
EVOLUTIONS IN LAST-MILE BROADBAND ACCESS

Not surprisingly, copper phone networks that were originally designed for voice calls have proven to be a valuable way to shift other types of information to and from households. From the first residential modem, which used an audio coupler between a computer and telephone, to high-speed VDSL2 technology, the traditional copper local loop continues to be a viable option to deliver faster data rates, for both high-speed Internet access and richer multimedia services. This type of access supports incumbent operators with a vast copper distribution network as well as retail operators that want access to the copper network in a wholesale environment. The wholesale environment could be *naked DSL* access, which provides a copper interconnect between the retail provider and the incumbent for each household. The environment could also be bit stream access using L2TP or ATM/Ethernet interconnects. This chapter describes the range of DSL types in common use among carriers today, from the 8Mbps speeds of ADSL, to ADSL2+, to the 100Mbps speeds of VDSL2.

Although twisted copper pairs are the most ubiquitous way of providing data access, cable providers using hybrid fiber/coaxial (HFC) is the next most popular way of delivering residential broadband. These were briefly described in Chapter 1, “A History of Broadband Networks”; however, this chapter deals almost exclusively with ADSL.

One pervasive data delivery method is wireless, which can take many forms: WiFi, CDMA, GPRS, UMTS, and others. Because this is such a large topic, Chapter 9, “The Future of Wireless Broadband,” is dedicated to wireless broadband access.

ADSL ACCESS

Chapter 1 explains the history of ADSL, beginning with the early line Carrierless Amplitude Phase modulation (CAP) line coding deployed under the standardization auspices of the American National Standards Institute (ANSI). The progression was toward the Discrete Multitone (DMT) method deployed today under the International Telecommunication Union (ITU) G.992 and G.993 family of standards. This section discusses the most important ADSL standards to date, from the first—G.992.1—to the latest—ADSL2plus (also commonly called ADSL2+)—which is documented in G.992.5.

Generally speaking, ADSL is a way to transmit digital signals over one or more analog carriers, much in the same way that many other digital transmission standards work. The majority of ADSL-based services can coexist on the same physical cable as existing telephonic services, such as Plain Old Telephone Service (POTS) or ISDN. It achieves this using a technique called Frequency Division Multiplexing (FDM), which splits the two services into two frequency domains. POTS shares the lower part and ADSL, the upper part. In the ADSL domain, duplexing of the upstream and downstream parts occurs; the upstream is on the lower frequency range, and the downstream uses the upper frequency band.

ADSL stands for Asymmetric Digital Subscriber Line. Its name comes from the difference between upstream and downstream data rates. The downstream rate is always higher than the upstream speed. One reason for this is because one of the early motivations for ADSL was to transport streaming video to the home, which does not require a large amount of upstream bandwidth. A further application—residential Internet broadband—has historically utilized more traffic in the downstream direction. Also, the asymmetric nature of ADSL reduces Near-End Cross Talk (NEXT), which allows for greater density of services, longer loop lengths, and higher bit rates than would otherwise be achievable. This section

covers all members of the ADSL family. It also describes SHDSL, which is a symmetric DSL service that better suits the needs of business connections that need more upstream bandwidth.

Technology Note: Crosstalk

In telecommunications, crosstalk is where one circuit, commonly a local loop, causes undesired interference to an adjacent circuit. This could be in the form of inductance or capacitance. A common example of crosstalk is where another party's voice call can be heard in the background of your telephone call. This is different from a party line, which intentionally wired several households together on a single circuit.

Crosstalk also causes issues for data transmission over copper pairs, including POTS ones. Humans can tolerate voice interference on POTS lines better than machines can. Because machines transmit digital data over copper pairs by using analog approximations of digital signals, any RF interference from other copper pairs can corrupt the information. There are two types of crosstalk discussed in the context of DSL systems: Near-End Crosstalk (NEXT) and Far-End Crosstalk (FEXT). NEXT occurs when the transmitting signal from one wire interferes with the signal in an adjacent wire. This is the opposite of FEXT, which occurs at the opposite, or far-end, end of the circuit. FEXT has less impact than NEXT because the signals are weaker at the other end.

One of the hallmarks of ADSL services is the effect that distance has on line speed. ADSL, for example, can reach speeds of up to 8Mbps on loop lengths of up to 6,000 feet (1.8km), but above that, the upper speed starts decaying. This happens because resistance, which attenuates signal, increases proportionally to both cable length and frequency. So as the cable length increases, so does the resistance, which also attenuates (reduces) the power of the higher frequencies more than the lower frequencies. Because the downstream portion of the line occupies the upper frequency band, this is the first to be affected, reducing the upper limit of the line speed.

Figure 6.1 shows how FDM lets the ADSL line share the same cable pair as POTS or an ISDN service. One semantic note: Frequency Division Duplexing (FDD) is where two frequency ranges coexist, such as an ADSL service, which duplexes the upstream and downstream frequency onto a single cable. Multiplexing is where more than two ranges share the same cable.




Figure 6.1 How ADSL coexists with telephone services sharing the same line.

Note that the low-pass filter is between the line and the POTS device. The DSL modem is connected directly to the line, not from behind the filter, as shown in the figure. The DSL modem has a high-pass filter that prevents POTS signalling from interfering with the DSL link and prevents the DSL modem from emanating frequencies in the voice or guard bands. This is part of the output Power Spectral Density (PSD) shaping, which is described later.

