subscribers.[9]

As illustrated in Figure 1.1, the CATV network uses a shared coaxial copper cable medium to connect to its subscribers. The coaxial segments are connected to remote equipment that provide the conversions to/from a fiber connection with a head-end office. These networks are called hybrid fiber coax (HFC). The bandwidth of the shared coaxial cable is divided into frequency bands, with one or more frequency bands being allocated for upstream transmission. Individual subscribers compete for the shared upstream bandwidth through a medium access control (MAC) protocol.

The most popular protocol for providing voice and data access is the Data Over Cable Service Interface Specification (DOCSIS™) protocol developed by CableLabs, a laboratory that is jointly funded by multiple cable network providers. The downstream data is modulated into the RF channel slots that would otherwise have been used for carrying video signals. The upstream data is similarly transmitted using RF

[8] The assumption is that if the power is out at the DSLAM location, the power is also out at the subscriber premises, and hence there is no subscriber equipment operational to use the DSL line.

[9] An early challenge for most CATV systems was that their signal repeaters only worked in the downstream direction, and required upgrades to support upstream traffic. The deployment of fiber reduced the number of repeaters requiring upgrade, since the fiber systems were designed to support bidirectional traffic.

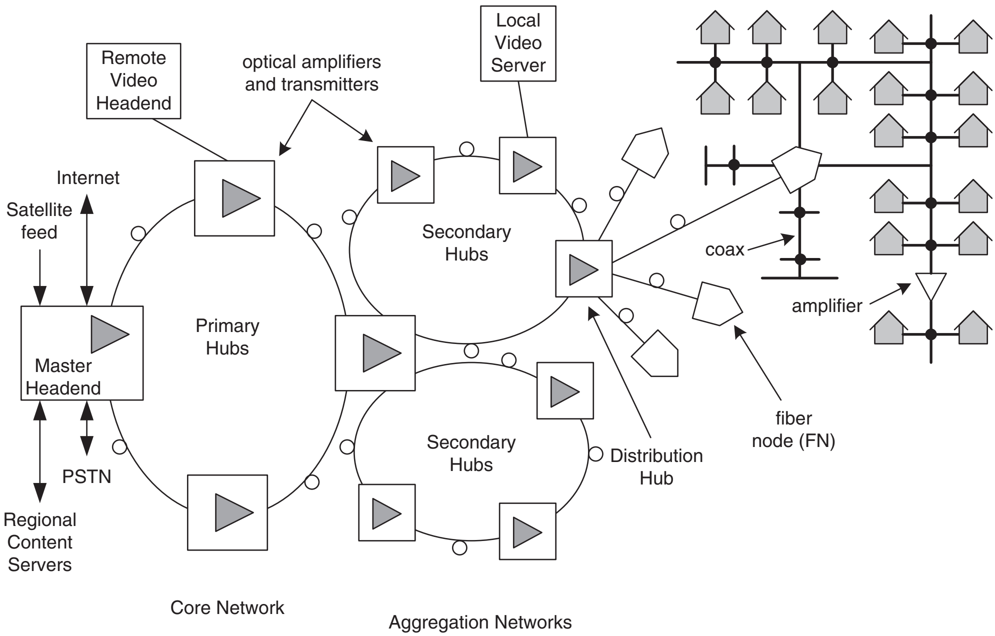

Figure 1.1 CATV network illustration

modulation into dedicated upstream frequency slots. The DOCSIS protocol assigns the frequency bands that are used by the cable modems at each customer's premises, and uses a shared-medium MAC protocol to determine the time slots in which cable modems can transmit their upstream data. The DOCSIS protocol also addresses the security issues associated with having a transmission shared among multiple subscribers where each can see the others' signals. The DOCSIS protocol is described in detail in Chapter 11.

Having been optimized for delivering video, the CATV networks are much better suited for delivering broadband video, including high-definition TV (HDTV) than are the telephone networks. A coaxial segment has typically been shared among several hundred subscribers. This degree of sharing inherently limits a CATV network's ability to provide high per-subscriber upstream bandwidth, and it also limits the number of video-on-demand (VOD) channels that can be provided. Reducing the coax sharing obviously increases their flexibility but also increases the cost of the CATV network. This is the primary tradeoff faced by the CATV providers in offering broadband access. To win in the market place, the cost and performance of the HFC networks must compete with the telephone company DSL networks and fiber to the home/curb networks (FTTH/C). One advantage HFC has over FTTH systems is that the copper coaxial cable allows a CATV company to provide power to the home telephone in the same manner that the telephone company does today.

1.5 Power Line Communications (PLC)

The use of power lines as a communication medium has been around for at least 100 years. This technology is generally referred to as Power Line Communications (PLC) and has sometimes enjoyed some degree of success over the years depending on the application it was used for.

The attractive feature of PLC is the high penetration of electrical infrastructure in the world, which in many areas is much higher than any other telecommunication infrastructure. As access to the internet today is becoming as indispensable as access to electrical power, and since devices that access the internet are normally plugged into an electrical outlet, the unification of these two networks always appeared to be a compelling option, despite the various technical challenges. As virtually every line-powered device can become the target of value-added services, PLC may be considered as the technological enabler of a variety of future applications that would probably not be available otherwise.

Among the various applications, today's interest for PLC spans several important applications: broadband internet access; indoor wired LAN for residential and business premises; in-vehicle data communications; smart grid applications (advanced metering and control, peak shaving, mains monitoring); and also municipal applications, such as traffic lights and lighting control and security.

In particular, smart grid applications have been and continue to be today a successful and promising area for PLC. Similarly, the interest in using PLC for home networking is increasing rapidly and, despite today's low penetration, many believe that home networking will be one of the most important areas of success for this technology. On the other hand, the great interest in the late 1990s for using PLC for providing broadband access to households has encountered many disappointments over the last two decades. Higher than anticipated costs in deploying PLC, growing EMC (Electro-Magnetic Compatibility) issues for the interference caused to radio services in the HF bands, its smaller capacity compared to DSL and cable, and the availability of other (and often cheaper) means to provide broadband access to consumers have made the initial enthusiasm in PLC for broadband access greatly diminish if not vanish.

There are very few PLC deployments in the world for broadband access and its use in industrialized countries, where the availability of other broadband access technology is abundant and cheaper and has made PLC a marginal technology. Perhaps the area where broadband access via PLC may still have some possibility of success is in third world countries, where access to the internet is essential to economic growth but there is no or very little telecom infrastructure. Similarly, rural areas in industrialized countries where it is very uneconomical to provide broadband services at competitive prices could also benefit from the deployment of PLC as most of these areas lack traditional telecom infrastructure but nevertheless have access to power.

Despite its failure to become a successful technology for broadband access, PLC will be addressed in Chapter 13. Because of its widespread use as a Smart Grid technology, the use of PLC in the power grid will also be addressed and its unique benefits for this application will be highlighted.

1.6 Fiber in the Loop (FITL)

Telephone company revenue from plain old telephone service (POTS) is declining as the result of losing some of their POTS customers to mobile phones and CATV companies. In order to increase their future revenue potential, the telephone companies believe that they need to be able to offer the best "triple-play" services, consisting of telephone, data (especially internet access), and video. At one time, telephone companies considered the idea of deploying the same type of HFC networks used by CATV providers. One drawback to this approach is that the coax networks typically have inferior reliability for voice service. The other main drawback is that they would only be "me-too" for video and data, thus providing no advantage over the CATV companies.

The model now preferred by telephone companies is based on their traditional approach of either avoiding fully shared media or limiting the amount of sharing. For non-shared subscriber medium, triple-play services typically will be provided through a fiber to the node (FTTN) architecture. With FTTN, as shown in Figure 1.2, a high-speed fiber connection[10] exists between the CO and a remote node that is close enough to the subscriber to allow individual VDSL connectivity to each subscriber served by that node. Variations on FTTN are Fiber to the Curb/Cabinet/Building (FTTC/FTTCab/FTTB). FTTN is very attractive when many subscribers are close enough together to be reached by VDSL (e.g., dense housing

[10] The high-speed fiber connection can either be a SONET/SDH link or a Gbit Ethernet link.