G.DMT

Several features of G.DMT are of interest to service providers. First, native STM or ATM data paths can be transmitted in the ADSL signal. Also, an Embedded Operations Channel (EOC) is used for statistics and diagnostics on the line. An interesting feature of the EOC is autonomous data transfer mode, which provides a 15kbps data channel that can transfer information such as the Customer Premises Equipment (CPE) vendor name, software or hardware revision number, serial number, and other informational parameters to the DSLAM. This is a

useful way to track the location of CPE in the network based on the serial number transmitted in the EOC. Finally, a “dying gasp” is a message that can be sent from the CPE when a loss of power occurs at the CPE location. The CPE must have sufficient operating power long enough to be able to transmit this message. This gives information to the DSLAM. If there were a power fault instead of a line break, a dying gasp message would not be sent. The specification additionally allows ISDN and POTS frequencies to coexist with the ADSL signal.

Standards and Spectrum

The first ITU DMT-based standard to precede ANSI CAP was G.992.1, released in mid-1999. DMT-based ADSL allows simultaneous Plain Old Telephone Service (POTS) (including AC-based telephone ringing signals [25 to 50Hz]) and voice frequencies (300 to 3400Hz), and ADSL to share a single copper line. Some documents refer to the voice range as 0 to 4kHz, including the ITU Technical Recommendation itself. The use of the frequency range between 3400 and 4000Hz is to cover any additional spectrum that non-POTS services might need. For example, G.992.1 notes that the V.90 modulation standard might need frequencies between 3400Hz and 4000Hz as a guard band. In the standard, anything between the frequency ranges of approximately 25kHz to 1.1MHz is considered part of the ADSL spectrum, but the exact usage varies according to hardware manufacturer and local deployment requirements. Between 25kHz and 138kHz is allocated to the upstream direction, and 138kHz to 1.1MHz is in the downstream direction.

Data Encoding

The standard also explains how to encode data bits at the provider’s ADSL modem, which is called an ADSL Transceiver Unit-CO (ATU-C), and at the remote CPE end, the ADSL Transceiver Unit-Remote (ATU-R). The encoding uses a constellation scheme. In the context of ADSL, this allows a chunk of bits (also called a symbol) to be represented by changes in frequency of the analog samples on the wire, as one example. Other ways to increase the line’s information rate are to use phase shifts, amplitude shifts, or a combination of the two. To improve throughput, the use of Trellis Code Modulation (TCM) is strongly advised.

Technology Note: Shannon-Hartley Theorem

In reference to a scientific innovation, Isaac Newton once said, “. . . If I have seen a little further it is by standing on the shoulders of Giants.” Many advancements in science have not been radical breakthroughs, but additive enhancements. The Shannon-Hartley theorem is built on two important information theories: the Nyquist rate and Hartley’s law.

Until the Shannon-Hartley theorem, a channel’s capacity (a piece of copper) was modeled on Hartley’s law, which was a way to calculate the rate at which information could be sent down a medium, given its bandwidth, suitability of transport for a given signal, and the reliability of re-creating the information at the other end. Of course, this acknowledged that noise is a limiting factor in a line’s information capacity, but not how much of a factor it is. Shannon came up with a model that approximates the noise that is always present, in the form of white, or Gaussian, noise. Shannon incorporated this white noise into this theorem and achieved channel reliability through error-correction coding (see the section “Error Correction and Detection”).

The theory is that a medium has a maximum capacity based on its signal and noise levels, the channel’s bandwidth, and its frequency. The rate of information transfer has an upper bound, and no more information than this theoretical maximum can be transmitted.

Shannon’s law is still relevant today and can be found in many spheres of telecommunications influence, including IETF mailing lists and mailing list discussions.

Technology Note: Trellis Coding

Trellis Coding (TC) is an efficient way of to use a data transmission medium. Most contemporary data transmission methods involve taking a series of bits and preparing them in some way before being transmitted.

A simple example of transmitting data over a copper wire could mean when a current is at a certain positive voltage, it represents a 1 and the same voltage in the negative could be zero. A more complex scheme, such as QAM, uses sinusoidal waves to represent the bits. QAM makes use of two carrier waves that are combined, which allow both amplitude and phase changes to be carried as information to the receiver. The decoder can take these amplitude and phase changes and work out the encoded bit stream. Often these combinations are represented in a constellation diagram. Figure 6.2 shows an example of a 16-QAM diagram, which uses discrete points to represent the various combinations of bits that can be sent with particular phase and amplitude changes. Behind these diagrams are relatively complex mathematics, especially when dealing with more than two dimensions (as shown in the figure).

In simple terms, TC modulation is a way to improve a given modulation scheme by adding an extra constellation points (for reference, there are 16 points in the figure) to improve the signal to noise ratio of a channel.




Figure 6.2 A 16-QAM constellation diagram.

Each point represents a series of bits and is based on a combination of signal phase and amplitude. When the receiving device demodulates the analog signal, it takes the signal's amplitude and phase and converts them to a series of bits and vice versa on the encoding process.

This is one such example of a Trellis Coding scheme. There are many others, including QPSK, 64-QAM, and 128-QAM.

Data Modulation

As soon as the data bits are encoded, they are modulated onto the wire into multiple bins or subcarriers, each capable of transporting a stream of at least 32kbps

of data. Each of these subcarriers is in its own frequency band, which is just over 4kHz in width. Up to 255 subcarriers are specified, although some are not used, giving a total downstream throughput of just over 8Mbps. One of these subcarriers is used for the pilot signal, reducing the throughput by 32kbps. The subcarriers can transport a variable number of bits (as long as the bins are a multiple of 32kbps); they usually are limited by the bit error rate. If it becomes too high (more than 1 in 10^7), the number of bits per symbol is reduced. Recall that a symbol is based on some constellation encoding, and the use of Trellis Code modulation with ADSL is advised.