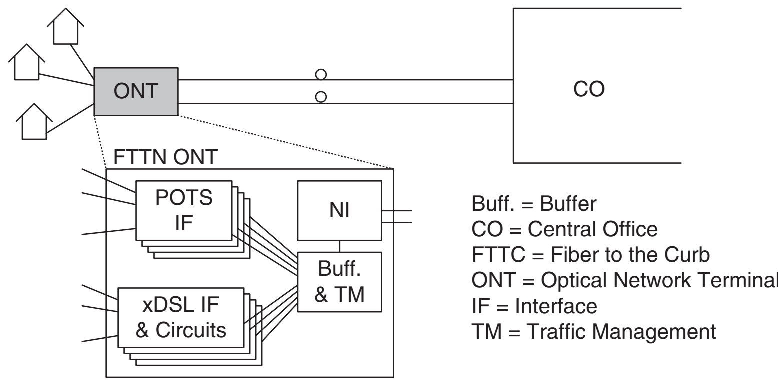

Figure 1.2 FTTN network illustration

neighborhoods or multi-tenant buildings). When subscribers are spaced further apart or require very high upstream bandwidth, fiber to the premises/home (FTTP/FTTH) becomes more attractive.

The passive optic network (PON) is the most attractive technology for FTTH/FTTP. PON systems share the fiber medium among a limited number of subscribers. Due to the directional nature of fiber optic transmission, only the downstream signals are visible to all subscribers on that PON. This simplifies the encryption processes required to ensure privacy relative to those required for shared coaxial cable or wireless networks.

Due to the relatively high cost of optical components (especially lasers and optical receivers), it is not cost effective to give each subscriber a separate fiber connection to the CO. The best way to reduce the number of optical components, as well as reducing the amount of fiber, is to have multiple subscribers share the same passive fiber network for their connection to the optical line terminal (OLT)[11] in the CO. The PON is illustrated in Figure 1.3. The terminal at the subscriber premises is typically called an optical network unit (ONU) or optical network terminal (ONT). Different generations of PON technology allow different numbers of ONUs to be connected to an individual PON, but 16 and 32 are typical numbers, with some systems connecting up to 64 and future systems being capable of higher numbers. Since passive optical splitters are used to divide (and merge) the optical signal among the ONTs, the number of ONTs connected to a fiber is often called the split ratio (e.g., 32-to-1).

PON systems typically transmit both upstream and downstream data over the same fiber. In some cases, only directional couplers are used to separate the upstream and downstream traffic, but higher speed systems typically use different wavelengths in each direction. The most common is course wave division multiplexing (CWDM), in which 1490 nm is used for the downstream direction and 1310 nm for the upstream. This wavelength assignment has the advantage of putting the less expensive 1310 nm lasers at the ONTs.

In the downstream direction, the OLT broadcasts the data for all ONUs. This downstream signal is comprised of the downstream data for all the ONTs and synchronization information for the upstream transmissions. The ONTs extract their downstream data based on either time slots or cell/packet address information.

In the upstream direction, the ONUs need a medium access control (MAC) protocol to share the PON. The most common MAC protocol is time domain multiple access (TDMA), which is similar to the protocols used by broadcast television satellites. With TDMA, the nodes are granted time slots in which to transmit their upstream data. In basic PON systems, each ONT is preassigned a fixed portion of the upstream bandwidth, and transmits its data at the appropriate time. In order to achieve greater efficiency,

[11] Another popular name for the OLT was host digital terminal (HDT). The OLT can either be located in the CO or at a remote (RT) site.