Error Correction and Detection

To be tolerant of bit errors that might occur on the local loop, Reed-Solomon Forward Error Correction (FEC) is used to detect and correct such externalities. FEC works by sampling the original data, much like constellation coding used in TCM. Accompanying the original data are parity symbols, so if any part of the data stream becomes corrupted, it can be reconstructed using the additional parity data. This works much like any other parity-checking mechanism, such as a RAID-5 data storage array. If one drive is lost from the array, the data can be faithfully reconstructed using parity information that is spread across all drives. Of course, such data protection is not for free. An overhead, or tax, is incurred by sending redundant information, which reduces the real data rate. The amount of overhead varies according to implementation. More FEC data can be added to the original payload if the line is more prone to error. These FEC-encoded symbols are called code words, and they are fed into the next stage of the ADSL modem system, the modulator.

Reed-Solomon FEC codes are very good at correcting isolated errors but not very good at correcting a run of errors, such as those due to a transient burst of noise on the line. G.992.1 provides for interleaving, which helps spread out a code word in time such that channel burst errors are also spread through a block of data. This introduces extra latency to the ADSL line, because data must be queued in blocks, typically from 1 to 32ms, before both encoding. Interleaving is necessary to provide noise immunity for ADSL connections, particularly given the density of DSL lines now in service. Provision for a noninterleaved fast path is also provided, allowing providers to accept a trade-off between noise immunity and latency for latency-sensitive applications such as VoIP and gaming.

G.LITE

Defined in G.992.2, G.lite is a lightweight implementation of the ADSL specification from G.992.1, but it comes with some enhancements. G.lite has removed references to STM transport and defines only ATM data transport. An added feature is fast retrain, which allows an ADSL line to quickly recover from induced noise, such as a telephone on-hook/off-hook condition or other types of noise. It does have a significantly lower maximum downstream bandwidth—1.5Mbps. On the modulation level, the number of frequency subcarriers is limited to 127 in the downstream path and 31 in the upstream direction.

One of the advantages of G.lite is its lack of requirement for a central splitter at the customer premises, thereby eliminating a provider truck roll. At the time the standard was written, another advantage was a cheaper CPE. This was achieved in part by its reduced maximum frequency, which limited the amount of RF interference generated due to unbalanced customer premises wiring. In practice, most providers have skipped G.lite and run full-rate ADSL with individual filters per phone and a splitterless installation wherever possible.

ADSL2 AND ADSL2+

The next jump in the evolution of ADSL standards was ADSL2, defined in G.992.3. It allows for download speeds of just over 12Mbps for short loop lengths. This is achieved through improved line-encoding techniques, signal-processing algorithms, and various other low-level enhancements over the first-generation ADSL standard. On the electrical level, most of the characteristics are the same as ADSL. There is still a 4.3kHz subcarrier, and the operating range is also between 25kHz and 1.1MHz. ADSL2 still allows the DSL signal to coexist with either a POTS or ISDN system. The one exception is when the line is configured to be in an all-digital mode. This enables the ATU-R (CPE) equipment to use frequencies starting as low as 3kHz, which overlaps with the POTS frequency band. The benefit of being able to use these lower frequencies is that the maximum upstream rate is increased by 256kbps. If POTS is not required on the line, either because voice services are delivered via some other mechanism (VoIP) or because the channelized voice mode is used, 256kbps can be an ideal bandwidth increase, especially to try to deploy as close to symmetric services as possible (the same rate upstream as downstream). However, another technology is more suited to symmetric requirements—SHDSL, which is described later in this chapter.

ADSL2 FEATURES

After spending some time deployed in the field, the industry found several ways to improve asymmetric DSL services. These include a longer loop length and reducing the synchronization time (the time to train the two modems) to a few seconds. An equally desirable feature is the ability to dynamically adapt to changes in line conditions, called Dynamic Rate Adaptation (DRA).

The specification also explains how to support the simultaneous transmission of STM, ATM, and voice over ADSL. Voice over DSL enables concurrently transporting channelized voice services and cell-based ATM streams. Voice and data streams can be converged on a single platform if a provider wants to use the same equipment on the DSLAM to deliver both services rather than splitting them off to a separate POTS switch or soft switch.

Link bonding using the ATM Forum's inverse multiplexing (IMA) technique is also supported. IMA allows two or more copper links to be inversely multiplexed together to achieve higher data rates compared to just a single link. IMA enables cells to be distributed evenly across the links at the transmitting end, then reassembled at the other end in the correct order. Of course, two adjacent pairs both transmitting down the wire with a given power can cause interference with each other, so a conservative approach to spectral usage is employed on the lines. So as long as interference does not present a problem, it is possible to achieve a rate of the sum of the member links.

Longer Loop Length

ADSL2 extends the upper usable limit of the loop length to about 15,000 to 18,000 feet (4.5km to 5.5km), depending on the gauge of the local loop. This works by introducing variable framing overhead. ADSL requires a fixed framing overhead of 32kbps, whereas ADSL2 can have anywhere from 4kbps to 32kbps. This is useful in situations where the overhead-to-payload ratio is high, such as on 128kbps lines to squeeze a few extra kilobits per second from the line.

Faster Initialization and Dynamic Rate Adaptation

Users will be pleased with the faster initialization time for the line to train from a completely down state. Previously, the DSL signal light on an ADSL system would blink on the modem for about 10 seconds during the handshaking

process. This process is significantly faster on an ADSL2 line and reduces the training time to about 3 seconds.