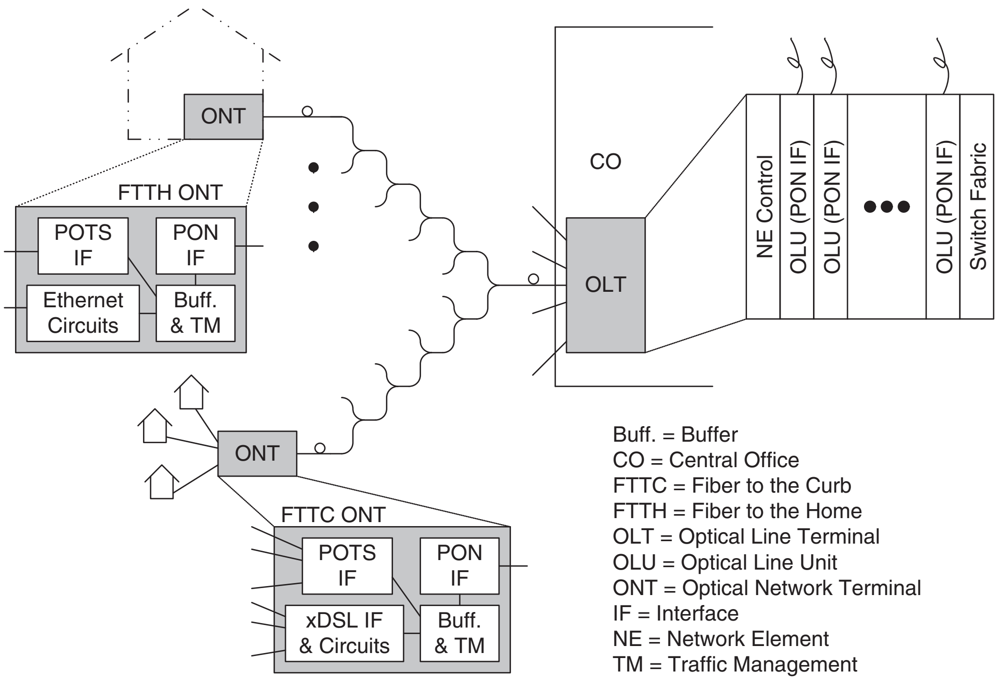

Figure 1.3 Illustration of a PON

PON systems now typically allow dynamic bandwidth allocation (DBA) among the ONTs. With DBA, each ONT uses part of its upstream transmission to inform the OLT of its bandwidth requirements.

For example, this information could be based its input queue fill level, including the levels for data in different classes of service. The OLT evaluates the requests from the ONTs, and assigns the bandwidth for the next upstream transmission frame. This bandwidth is typically communicated as a transmission start time and either a stop time or transmission duration time within the upstream frame. These bandwidth assignments are sent in the downstream transmission frame. The information used by the OLT in determining the appropriate bandwidth allocations can include the service level agreements (SLAs) associated with the ONT data flows. In some systems, the ONT is responsible for determining how to accommodate the relative priorities of its transmit data within the granted upstream transmission slot. The most popular TDMA PON protocols are described in detail in Chapters 3 and 4.

One alternative to TDMA is wavelength division multiple access (WDMA) in which each ONU has its own upstream and downstream wavelength for communication with the OLT. In other words, the separate wavelengths allow each ONU to have a point-to-point connection to the OLT over the shared PON fiber. The main drawback to WDM is that each ONU needs a unique wavelength, which would be very hard to administer if subscribers are allowed to buy their own ONUs. Tunable lasers would alleviate this problem, but they are currently too expensive. Other frequency selective technologies are being researched and developed for use at ONUs, but to date they have not been cost effective relative to TDMA technologies.

Another alternative is code division multiple access (CDMA). CDMA uses a spread spectrum approach where the subscriber bit stream modulates a code sequence, essentially in the same manner as is used for mobile phones. CDMA is very attractive since it can be implemented with entirely passive components at the transmitter and receiver. A further advantage of CDMA is that each subscriber can use a different native client interface. CDMA circuits, however, typically require optical amplifiers and precision receiver discriminator circuits to achieve the required signal to noise ratio. Other optical-domain medium access methods are also possible, but WDMA appears to be the most likely long-term approach. These optical domain technologies are described in Chapter 5.

There are also technology combinations that use a PON infrastructure in combination with a different technology, such as carrying the radio-frequency modulated CATV signal over a PON. These hybrid PON protocols and technologies are covered in Chapter 6.

1.7 Wireless Broadband Access

The mobile computing paradigm has seen phenomenal growth in the first decade of this century. Most services traditionally accessed on desktop PCs or dedicated networked hardware are being augmented with, or completely supplanted by, mobile access on tablets and smartphones. Mobile devices are clearly being powered by wireless technologies. However, before delving into the different wireless technology options, one must first establish a clear understanding of the role played by both wired and wireless technologies in delivering broadband access to the untethered end user.

Let us take a simplistic view of wireless access technologies initially, and divide wireless technologies into "long range" and "short range". Long-range wireless links (such as those used by cellular technologies) can serve users over a widely distributed geographical area, and can therefore be seen as a true alternative to the wired access options introduced in the previous sections. On the other hand, short-range wireless links only cover a small area such as a home or an office. Short-range wireless technologies therefore need to be augmented by wired backhaul access technologies in order to provide a complete solution for broadband access to the end user.

It is important to understand in this configuration that the "speed" of the broadband connection is actually determined by the smaller of the access rates of the wireless portion and the backhaul portion. To give a concrete example, a WiFi installation at a cafe may use the latest and fastest version of the standard, providing several hundred megabits of throughput. However, the cafe may use a DSL backhaul connection providing only a few tens of megabits of throughput, due to cost or availability considerations. In this example, the end user experience will be limited by the backhaul speed. The converse is also possible, where the wireless access speeds can limit the overall user experience, as we shall see.

Another possible distinction between wireless access technologies can be based on whether they provide "fixed" access or "mobile" access. In the early 2000s, several vendors developed systems based on a DOCSIS-like protocol for fixed access, where equipment would be installed at customers' premises and provide the long-range backhaul for Ethernet-based LANs. Early on, these systems were proprietary, but the need for a common standard soon became apparent. This led to the development of the IEEE 802.16 standard to provide long-range, fixed wireless access under the title "Wireless Metropolitan Area Networks" (Wireless-MAN).

However, fixed wireless systems, whether they were proprietary, or based on 802.16, were mainly restricted to smaller deployments in low-density population centers where the cost of installing wired access was seen as expensive for the corresponding revenue potential. Moreover, it also became apparent that a single technology, developed for both fixed and mobile access, would result in a more robust ecosystem with more applications and adoption potential. As a result, the 802.16 standard evolved to provide mobile access, but due to the deployment of competing cellular technologies, 802.16-based systems have not been able to gain any significant market share.

The two types of technologies that are most popular for providing broadband access today are the IEEE 802.11-based Wireless LAN (WLAN) standard, popularly known as WiFi, and the third and fourth generations of cellular technology. WLANs use unlicensed spectrum with restrictions on the transmit power and are, therefore, mostly used as a short-range technology. In addition, WiFi is expected to coexist with other systems in an unregulated environment, so it has been designed to be able to coexist and be robust in the presence of interference. Furthermore, the technology is designed to use a simple architecture that is easy to configure and install. The basic topology consists of an access point providing the broadband connectivity to several associated wireless clients that represent end users. It is this combination of the use of unlicensed spectrum, ease of installation and interference-robustness that led to the rapid adoption of this technology. It is ubiquitous today as the predominant access technology in low-mobility environments such as homes, offices, campuses, and other public spaces.

In contrast, cellular technologies took a very different evolution path to becoming an alternative for broadband access. While WiFi was designed from the start to provide access to data networking services, cellular technologies were initially designed with the sole goal of providing mobile voice service. The requirement to support seamless mobility for voice via handovers resulted in a more complex and expensive system. Furthermore, they were deployed in licensed spectrum to ensure that interference can be managed in a regulated fashion, thereby ensuring high reliability for voice services. As such, these systems are owned and operated by service providers, not by individuals or small enterprises. With transmit power not severely restricted, as in the case of unlicensed spectrum, cellular systems can cover much larger areas.