Equally beneficial to customers and providers is the ability to dynamically modify the synchronization rate in real time if the line has too many errors. This works by monitoring the bit errors in a subcarrier. If the rate becomes too high, it reduces the number of bits-per-symbol in use in that subcarrier. The bits-per-symbol rate continues to decrease to 0, effectively disabling that particular DMT subcarrier if the error rate does not decrease. As a result, it lowers the line rate without needing to drop the entire DSL line. If the noise was transient, some DSLAMs can train back up to a higher rate. This is a change from G.992.1, in which the whole line had to retrain when the bit error rate went above a certain threshold. To communicate these real-time changes to the rest of the network, Access Node Control Protocol (ANCP) in rate-adaptive mode (RAM) is an ideal protocol to handle the communication of the line changes from the DSLAM to the BNG. It works by using an out-of-band channel between the DSLAM and BNG. Whenever a line changes its synchronization rate, the new details are communicated back to the BNG. ANCP is described in more detail in the section “ANCP and the Access Network” in Chapter 8, “Deploying Quality of Service.”

ADSL2+

The ITU then released the standard for ADSL2+; it significantly increases the downstream rate to about 24Mbps. It also increases the upper limit of the frequency range to 2.2MHz and increases the number of subcarriers to 512. All these extra subcarriers are allocated to the downstream transceivers. Because much of this specification builds on ADSL2 (G.992.3-4), the peak upstream rate is still roughly 1Mbps. An exception to this is Annex M, which increases the upstream peak rate to 3.5Mbps and increases the maximum upstream frequency from 138kHz to a potential 276kHz. Annex M is not too widespread due to the existing base of Annex A and B systems, which use 138kHz to 276kHz for the downstream path. NEXT is normally avoided in DSL by using a separate spectrum for transmitters and receivers; thus, there is potential for NEXT if the three types are mixed.

As described earlier, as both loop length and frequency increase, resistance also increases. This causes attenuation of the higher frequencies ahead of the others. As ADSL speeds increase, they become more dependent on a short loop length to keep the current loss low enough to support such high frequencies. For example, ADSL2+'s synchronization rate starts dropping from a peak of about 24Mbps at about 3,000 feet (900 meters). ADSL starts decaying at 6,000 feet (1.8km).

VDSL AND VDSL2

VDSL and VDSL2 dramatically increase transmission rates up to a potential 104Mbps. This is achieved similar to the step that ADSL2+ made from ADSL2—the frequency range on the copper is increased. The upper limit is now 12MHz. This section covers changes that VDSL will bring to the market and how service providers should think about designing their access network. For example, the change from ADSL2 to ADSL2+ meant that shorter loop lengths were needed to realize the benefit of higher line speeds—and even more so with VDSL and VDSL2. This also has implications for frequency spectrum planners, for both service providers and local regulators, because the increase in line frequency to 12MHz has the potential to cause much more interference than previous DSL standards.

One of the changes is that the ATU-C is now called the VDSL (or VDSL2) Transceiver Unit at the Optical Network Unit (VTU-O), and the ATU-R is now called the VDSL (or VDSL2) Transceiver Unit-Remote Site (VDSL-R).

VDSL PHYSICAL AND ELECTRICAL CHARACTERISTICS

As mentioned, the frequency band has been increased to 12MHz with the release of ITU VDSL recommendation G.993.1. To limit the degree of interference with ADSL systems, VDSL uses the frequency range in a different way than its predecessors. Recall that with ADSL2+ Annex M, the extra upstream bandwidth capacity was achieved by increasing the upstream frequency range to 276kHz (if needed). This meant that it overlapped with part of the downstream spectrum of non-Annex M systems, creating the possibility of unwanted interference between the two Annex systems. The problem is, an increase in upstream frequency is required to boost the speed. However, if the upstream starts encroaching in the

range up to 276kHz, this would cause the same compatibility issue that ADSL2+ Annex M has. The approach of VDSL and VDSL2 is to shift up the lowest-frequency band to start at 138kHz and make it the first *downstream* carrier. This homogenizes the local loop so that adjacent cables in the binder are more spectrally compatible. The next band up is used for the first upstream part, followed by two more bands, which are downstream and upstream, respectively. Figure 6.3 shows this more clearly, including how it interoperates with VDSL.




Figure 6.3 Frequency usage on a VDSL line that is shared with a POTS service.

To accommodate different regional requirements, these groups of frequencies can be split in three different ways using band plans. Even so, these band plans still follow the principle of four frequency bands, which are composed of two upstream and two downstream parts.

Ethernet Encapsulation

Expanding further on the advantage of using Ethernet in the first mile (EFM, described in 802.3AH), it reduces the overhead and complexity in both the CPE and the DSLAM. By eliminating a fixed 5-byte-per-cell overhead and CPCS padding, it optimizes bandwidth usage on the local access. At times, this overhead (also called cell tax) can be as much as 20%, depending on the packet size. Also, the complexity of the hardware at each end of the line is reduced, because an ATM SAR function is no longer needed if Ethernet encapsulation is required. However, in practice DSL chip vendors will continue to integrate both ATM and Ethernet capability in a single chip in both the CPE and DSLAMs for the foreseeable future. The active development and market shift are likely to precede a move to Ethernet encapsulation as VDSL becomes more popular.

Figure 6.4 shows an extreme (but possible) case of an ATM last-mile encapsulation producing more than 20% overhead.




Figure 6.4 IP packet requiring 20% overhead for transportation.

QoS features are equal to, if not better than, ATM for Ethernet-encapsulated local loops. With an ATM-encapsulated local loop, devices at each end perform prioritization between PVCs based on each PVC's service class—CBR, UBR, VBR-nrt, and so on. However, no capability in the AAL5 protocol allows prioritization within a class. For example, with a single PVC access architecture, prioritization needs to be done some other way. A single virtual circuit or VLAN on a local loop using Ethernet and 802.1Q can contain up to eight different priority markings in the frame header, which allows priority multiplexing at the data link layer within a single virtual circuit. Of course, if multiple VLANs are used on the local access, each VLAN can also be assigned a weighting relative to the others, which provides the same inter-PVC prioritization as the ATM model.

VDSL2

In addition to increased frequency spectrum usage, the biggest benefit of VDSL2 is the ability to encapsulate native Ethernet on the local loop, rather than remain with traditional ATM encapsulation; of course, ATM (and STM, for that matter) can still be used.

Standards and Spectrum

G.993.2 pushes the upper frequency limit to 30MHz and increases the number of upstream bands to three and the downstream to potentially four. In the most optimal case, it is possible to reach speeds of up to 100Mbps upstream and downstream (200Mbps total), with loop lengths of up to 1,600 feet (500 meters). Of course, the rate drops dramatically as the distance from the DSLAM increases, as with G.993.1 VDSL. However, due to the mandatory use of Trellis Coding, an optional extra frequency usage in the first upstream band, and some other line optimizations, VDSL can continue to provide reliable operation on loop lengths up to 8,200 feet (2,500 meters). The first upstream band is similar to ADSL2+, Annex M, which extends the upper bounds of the upstream channel to 276kHz. Of course, the potential for interference still applies and needs to be considered by spectrum planners in each region.

Band Plans

With the greater frequency range and increased power usage in the downstream direction (up to 20dBm) used for VDSL2, it's important to consider compatibility with existing DSL-based installations and local regulations. For example, in the UK, the regulator permits a maximum frequency of only 7.05MHz on the local loop, so a special set of frequencies was created that satisfied local spectral requirements. Such requirements are captured as part of *band plans* or *profiles* in the ITU specification, which record various details, including frequency ranges and DMT subcarrier spacing (one of the profiles changes the subcarrier range from 4.3kHz to 8kHz). Eight profiles are defined. Some address regional spectral requirements, and others address compatibility between two local ADSL system installations.

The specification describes two modes of access node (AN) system configurations: cabinet-based and CO-based. The section "Copper Network Reticulation Considerations" describes the changes that need to occur in the local access

network to support such high bit rates. One point that needs to be covered is that when VDSL2 systems are installed in a remote cabinet, if the system is configured to run in a certain band plan, existing ADSL loops may be adversely affected due to spectral incompatibility. Thus, several modes coexist better with ADSL. Therefore, it is important that this be incorporated into planning of the local loop.

COPPER NETWORK RETICULATION CONSIDERATIONS

When looking at graphs of theoretical VDSL access speeds versus loop length, it may seem that the high bandwidths needed to support Standard-Definition Television (SD-TV) and High-Definition Television (HD-TV) services can be extended along significant copper distances. In real deployments, the speed drops significantly after 4,000 feet (1,200 meters). Because POTS can reach distances of 18,000 feet (5.5km) or longer with bridge taps, carriers have exploited this distance capability. This means that most existing copper loops are a lot longer than the shorter lengths needed for high bit rate services. Therefore, VDSL deployments need to incorporate the deployment of fiber-fed DSLAMs in remote cabinets that lie close to customer premises. Delivering fiber to the DSLAM is called fiber to the node (FTTN). These nodes can be in either new or existing roadside cabinets. In some scenarios, however, the heat dissipation requirements, environmental noise-level bylaws, and space available in existing cabinets can restrict the deployment of some hardware types. This can mean relegating the VTU-O to a larger central office. The danger is that this reduces the coverage footprint of high bit-rate VDSL services. It also increases the copper loop length and the density of pairs in the cable binder, because the CO must serve many more subscribers compared to a remote node. Thus, an increased amount of crosstalk occurs, and hence lower achievable speeds in the local-access network.

Multi-Unit Dwellings (MUDs) are ideal for this kind of distributed deployment. Such buildings could be apartment blocks, where an access provider installs DSLAM subracks in the basement for the rest of the building. Hotels can install such subracks to deliver high-speed access to guest rooms. They also can be useful in older buildings where it is infeasible to rewire rooms and conduits with CAT-5 or CAT-6 cabling for Ethernet services. Of course, the per-room installation cost then goes up, because modems need to be installed in each room.

In any case, the new VDSL services delivered over cabinet-fed DSL lines (rather than CO-fed) need to be spectrally compatible with existing DSL deployments. One example of when problems arise is when a short VDSL loop is deployed adjacent to a long ADSL loop; the power differences on each line are bound to cause interference in the form of NEXT. Power Spectral Density (PSD) shaping helps by reducing the power level of certain frequency subcarriers and potentially reallocating power elsewhere in the spectrum. Also, carriers can elect to set power levels lower than the upper limit of the PSD mask. This is useful in specific deployment scenarios where there is the potential for spectral interference at a given power level. Typically DSL modems of all types lower output power to the minimum required to maintain a given performance. As mentioned earlier in the chapter, having multiple band plans in the VDSL specification allows service providers to choose the best profile for their environment.

SHDSL

Single-pair High-Speed DSL (SHDSL) can be both a symmetric and asymmetric DSL technology and is ideally suited to business applications. All the DSL services, with the exception of VDSL2, are designed to be asymmetric. The applications originally envisioned for ADSL—Video on Demand (VoD) and Internet services—need more bandwidth in the downstream direction toward the customer. In the residential market, applications such as YouTube, Google video, Slingbox, and other subscriber-produced content are changing this model. Nevertheless, the traffic ratio is still tipped in favor of the downstream direction.