The basic topology of a cellular system is based on the concept of cells and frequency reuse. Early systems were not designed to be robust to interference, and therefore needed to use frequency separation and cells to manage interference, as shown in Figure 1.4. The figure shows a frequency reuse factor of three, because a separate carrier is needed in each cell in the repeating cluster of three cells in order to maintain a minimum interference separation. Other reuse factors are possible, with varying degrees of interference separation. Early cellular technologies, in the first and second generations required frequency reuse and supported mainly voice services. However, as the need for mobile data services grew with the growth in internet traffic, combination of scarce licensed spectrum resources, and the greater capacity needed for data services led to the design of reuse-one systems where all cells could use the same frequency, and the interference mitigation was carried out by a more robust physical layer designed to operate at lower signal-to-noise ratios.

In the wireless section of this book, comprising of Chapters 14–17, we take an in-depth look at the three most widely deployed technologies for mobile broadband access. Before delving into the details of the technologies, we first try to establish, in Chapter 14, the fundamental concepts that apply to all wireless systems. In addition, the various basic building blocks that are part of the air-interface of any broadband access technology are explained. Next, in Chapter 15, we discuss WiFi based on the IEEE 802.11 standard. Lastly, in Chapters 16 and 17, we discuss third and fourth generation cellular technologies. In Chapter 16, we focus on the technology based on Wideband Code Division Multiple Access (W-CDMA), and briefly discuss how it contrasts with the other third generation system based on CDMA-2000. In

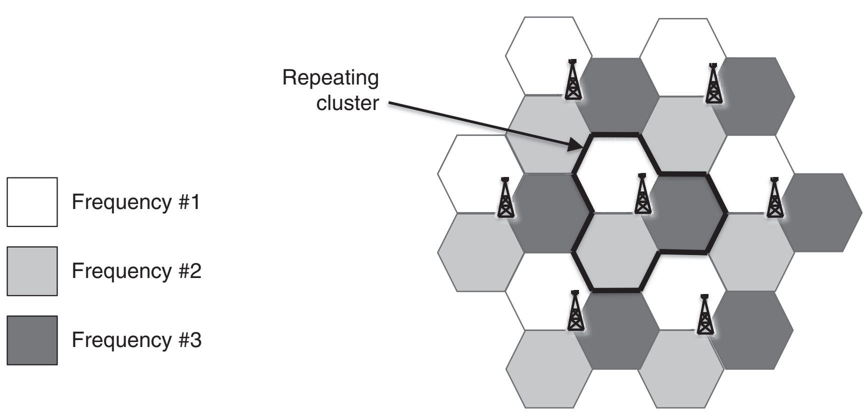

Figure 1.4 Cellular system with frequency-reuse factor = 3

Chapter 17, we discuss the fourth generation systems based on LTE and LTE-Advanced that are expected to be widely deployed, with a brief mention of WiMAX which is based on the IEEE 802.16.

All of these technologies are developed and implemented on the basis of specifications set forth by standards development organizations. As such the evolution of these technologies can be tracked by examining each new release of the specification in sequence. In discussing these technologies, we will first discuss the baseline features, network and protocol architecture of the technology. Next we will see how the technology evolved by taking a detailed look at each significant release of the standard and pointing out the key features and capabilities that were introduced in that release. Each chapter concludes with a summary that condenses all the material covered in the chapter into a few paragraphs to provide a quick review of the most noteworthy elements of the technology.

1.8 Direct Point-to-Point Connections

While direct point-to-point connections are not cost effective for residential subscribers, they will continue to be used for large corporate subscribers. Copper wireline connections can be DS1, E1, DS3, or Ethernet. Fiber connections include SONET/SDH, 1G, 10G, or 40G Ethernet, dark fiber, or a WDM wavelength. Wireless point-to-point connections are typically microwave radio links. The primary advantages of these direct connects are guaranteed bandwidth and security (since there is no shared medium).