For providers that have deployed only ADSL services, the maximum upstream rate is about 800kbps. This is unsuitable for customers who want to deploy IP private branch exchanges (PBXs) with more than seven or eight lines or host a small to medium-sized data center or any other data services that need more than 1000kbps of upstream capacity. This bit rate is also the upper limit of ADSL. Just as 56kbps is the upper limit of analog modems running over the PSTN, both technologies rarely achieve their advertised maximums except in perfect lab conditions.

SHDSL addresses many of these limitations by providing a more business-suited service by allowing the line to operate with symmetric bit rates. Other features include support for a multiwire configuration, either to extend the service's reach

or to increase the total bandwidth available on shorter loops. Signal repeaters can be used on the line to increase the signal distance if there are not extra cables for a multipair system that is being used for reach extension. Remote powering by the CO of the modem at the customer premises is supported. There are also several supported payload types—Transport Protocol-Specific Transmission Convergence (TPS-TC)—such as channelized and unchannelized T1 and E1 circuits, ATM, and packet mode. Finally, to improve resilience against line clocking slips, bit stuffing can be performed on the path.

However, to be a completely business-grade service, the SLA, in terms of Time-To-Restore (TTR) and service uptime, must be equivalent to other leased-line services, such as ATM and Frame Relay.

STANDARDS AND SPECTRUM

In North America, 1.544Mbps was the most common circuit of choice in the pleiochronous (PDH) transmission hierarchy. For a long time, it was delivered using Alternate Mark Inversion (AMI), which created a well-established base of equipment that required repeaters about every mile (1.5km). The successor, High-Speed Digital Subscriber Line (HDSL), was originally developed in the U.S. by Bellcore. It used 2B1Q line coding, which is much more spectrally efficient. Like AMI-provisioned lines, HDSL lines required two pairs to reach the full 1.544Mbps line speed. ANSI adopted it as an American standard. HDSL2 and SDSL soon followed. HDSL2 is more forgiving of poor line conditions and enables a T1-speed line to be provisioned over a single cable pair. To get rates slower than T1 speeds, the line can be channelized into multiple 64kbps time slots. It also makes line coding more efficient by adopting TC-PAM. In the North American market, SDSL has been deployed using proprietary 2B1Q line coding techniques to achieve variable bandwidth rates up to 1.544Mbps over a single pair. This technology is giving way to the international standard adopted by the ITU, SHDSL, defined in G.991.2. This is because equipment can be made with a single standard in mind (apart from annex differences), and economies of scale bring down the price of SHDSL chipsets.

A similar process of evolution occurred in Europe. E1 lines were first provisioned using the inefficient HDB3 line coding method. After HDSL was developed in

the U.S., it was adapted to the European environment and was adopted as a standard by the European Telecommunications Standards Institute (ETSI). The main modification was to support the full 2048kbps on a dual-pair system compared with the 1544kbps system across the Atlantic. SDSL technology did not catch on in Europe, except by name. SHDSL, standardized by the ITU, was adopted by ETSI and renamed SDSL. Luckily, the two are compatible.

PHYSICAL AND ELECTRICAL CHARACTERISTICS

SHDSL, also known as G.SHDSL, is defined in ITU technical recommendation G.991.2. The specification went through two iterations. The first version defined support for a two-wire, single-pair service and an optional four-wire mode. The second version added an M -pair mode, which allows up to four pairs to be used concurrently. A cable pair is also called a span in G.SHDSL terminology. The use of M pairs increases the usable bandwidth in a linear fashion, according to the number of pairs (that is, $n \text{ kbps} * M$). Depending on the Annex used, n could be anywhere between 192kbps and 5696kbps. The different data rates are described in the section “Transport Capacity.” In the M -pair approach, the data are spread among each of the pairs just before modulation onto the wire. Each SHDSL frame is divided into four blocks, each of which is subdivided into 12 subblocks. Parts of each subblock are distributed among the spans and are reassembled at the other end. To compensate for differences in wire diameter, length, use of bridge taps, and any other impairments, each SHDSL Termination Unit (STU, similar to a VTU or ATU) implements a delay buffer of just over 50 microseconds to absorb propagation delay differences between the spans. This data interleaving is different from ADSL-based systems, which use ATM IMA to inversely multiplex multiple pairs to get increased data throughput. This involves spraying cells across the constituent links to evenly distribute data.

SHDSL also includes a transport-convergence-specific (TC-specific) sublayer for packet transfer, similar to how VDSL includes a TC specification for packet transfer. Because SHDSL is very liberal with its use of the radio frequency spectrum, it does not support sharing the same cable pair with either POTS or ISDN. The intention is that such services are handled by higher-layer protocols. However, SHDSL does support synchronous ISDN through the use of an ISDN-specific TC application sublayer. This is in comparison to VDSL and ADSL, which allow ISDN to share the same cable pair by using FDD.

SHDSL streams occupy the spectrum between 0 and 400kHz and place various limits on the maximum power usage at several ranges of the frequency spectrum and use PSD shaping. Of course, local spectrum usage dictates the ultimate frequency configuration.

One useful feature in the specification is the usage of signal regenerators. Up to eight signal regenerators can be used on a span, increasing the reach if the signal-to-noise ratio (SNR) becomes too great for non-regenerated line segments. To faithfully regenerate any such signal, such a piece of equipment needs to be active so that it can decode the SHDSL signal, reclock it, and provide an SHDSL repeater unit (SRU) to both sides of the regenerator. Such a regenerator can be remotely placed in diagnostic mode if required, by setting the necessary bit in the EOC channel. One important point is that SHDSL uses more optimized line coding techniques, so it uses fewer spectra than a T1/E1 line, which usually uses the spectrally inefficient AMI or High-Density Bipolar level 3 (HDB3) coding respectively. As a result, the loop length before needing a repeater is much longer with SHDSL. For example, a regenerator is not needed for loops shorter than 12,000 feet (3.6km), compared to every 5,000 feet (1.5km) on a typical AMI-encoded T1 line.