While direct fiber connections are often not available to enterprise subscribers, DS1/E1 connection availability is ubiquitous. In North America, the regulatory environment can also create a price advantage for services providers to lease DS1/DS3 connections rather than fiber connections through the local exchange carrier networks. With the addition of virtual concatenation support for DS1/E1/DS3/E3 signals, copper connections through the traditional telecommunications infrastructure have become much more flexible. GFP then provides the transparent mapping for packet data services (see PMC-Sierra white paper PMC2041096). Previously, providing copper connections between the DS1 and DS3 rates required fractional DS3 or some relatively inflexible or inefficient method of combining DS1 s. These methods included inverse multiplexing with ATM (IMA), packet-specific techniques such as the IETF Multi-Link Point-to-Point Protocol (ML-PPP), or proprietary solutions.

1GE and 10GE fiber connections are becoming increasingly important as the UNI to enterprise subscribers. The telecommunications network provider may, in turn, use WDM for increased utilization, or map this data into its SONET or OTN infrastructure where TDM multiplexing allows even greater fiber utilization.

Appendix 1.A: Voiceband Modems

Voiceband modems began by using dual-tone frequency-shift key modulation for rates of 300 bit/s. As technology advanced, it became practical to us phase-shift key modulation and combinations of the amplitude modulation and phase modulation such as quadrature amplitude modulation (QAM) for greater efficiency. The capacity of any information channel is determined by the Shannon channel capacity theorem:

$$C = B \log_2(1 + S/N)$$

The capacity limits on the data rates for voiceband modems are primarily determined by the analog-to-digital conversion that takes place when the modem signal from the subscriber reaches the telephone network equipment (DLC or central office switch). Specifically, the 8 kHz sampling rate, and the quantization noise introduced when converting a voiceband signal to a 64 kbit/s digital signal determine the channel bandwidth (B) and the noise (N) terms of the Shannon capacity equation. The modem signal power (S) is limited by both the dynamic range of the analog-to-digital conversion, regulation, and the need to avoid crosstalk into other subscriber loops in the cable. The resulting capacity limit (C) for a voiceband modem is approximately 34 kbit/s, considering data transmission over voiceband channel with additive white Gaussian noise and assuming a nominal bandwidth of about 3.5 kHz and a signal-to-noise ratio of about 30 dB. Using efficient modulation techniques and error correction technologies, such as trellis coding, allowed standard voiceband modems to approach this limit with 33.6 kbit/s.

However, the value of 33.6 kbit/s was still far from the theoretically possible DS0 data rate of 64 kbit/s that could have been achieved with the same bandwidth but higher signal-to-noise ratio. The 64 kbit/s maximum value depended on the use of a 8ksample/s sampling rate and of 8 bits/sample in the analog-to-digital conversion.

In some circumstances, modems are indeed capable of approaching the theoretical maximum of 64 kbit/s if certain conditions of low quantization noise are met, for example, when a subscriber is connected via an analog line to a switched digital network and thus only one analog-to-digital conversion takes place. In some cases, the source of the data sent to a subscriber has a digital connection to the network (e.g., a DS1/T1 link) rather than a modem connection. Examples of such data sources include internet service providers. The digital-to-analog converter connecting to the subscriber loop creates a downstream signal that has none of the quantization noise that would have been created by an analog-to-digital conversion. If the other noise sources affecting that subscriber loop are small enough, then the channel capacity of the loop can be approached. In these circumstances, and sill using a sampling rate of 8ksample/s but encoding with data only seven bits of the 8-bit word in the analog-to-digital conversion,[12] then modems can achieve 56 kbit/s downstream rates. Since the upstream signal from the subscriber must go through the telephone network equipment's analog-to-digital conversion, the upstream signal rate of these modems is still limited to the standard 28.8 or 36.6 kbit/s rates.

[12] To improve the probability of error, only 128 PCM values are used.