Modulation

The physical DSL level is called the Medium-Specific TC (PMS-TC) layer. Before an SHDSL data stream is encoded onto the wire, it is placed in a PMS-TC frame. Each frame contains 4 bits that can be used as stuffing bits to guard against frame alignment problems during a burst of noise on the line.

The ITU standard uses a 16-TC-PAM scheme, which makes efficient use of the spectrum by representing several bits in a single symbol. Other symmetric systems, such as HDSL, SDSL, and ISDN, use 2B1Q line coding. 2B1Q is less spectrally efficient because it can encode only two bit-patterns per symbol, compared to 16-TC-PAM, which encodes 4 bits per symbol.

The standard also allows remote equipment (STU-R or SRU) to be supplied with a 200V, 15W power source from the CO using the same cable pair as that used for data.

TRANSPORT CAPACITY

When SHDSL lines are operating in symmetric mode on a single pair, the maximum transport capacity can be anywhere between 192kbps and 5696kbps, depending on the mode and annex used. For example, Annex A is intended for use in North America and has a minimum rate of 192kbps and a maximum rate of 2304kbps, with steps of 64kbps in between. All implementations supporting this and the other annexes have limits placed on various radio frequencies through the use of PSD shaping. Annex B typically is used in European networks and specifies different frequency and power requirements than Annex A. Annex B also has bit rates up to 2304kbps. If the line is set in asymmetric mode, increments can be in 8kbps steps. Annex F extends the maximum data rate to 5996kbps.

SUMMARY

This chapter has discussed last-mile access using DSL as the primary technology. Although several other types of access technologies exist, including wireless (covered in Chapter 9), this book has focused on DSL. Other techniques, such as Ethernet over Power, fiber to the home/curb/house, HFC, are also suitable candidates to use instead of, or in addition to, DSL. In the interest of space, this chapter has focused solely on DSL-related access. Table 6.1 summarizes the ITU standards mentioned in this chapter, their maximum bit rates, and the frequency ranges used by each.

Table 6.1 ITU DSL Recommendations and Their Speed Capabilities

Recommendation Number (Informal Name in Parentheses)	Recommendation Title	Approximate Maximum Speed Capabilities	Passband Frequencies
G.991.2 (G.SHDSL)	Single-pair High-speed Digital Subscriber Line (SHDSL) transceivers	2304kbps upstream 2304kbps downstream	0kHz to 400kHz upstream 0kHz to 400kHz downstream
G.992.1 (G.DMT)	Asymmetric Digital Subscriber Line (ADSL) transceivers	1Mbps upstream 8Mbps downstream	25kHz to 138kHz upstream 138kHz to 1104kHz downstream

continues

Table 6.1 continued

Recommendation Number (Informal Name in Parentheses)	Recommendation Title	Approximate Maximum Speed Capabilities	Passband Frequencies
G.992.2 (G.lite)	Splitterless Asymmetric Digital Subscriber Line (ADSL) transceivers	1Mbps upstream 1.5Mbps downstream	25 to 138kHz upstream Annex A 138 to 552kHz downstream or Annex B 25 to 552kHz downstream Note: Annex B uses overlapped upstream and downstream spectrum.
G.992.3	Asymmetric Digital Subscriber Line transceivers 2 (ADSL2)	1Mbps upstream 12Mbps downstream	Annex A ~26kHz to 138kHz upstream Annex B 120kHz to 276kHz upstream Annex A ~26kHz to 1104kHz downstream (overlapped upstream and downstream spectrum) or ~138kHz to 1104kHz downstream (nonoverlapped upstream and downstream spectrum) Annex B 120kHz to 1104kHz downstream (overlapped upstream and downstream spectrum) or 254kHz to 1104kHz downstream (nonoverlapped upstream and downstream spectrum)

Recommendation Number (Informal Name in Parentheses)	Recommendation Title	Approximate Maximum Speed Capabilities	Passband Frequencies
G.992.4	Splitterless Asymmetric Digital Subscriber Line transceivers 2 (ADSL2)	1Mbps upstream	~25kHz to 138kHz upstream Annex A 138kHz to 552kHz downstream (nonoverlapped upstream and downstream spectrum) or ~25kHz to 552kHz downstream (overlapped upstream and downstream spectrum)
G.992.5	Asymmetric Digital Subscriber Line transceivers Extended Bandwidth ADSL2 (ADSL2+)	1Mbps upstream 24Mbps downstream	Annex A ~26kHz to 138kHz upstream Annex B 120kHz to 276kHz upstream Annex A ~26kHz to 2208kHz downstream (overlapped upstream and downstream spectrum) or ~138kHz to 2208kHz downstream (nonoverlapped upstream and downstream spectrum) Annex B 120kHz to 2208kHz downstream (overlapped upstream and downstream spectrum) or 254kHz to 2208kHz downstream (nonoverlapped upstream and downstream spectrum)
G.993.1	Very high-speed Digital Subscriber Line (VDSL)	13Mbps upstream 22Mbps downstream	Various band plans from ~26kHz to 12MHz. Implementation depends on region.
G.993.2	Very high-speed Digital Subscriber Line 2 (VDSL2)	Up to 100Mbps upstream Up to 100Mbps downstream	Various band plans from ~26kHz to 30MHz

Table 6.2 summarizes the technologies and acronyms used in this chapter.

Table 6.2 Terms Used in This Chapter

Protocol or Term	Acronym	Definition
Access node	AN	A term adopted by the DSL Forum to generalize a device that implements an ATU-C function and optionally a switching function from other ANs.
ADSL Transceiver Unit, CO	ATU-C	The device at the CO or remote node terminating the ADSL signal from the ATU-R.
ADSL Transceiver Unit at the Remote End	ATU-R	The device at the customer that terminates the ADSL signal from the ATU-C.
Annex A, B, C, M		Examples of sections added to specifications to capture location-specific requirements. For example, Annex A in an ADSL context refers to an ADSL service delivered on the same cable as POTS.
Asymmetric Digital Subscriber Line	ADSL	A technology standardized by various bodies, including the ITU, which permits simultaneous high-speed data rates and POTS usage on an ordinary telephone line.
Bit swapping		A mandatory technique used in the low-level transmission layer of DSL transceivers. Changes data and power rates of individual subcarriers to compensate for induced noise but aims to keep the overall line rate the same.
Bridge tap		An open-circuit, unused twisted-pair section connected at some point along a subscriber line. A legacy from the days when copper pairs to households were often oversubscribed.
Carrierless Amplitude-Phase modulation	CAP	A single-carrier modulation technique used in the DSL market that uses adaptive equalization to compensate for attenuation and phase errors.
Digital Subscriber Line	DSL	A way to transmit data over one or more copper pairs. The two modulation standards are CAP and DMT. Current speeds are up to 100Mbps.
Digital Subscriber Line Access Multiplexer	DSLAM	A node that can be at either a telco's CO or a remote node in a roadside cabinet. It terminates many customer DSL lines by implementing multiple ATU-C or VTU-O functions.
Discrete Multitone	DMT	A technique used to subdivide the spectrum on a DSL line into many autonomous transceivers. Transceivers operate in parallel to deliver a fast data stream to higher layers.
Frequency Division Duplexing	FDD	When two frequency ranges are duplexed onto a single medium by a radio application. ADSL is an FDD system because the upstream and downstream portions are in separate frequency ranges and normally do not interfere with one another.

Protocol or Term	Acronym	Definition
Frequency Division Multiplexing	FDM	When more than two frequencies are used by a radio application on a single medium. VDSL uses FDM because it has several frequency ranges, for upstream, downstream, and POTS.
Layer 2 Control Protocol (now Access Node Control Protocol)	L2CP (now ANCP)	A way of controlling an AN remotely for operations and administration management. Also can be used by a BNG or other device to retrieve statistics from the AN.
Multi-Unit Dwelling	MUD	A building such as an apartment block where many dwellings share the same building.
Online reconfiguration	OLR	Allows subtle changes at the subcarrier level (for example, bit swapping) or more drastic changes such as changing the overall line rate. All OLR changes can be done without resynchronizing the line.
Packet Transfer Mode, Transmission Convergence	PTM-TC	The ability of a DSL line to support native packet transfer rather than through an intermediate protocol, such as ATM. Defined in G991.2 and the VDSL standards.
Plain Old Telephone Service	POTS	An analog telephone service operating in the voice band between 300 and 3400Hz. For most of its existence, it has been delivered using analog last-mile loops.
Power Spectral Density shaping	PSD	Controls the power used across different ranges of the DSL frequency spectrum.
Public Switched Telephone Network	PSTN	A general term for the telephone network supporting voice call switching. Voice calls include modulated data signals coming from analog modems using V.90 or V.24 modulation.
Reed-Solomon coding		A FEC scheme common in DSL systems. Check bytes are added to transmitted data to allow correction of data corrupted on the line. Reed-Solomon coding is a cyclic block-coding scheme first described in the early 1960s.
Seamless Rate Adaptation	SRA	A way to modify the overall DSL line rate without resynchronizing the modems at each end. A feature of the overall OLR capability.
Subcarrier		A fundamental building block of the DMT-based DSL system. Each ADSL subcarrier has a frequency range of ~4kHz and supports an ADSL transceiver.
VDSL Transceiver Unit-Central Office	VTU-O	A VDSL access node that provides VDSL service to a VTU-R.
VDSL Transceiver Unit-Remote End	VTU-R	A VDSL access node that terminates the transmit signal from the VTU-O and transmits a VTU-R signal to the transceiver at the other end. Often called CPE for simplicity.
Very high-speed DSL, Very high-speed DSL 2	VDSL	Standards defined in G.993.1 and G.993.2, respectively. Throughput rates are up to 22Mbps for the first standard and 100Mbps in one Annex of the second standard.

Cable access operators such as Time Warner have offered triple-play services for some time. Video and Internet services run over the coaxial cable between the head-end and the cable modem. Voice services may run using a reserved voice range in the coaxial, a separate twisted pair, or as packetized voice using a SIP or H.323 client connected to the modem. One of the big advantages of DSL over other broadband systems is the pre-existing access network. Copper loops run to any location where there is a POTS handset.

But working against DSL's ubiquity is fiber-to-the-premises rollouts. Even with VDSL2 potentially reaching rates of 100Mbps, fiber has the advantage of having a future-proof medium capacity. Upgrading its capacity is relatively easy by changing transponders at each end. Inherent in the solution are services being asymmetric in bandwidth. One advantage that it does not offer yet is low-cost installation; this is the major inhibiting factor of more aggressive fiber access rollouts. DSL has the advantage of being able to reuse infrastructure that has been over a century in the making—the local loop. By running ever higher bit rate services over the existing copper infrastructure, cheaper DSL services shift the need for fiber deployments further into the future. In fact, the two can coexist to a large degree. Fiber-to-the-node is a common configuration when running VDSL services to remote cabinets or DSLAMs to the basements of MUDs. Fiber has the necessary capacity to support the large bandwidths needed for triple-play services, and DSL can provide sufficient bandwidth and sufficiently low cost in the last segment to the household—for now.